

Traceable Secret Sharing: Strong Security and Efficient Constructions

Dan Boneh, Aditi Partap, and Lior Rotem

Stanford University

{dabo,aditi712,lrotem}@cs.stanford.edu

Abstract. Suppose Alice uses a t -out-of- n secret sharing to store her secret key on n servers. Her secret key is protected as long as t of them do not collude. However, what if a less-than- t subset of the servers decides to offer the shares they have for sale? In this case, Alice should be able to hold them accountable, or else nothing prevents them from selling her shares. With this motivation in mind, Goyal, Song, and Srinivasan (CRYPTO 21) introduced the concept of *traceable secret sharing*. In such schemes, it is possible to provably trace the leaked secret shares back to the servers who leaked them. Goyal et al. presented the first construction of a traceable secret sharing scheme. However, secret shares in their construction are quadratic in the secret size, and their tracing algorithm is quite involved as it relies on Goldreich-Levin decoding.

In this work, we put forth new definitions and practical constructions for traceable secret sharing. In our model, some $f < t$ servers output a reconstruction box R that may arbitrarily depend on their shares. Given additional $t - f$ shares, R reconstructs and outputs the secret. The task is to trace R back to the corrupted servers given black-box access to R . Unlike Goyal et al., we do not assume that the tracing algorithm has any information on how the corrupted servers constructed R from the shares in their possession.

We then present two very efficient constructions of traceable secret sharing based on two classic secret sharing schemes. In both of our schemes, shares are only twice as large as the secret, improving over the quadratic overhead of Goyal et al. Our first scheme is obtained by presenting a new practical tracing algorithm for the widely-used Shamir secret sharing scheme. Our second construction is based on an extension of Blakley’s secret sharing scheme. Tracing in this scheme is optimally efficient, and requires just one successful query to R . We believe that our constructions are an important step towards bringing traceable secret-sharing schemes to practice. This work also raises several interesting open problems that we describe in the paper.

1 Introduction

Secret sharing [64, 6] is one of the most foundational concepts in cryptography. Of particular use are threshold secret-sharing schemes in which a secret s is split into n shares sh_1, \dots, sh_n so that any t shares suffice to reconstruct the secret, but any $t - 1$ of them reveal nothing about it. A fundamental use case of threshold secret sharing is storing confidential data on untrusted servers. Concretely, say that Alice stores the key to her crypto wallet (or any other piece of sensitive information) on n servers or cloud storage providers using a t -out-of- n secret sharing scheme.

Now suppose that a subset of f servers collude and sell the key shares in their possession to whomever is willing to meet their price. In this worrisome scenario, we would like some mechanism to enable Alice to hold these servers accountable. Otherwise, they would have a risk-free incentive to sell Alice’s secret information with no way to hold them accountable for their action.

Traceable secret sharing. This question was recently considered by Goyal, Song, and Srinivasan [42], who introduced the concept of *traceable secret sharing*. Roughly speaking, these are secret sharing schemes where leaked secret shares can be traced back to the servers who leaked them. We will make this requirement precise shortly.

In some secret sharing schemes, well-formed shares can be easily linked to the server to which they were issued. For example, if Shamir secret sharing is implemented such that server i gets the secret share $(i, f(i))$ for some polynomial f , then this share, at least as is, trivially identifies the i th server. However, the corrupted servers could potentially pull their shares together and obfuscate them in a manner that prevents their trivial identification. Hence, what we want is a secret-sharing scheme with the following two guarantees. First, if Alice obtains the leaked information, she should be able to trace it back to the corrupted servers. Moreover, to hold them accountable she should be able to produce a proof that these servers indeed leaked her secret information. This property is called *traceability*. The scenario to consider here is this: Alice, or some law enforcement agency, obtains the leaked information (for example, by buying it from the corrupted servers who offer it for sale) and traces it back to the corrupted servers. They produce a proof implicating the corrupted servers, which Alice may use in court. In this sense, traceability does not aim to, and cannot, prevent the leaking of secret shares with absolute certainty. Rather, it aims to serve as a deterrent, discouraging parties from leaking secret shares. This is similar to the logic underpinning traitor tracing schemes [22].

This leads us to the second requirement of traceable secret sharing, called *non-imputability*. This property stipulates that Alice should not be able to generate a false proof, implicating an honest server who was not involved in any leak. In our example, this prevents Alice from collecting unlawful damages in court.

In their work, Goyal et al. [42] constructed a new secret-sharing scheme, that satisfies both traceability and non-imputability. In their scheme, however, traceability comes at the cost of efficiency. First, the shares in their schemes are very large; they are λ^2 -bit long where λ is the security parameter and secrets lie in $\{0, 1\}^\lambda$. That is, shares are quadratic in the secret size. This should be contrasted with other secret sharing schemes (e.g., Shamir) in which shares are the same size as the secret. Second, their tracing algorithm is quite involved and relies on Goldreich-Levin decoding [35]. These sources of inefficiency leave open the problem of constructing a *practical* traceable secret-sharing scheme, in which shares are small and tracing is efficient. Additionally, the scheme of Goyal et al. is a custom scheme, and is quite different from standard secret-sharing schemes, such as Shamir’s [64] or Blakley’s [6] schemes. It is thus an interesting question to trace misbehaving servers in widely-used secret sharing schemes, such as Shamir secret sharing.

1.1 Our results

In this work, we present very efficient constructions of traceable secret-sharing schemes based on the classic schemes of Shamir [64] and Blakley [6]. We believe that our constructions are an important step towards bringing traceable secret-sharing schemes to practice. In more detail, our contributions are threefold:

1. We give new security definitions for traceable secret sharing, strengthening the definitions of Goyal et al. [42], with one caveat on which we elaborate below.
2. We present a new and practical tracing procedure for standard Shamir secret sharing.
3. We put forth a variant of Blakley secret sharing, and present an optimally-efficient tracing procedure for it.

We now elaborate on each of these contributions.

A stronger security notion. We revisit the definition of traceable secret sharing put forth by Goyal, Song, and Srinivasan [42]. In their model, a subset $\mathcal{I} \subseteq [n]$ of servers may leak $f_{\mathcal{I}}(\mathbf{sh}_{\mathcal{I}})$

where $f_{\mathcal{I}}$ is some adversarially chosen function and $\mathbf{sh}_{\mathcal{I}}$ is the set of all shares held by servers in \mathcal{I} . There is also a reconstruction box R that takes in the leaked information $f_{\mathcal{I}_1}(\mathbf{sh}_{\mathcal{I}_1}), \dots, f_{\mathcal{I}_k}(\mathbf{sh}_{\mathcal{I}_k})$ from several disjoint subsets and, if $|\bigcup_{i \in [k]} \mathcal{I}_i| \geq t$, outputs the underlying secret s . For tracing, they require the existence of a tracing algorithm Trace that solves the following problem: given $f_{\mathcal{I}_1}(\mathbf{sh}_{\mathcal{I}_1}), \dots, f_{\mathcal{I}_k}(\mathbf{sh}_{\mathcal{I}_k})$, a description of the functions $f_{\mathcal{I}_1}, \dots, f_{\mathcal{I}_k}$, and access to R , the tracing algorithm should output the identity of at least one corrupted server along with a proof π that this server is indeed corrupted.

In Section 2 we present a definition which requires tracing with less information, making it more realistic in our view. Concretely, in our model, a subset \mathcal{I} of $f < t$ colluding servers outputs a reconstruction box R that depends on the shares held by the servers in \mathcal{I} . This reconstruction box takes in as input additional $t - f$ shares $\mathbf{sh}'_1, \dots, \mathbf{sh}'_{t-f}$ of servers not in \mathcal{I} , and outputs the secret that is reconstructed from the shares provided as input and the shares of the servers in \mathcal{I} . Then, the definition requires that there is a tracing algorithm, that given *black-box access* to R , traces it back to \mathcal{I} . Indeed, giving the tracing algorithm the code of R seems unlikely to help, since the corrupted servers might apply some heuristic code cloaking to R before publishing it. In any case, our tracing algorithm will only require black-box access to R . Also observe, that if $|\mathcal{I}| \geq t$ then tracing R is impossible, since the servers could reconstruct the secret s^* and sell it directly. This perfectly hides their identity. This is why our focus is on the case where $f < t$.¹

We believe that our definition has a couple of advantages:

- First, the corrupted servers may try to thwart detection by combining their shares and applying some unknown post-processing to them. Modeling this leak as a reconstruction box R captures *any* post-processing that maintains the functionality of the shares. Importantly, our definition does not require that the tracing algorithm learns this post-processing function in order to trace. This is in line with the realistic scenario described above, in which a subset \mathcal{I} of corrupted servers decides to sell their shares in the form of the reconstruction box R . In the case Alice or a law enforcement agency obtain R to trace it back to \mathcal{I} , it seems unlikely that they will learn the exact manner by which R was constructed from the secret shares of \mathcal{I} .
- Second, our definition is stronger than that of Goyal et al. in the following sense: given the information required to trace in their model, one can trace in our model as well. One caveat to that is that our definition requires that R output the complete reconstructed secret, whereas Goyal et al. only require that R learns any non-trivial information about the secret. A thorough comparison of the two definitions is presented in Section 2, where we argue that tracing a complete reconstruction box is a natural model.

In Appendix A we present a simple tracing algorithm for any threshold secret sharing scheme. However, this algorithm suffers from three main limitations: first, the tracer learns the secret from the tracing key, which violates the principle of secret sharing; second, it only works for perfect reconstruction boxes that always output the correct secret; and third, it does not seem to allow for non-imputability. These shortcomings lead us to our Shamir-based and Blakley-based constructions. Looking ahead, our Shamir-based scheme will perfectly hide the secret from the tracer, whereas our Blakley-based scheme will satisfy some weaker notion of secret hiding given the tracing key (which still, unlike the simple tracing algorithm from Appendix A, does not allow the

¹ In the model of Goyal et al. this restriction is captured by requiring that all colluding subsets $\mathcal{I}_1, \dots, \mathcal{I}_k$ are of size less than t .

tracer to reconstruct the secret; see Section 2 for details). Both our schemes can trace imperfect reconstruction boxes and satisfy non-imputability.

Tracing leaks in Shamir secret sharing. In Section 3 we present an efficient tracing procedure for tracing leaks in standard Shamir t -out-of- n secret sharing, where each server's share is the evaluation of a degree $t - 1$ polynomial on a point x that is drawn *at random* from a field \mathbb{F} satisfying $|\mathbb{F}| \gg n$. We explain in Section 3 why choosing the x points in this manner is necessary. In particular, shares in our schemes consist of just two field elements – just twice as long as standard Shamir shares.

The basic idea behind our tracing algorithm is the following. Suppose for simplicity that R is a perfect reconstruction box that has $f = t - 1$ shares $\{(x_i, y_i)\}_{i \in [t-1]}$ hardcoded in it, and it takes one additional share (x^*, y^*) as input. By perfect, we mean that R always outputs the result of Lagrange interpolation at the point 0, applied to $\{(x_i, y_i)\}_{i \in [t-1]}$ and (x^*, y^*) . We observe that the Lagrange coefficient of y^* in this interpolation can be treated as an inverse polynomial in x^* whose roots are exactly the x -values x_1, \dots, x_{t-1} of the shares hardcoded in R . Then by carefully querying R a total of $\Omega(t)$ times, we are able to obtain t evaluations of this polynomial. Hence, we can interpolate it and then find its roots, giving us the x values of the leaked shares used by R . Since all x values are chosen at random from a large field, this information identifies the corrupted servers with overwhelming probability. More generally, to trace $f < t$ corruptions, our tracing algorithm issues $\Omega(f)$ queries to the reconstruction box R .

When R is not perfect, and may err with some probability ϵ , we need to use list decoding for Reed-Solomon codes (for example, the Guruswami-Sudan algorithm [43]) to obtain a list of candidate polynomials, such that one of them is guaranteed to have x_1, \dots, x_{t-1} as its roots. For more details, see Section 3, where we also explain how to trace when the number f of corruptions is not known in advance, and how to make our scheme accommodate non-imputability.

Tracing leaks in Blakley secret sharing. In Section 4 we present a very efficient tracing algorithm for t -out-of- n Blakley secret sharing. We let the secret s live in some finite field \mathbb{F} . To share s , the dealer first randomly extends it to a point $\mathbf{x} = (s, x_2, \dots, x_t)$ in \mathbb{F}^t , by appending $t - 1$ random \mathbb{F} -elements to it. Then, the share of server i is a random hyperplane H_i in \mathbb{F}^t that passes through \mathbf{x} . The intersection of any t such hyperplanes uniquely defines \mathbf{x} , and hence s , with overwhelming probability, whereas every $t - 1$ hyperplanes are (almost) statistically independent of s . Suppose we have a perfect reconstruction box R that has f hyperplanes H_{i_1}, \dots, H_{i_f} hardcoded in it which it uses to reconstruct the secret. Moreover, suppose that instead of just the secret s , this box outputs the entire point \mathbf{x} that is the intersection of H_{i_1}, \dots, H_{i_f} and the $t - f$ hyperplanes that it receives as input. In this case, tracing becomes easy. To trace R back to servers i_1, \dots, i_f , we simply invoke it on $t - f$ random hyperplanes H'_1, \dots, H'_{t-f} . It is not hard to see that the intersection point that R outputs is a random point \mathbf{r} in \mathbb{F}^t , conditioned on it intersecting H_{i_1}, \dots, H_{i_f} . If \mathbb{F} is large, then \mathbf{r} will almost surely not lie on any of the honest servers' hyperplanes, and will therefore identify the corrupted servers. Hence, with a single query to R we can identify the corrupted servers.

The problem is that R does not output the entire intersection point \mathbf{r} , but just its first coordinate, which is the reconstructed secret. This is not enough information for the above tracing procedure, as it does not uniquely identify the corrupted servers. To remedy this situation, in Section 4.1 we extend Blakley's scheme such that the secret is a full point in \mathbb{F}^t . Doing so while still retaining the secrecy of the secret sharing scheme turns out to be non-trivial. Then, in Section 4.2, we present a tracing algorithm that relies on the above idea for our extended scheme. This tracing

algorithm still requires just one query to the reconstruction box R , and is hence optimal in this sense.²

In Section 5 we also discuss how the fact that we use just one query to R may be beneficial for tracing *stateful* reconstruction services (as opposed to just *stateless* reconstruction boxes). Both the scheme of Goyal et al. [42] and our Shamir-based schemes cannot be used to trace stateful reconstruction services.

We emphasize that our Blakley-inspired traceable secret sharing scheme is also very efficient in terms of its share size. In Section 4 we discuss how to represent shares in this scheme with just $\lambda + |s|$, where λ is the security parameter and $|s|$ is the secret size. This essentially matches our Shamir-based construction for secrets in $\{0, 1\}^\lambda$. The reader is referred to Section 4 for our full-fledged scheme that satisfies non-impurity, and for additional technical details.

Future work. Our work opens up many avenues for future work, which we discuss at length in Section 5. One such interesting direction is to explore the connection between traceable secret sharing and erasure codes. Both Shamir’s and Blakley’s schemes can be seen as first randomly embedding the secret in a higher dimension space, and then deterministically encoding it using some erasure code [50]. A fascinating question that our work leaves open is to come up with tracing procedures for other secret-sharing schemes that can be thought of in the same manner. These include schemes that are based on CRT-encodings [3, 51, 36], and schemes that are obtained from low-density parity-check (LDPC) codes [23, 2].

Another compelling direction for future research is to come up with a traceable secret-sharing scheme in which reconstruction is linear and tracing requires only degree-2 operations. As we explain in Section 5, this may have exciting applications for tracing traitors in threshold VUFs.

Finally, Goyal et al. [42] showed how the techniques underlying their traceable secret sharing construction can be applied to obtain traceability for leaked information in schemes for computation delegation. However, the construction is not black-box from any traceable secret sharing scheme, and it would be interesting to apply our techniques to traceable computation delegation as well.

1.2 Additional Related Work

There are several cryptographic primitives that resemble the notion of traceable secret sharing. We briefly discuss them here. The first notion that comes to mind in that respect is that of traitor tracing for encryption schemes, first introduced by Chor, Fiat, and Naor [22]. In such schemes, a central authority issues a single encryption key, and n decryption keys, one per receiver. If any coalition of receivers comes together and produces a pirate decoder box D that can decrypt well-formed ciphertexts, then this box can be traced back to at least one of the receivers who contributed to it. Traitor tracing, therefore, also deals with tracing leaked secrets back to the parties who leaked them. The techniques used to construct traitor tracing schemes are, however, fundamentally different the techniques we use here. Many traitor tracing schemes have been proposed over the years (see, for example, [48, 54, 7, 30, 63, 45, 52, 28, 17, 33, 14, 39, 20, 70, 69, 38] and the references therein). Almost all of these build on either private linear broadcast encryption (PLBE) [12] or fingerprinting codes [13] (on which we will elaborate shortly). Tracing using these techniques is very different from the tracing algorithms we construct in this paper. It is an interesting open question whether inspiration can be drawn from the rich traitor tracing literature to construct new traceable secret-sharing schemes.

² If R may sometimes err, then our tracing algorithm requires just one *successful* query to R , which is again optimal.

Several works have considered the setting in which the tracer in a traitor tracing scheme might be malicious, and the task of preventing it from falsely accusing an innocent receiver [57, 58, 59]. This is analogous to the notion of non-imputability in traceable secret sharing.

Boneh, Partap, and Rotem [11] recently generalized traitor tracing to the threshold decryption setting, in which t out of the n receivers are needed for decryption. They gave formal definitions and efficient constructions. Their constructions also rely on fingerprinting codes, and hence on fundamentally different techniques than our traceable secret-sharing constructions.

Fingerprinting codes, introduced by Boneh and Shaw [13] (see also [10, 68, 56, 8]), are used for fingerprinting digital content. Such codes are “traceable” in the following sense: if a new word is constructed from a subset of the codewords, by way of mixing and matching symbols from these codewords, then this new word can be traced back to at least one of the codewords that contributed to it. Very roughly speaking, our constructions for traceable secret sharing can also be seen as tracing symbols of a codeword back to a codeword, but in a very different sense. We discuss this point at length in Section 5.

Goyal [40] (see also [41]) has suggested the notion of Accountable-Authority Identity-Based Encryption, where a misbehaving key authority that leaks the secret key of a server can be proven malicious and held accountable.

Finally, Bellare, Dai, and Rogaway [4] recently proposed a new strengthening of secret sharing schemes, which they called *adept* secret sharing, offering enhanced secrecy and correctness guarantees. In particular, they considered schemes that ensure the authenticity of the reconstructed secret: the reconstruction algorithm may only output the original secret or fail, even if shares get corrupted. Bellare et al. did not consider accountability for leaking secret shares. It is an interesting direction for future work to construct schemes that satisfy both our traceability notion and the enhanced correctness notions of Bellare et al., as the techniques we use for tracing seem to be at odds with the above-mentioned notion of authenticity.

1.3 Notation And Basic Definitions

In this section, we present the basic notions and cryptographic primitives that are used in this work. For an integer $n \in \mathbb{N}$ we denote by $[n]$ the set $\{1, \dots, n\}$. For a distribution X we denote by $x \leftarrow \$ X$ the process of sampling a value x from the distribution X . Similarly, for a set \mathcal{X} , we denote by $x \leftarrow \$ \mathcal{X}$ the process of sampling a value x from the uniform distribution over \mathcal{X} . For a pair X, Y of distributions defined over the same domain Ω , we denote by $\text{SD}(X, Y)$ the *statistical distance* between them, defined as $\text{SD}(X, Y) = \frac{1}{2} \sum_{\omega \in \Omega} |\Pr[X = \omega] - \Pr[Y = \omega]|$. In Section 3, we make use of the standard cryptographic notion of a one-way function, and the reader is referred to [34] for a formal definition. All other notation and existing algorithms or cryptographic primitive that we rely on will be defined in their respective sections.

2 Traceable Secret Sharing

In this section, we present our definition of traceable threshold secret-sharing schemes (we may sometime call them traceable secret-sharing schemes for short). We follow the definitions of Goyal, Song, and Srinivasan [42], with several key changes that we discuss at the end of the section.

2.1 Syntax and Correctness

A traceable secret-sharing scheme is a 4-tuple $\text{TTSS} = (\text{Share}, \text{Rec}, \text{Trace}, \text{Verify})$ of PPT algorithms. The algorithms **Share** and **Rec** are the standard secret sharing and secret reconstruction algorithms of threshold secret sharing schemes. We do restrict the presentation, however, to secret sharing schemes that are **symmetric**. By that, we mean that the secret sharing procedure works in a particular fashion; instead of a secret sharing algorithm that samples all n shares at once, we consider a **Share** algorithm that produces one share at a time. To produce n shares, **Share** is invoked n times over. Note, however, that to ensure correctness, many secret sharing schemes require that the n shares are correlated beyond just the underlying secret s . For example, in Shamir secret sharing, all shares must lie on the same degree $t-1$ polynomial whose free coefficient is s . To account for this added correlation, we consider secret sharing algorithms that take in a **correlation string** ρ as an additional input. In Shamir secret sharing, for example, this ρ specifies the other $t-1$ coefficients of the said polynomial. Correctness then needs to hold for every choice of ρ , and security is defined over a random choice of ρ .

We now define the syntax of traceable secret-sharing schemes in more detail:

- $\text{Share}(1^\lambda, s, n, t, \rho) \rightarrow (\text{sh}_i, \text{tk}_i, \text{vk}_i)$ is the sharing algorithm.³ It takes as input the security parameter 1^λ , the secret s , the number n of parties, the threshold $t \leq n$, and a correlation string $\rho \in \{0, 1\}^\kappa$, where $\kappa = \kappa(\lambda, n, t) \in \mathbb{N}$.⁴ It outputs a share sh_i , a tracing key component tk_i and a verification key component vk_i .
- $\text{Rec}(\{\text{sh}_{i_1}, \dots, \text{sh}_{i_t}\}) \rightarrow s$ is the deterministic secret reconstruction algorithm. It takes as input t secret shares $\text{sh}_{i_1}, \dots, \text{sh}_{i_t}$ and outputs a secret s . We stress that we do not assume that the reconstruction algorithm has external knowledge regarding the origin of the shares it receives as input, since the scheme is symmetric and all shares come from the same distribution. We write the input as $\{\text{sh}_{i_1}, \dots, \text{sh}_{i_t}\}$ to emphasize that **Rec** is oblivious to the ordering of the shares given to it, but to simplify notation, we may sometime write $\text{Rec}(\text{sh}_{i_1}, \dots, \text{sh}_{i_t})$ instead.
- $\text{Trace}^R(\text{tk}, 1^{1/\delta}) \rightarrow (\mathcal{I}, \pi)$ is the randomized tracing algorithm. It takes as input a tracing key $\text{tk} = (\text{tk}_1, \dots, \text{tk}_n)$ and an error parameter δ , and it gets oracle (black-box) access to a reconstruction box R . The algorithm outputs a subset $\mathcal{I} \subseteq [n]$ of identities of leaking parties and an associated proof π .
- $\text{Verify}(\text{vk}, \mathcal{I}, \pi) \rightarrow \{0, 1\}$ is the deterministic verification algorithm. It takes as input the verification key $\text{vk} = (\text{vk}_1, \dots, \text{vk}_n)$, an alleged traitor subset \mathcal{I} , and an associated proof π , and outputs either 1, implying acceptance of the proof that the parties in \mathcal{I} are guilty, or 0, implying rejection.

Correctness. The correctness requirement for a traceable secret-sharing scheme is the standard correctness property of secret-sharing schemes. That is, any t -tuple of the shares should suffice to correctly reconstruct the secret.

³ In the original definition of [42], the sharing process is done by running a generic two-party computation protocol between the dealer and each of the shareholders. We chose to abstract this process away. We elaborate on the differences between the definitions below.

⁴ We allow κ to depend on n for generality, but in our constructions, κ will only be a function of λ and t . This allows for sampling of shares “on the fly”, without knowing n in advance. Throughout the paper, κ will always refer to the bit-length of the correlation string, and when clear from context, we will not mention this explicitly.

Definition 1. Let $\text{TTSS} = (\text{Share}, \text{Rec}, \text{Trace}, \text{Verify})$ be a traceable secret-sharing scheme and let $\epsilon = \epsilon(\lambda) \in [0, 1]$ be a function of the security parameter λ . We say that TTSS is ϵ -correct if for every $\lambda \in \mathbb{N}$, every secret s , every $n \in \mathbb{N}$, every $0 < t \leq n$, every correlation string $\rho \in \{0, 1\}^\kappa$, and every subset $\mathcal{J} = \{i_1, \dots, i_t\} \subseteq [n]$ of size t , it holds that

$$\Pr [s' = s] \geq 1 - \epsilon(\lambda),$$

where $(\text{sh}_i, \text{tk}_i, \text{vk}_i) \leftarrow \text{Share}(1^\lambda, s, n, t, \rho)$ for every $i \in [n]$ and $s' := \text{Rec}(\text{sh}_{i_1}, \dots, \text{sh}_{i_t})$.

2.2 Security

A traceable secret-sharing scheme should satisfy three security properties: secrecy, traceability, and non-imputability. The first is the standard secrecy property of secret-sharing schemes, which stipulates that any subset of less than t parties should learn nothing about the secret s . This is captured by the following definition.

Definition 2 (Secrecy). Let $\text{TTSS} = (\text{Share}, \text{Rec}, \text{Trace}, \text{Verify})$ be a traceable secret-sharing scheme for secrets in a secret space $\mathcal{S} = \{\mathcal{S}_\lambda\}_{\lambda \in \mathbb{N}}$. Let $\epsilon = \epsilon(\lambda) \in [0, 1]$ be a function of the security parameter λ . We say that TTSS satisfies ϵ -secrecy if for every PPT algorithm D , every $\lambda \in \mathbb{N}$, any two secrets s, s' , every $n \in \mathbb{N}$, every $0 < t \leq n$, and every subset $\mathcal{J} = \{i_1, \dots, i_{t-1}\} \subseteq [n]$ of size $t - 1$, it holds that

$$\left| \Pr [D(\text{sh}_{i_1}, \dots, \text{sh}_{i_{t-1}}) = 1] - \Pr [D(\text{sh}'_{i_1}, \dots, \text{sh}'_{i_{t-1}}) = 1] \right| \leq \epsilon(\lambda)$$

where $\rho \leftarrow \{0, 1\}^\kappa$, $(\text{sh}_i, \text{tk}_i, \text{vk}_i) \leftarrow \text{Share}(1^\lambda, s, n, t, \rho)$ for every $i \in [n]$, and $(\text{sh}'_i, \text{tk}'_i, \text{vk}'_i) \leftarrow \text{Share}(1^\lambda, s', n, t, \rho)$ for every $i \in [n]$.

In this work, we will primarily be interested in schemes in which secrecy is statistical and the distinguisher D may be unbounded.

For traceable secret-sharing schemes, it makes sense to require that secrecy is maintained even against the tracer. That is, that the tracing key tk does not help in breaking secrecy. This is made formal in the following definition.

Definition 3 (Tracer secrecy). Let $\text{TTSS} = (\text{Share}, \text{Rec}, \text{Trace}, \text{Verify})$ be a traceable secret-sharing scheme for secrets in a secret space $\mathcal{S} = \{\mathcal{S}_\lambda\}_{\lambda \in \mathbb{N}}$. Let $\epsilon = \epsilon(\lambda) \in [0, 1]$ be a function of the security parameter λ . We say that TTSS satisfies ϵ -tracer secrecy if for every PPT algorithm D , every $\lambda \in \mathbb{N}$, any two secrets s, s' , every $n \in \mathbb{N}$, every $0 < t \leq n$, and every subset $\mathcal{J} = \{i_1, \dots, i_{t-1}\} \subseteq [n]$ of size $t - 1$, it holds that

$$\left| \Pr [D(\text{tk}, \text{sh}_{i_1}, \dots, \text{sh}_{i_{t-1}}) = 1] - \Pr [D(\text{tk}', \text{sh}'_{i_1}, \dots, \text{sh}'_{i_{t-1}}) = 1] \right| \leq \epsilon(\lambda)$$

where $\rho \leftarrow \{0, 1\}^\kappa$, $(\text{sh}_i, \text{tk}_i, \text{vk}_i) \leftarrow \text{Share}(1^\lambda, s, n, t, \rho)$ for every $i \in [n]$, and $(\text{sh}'_i, \text{tk}'_i, \text{vk}'_i) \leftarrow \text{Share}(1^\lambda, s', n, t, \rho)$ for every $i \in [n]$, $\text{tk} = (\text{tk}_1, \dots, \text{tk}_n)$ and $\text{tk}' = (\text{tk}'_1, \dots, \text{tk}'_n)$.

Jumping ahead, the scheme obtained from the simple tracing algorithm in Appendix A does not satisfy tracer secrecy. There, the tracing key completely reveals the secret. Our Shamir-based scheme (Section 3), on the other hand, does provide tracer secrecy, as the tracing key does not depend on the secret. Our Blakley-based Scheme (Section 4) does not provide tracer secrecy: the

tracing key leaks information about the secret, and in particular, allows to test whether a “candidate secret” is really the secret or not. However, it is still better than the simple traceable scheme from Appendix A in that it does not completely leak the secret. In fact, one can prove that recovering the secret from tk and $t - 1$ shares amounts to breaking discrete log in some discrete-log-hard group.

Traceability. In addition, a traceable secret-sharing should provide *traceability*. Suppose a coalition $\mathcal{I} \subseteq [n]$ of parties, of size $f < t$, gets together and constructs a reconstruction box R using their shares. This R is an algorithm that takes in additional $t - f$ secret shares and outputs a secret s . Intuitively, if this R is a “good” reconstruction box, then it should be possible to trace it back to at least one of the parties who “contributed” its share to it. As we discuss in the introduction, tracing R back to the corrupted parties should be done given only black-box access to it.

In more detail, say that the secret shared among the parties is s^* , and the shares of all parties are $\text{sh}_1, \dots, \text{sh}_n$. Then, R is a good reconstruction box if for random $t - f$ additional shares for the secret s^* , denoted $\text{sh}'_1, \dots, \text{sh}'_{t-f}$, it holds that $R(\text{sh}'_1, \dots, \text{sh}'_{t-f})$ outputs s^* with high probability. Definition 4 below formally defines good reconstruction boxes.

Definition 4 (Good reconstruction boxes). Let TTSS be a traceable secret-sharing scheme. Let $\lambda \in \mathbb{N}$, let $n, t, f \in \mathbb{N}$ such that $0 < f < t \leq n$, and let $\kappa = \kappa(\lambda, n, t)$. For $\epsilon \in [0, 1]$, a secret s^* , correlation string $\rho \in \{0, 1\}^\kappa$, we say that a reconstruction box R is $(n, t, f, s^*, \rho, \epsilon)$ -good if

$$\Pr [R(\text{sh}'_1, \dots, \text{sh}'_{t-f}) = s^*] \geq \epsilon,$$

where the probability is taken over $(\text{sh}'_i, \text{tk}'_i, \text{vk}_i) \leftarrow \text{Share}(1^\lambda, s^*, n, t, \rho)$ for $i = 1, \dots, t - f$ and the random coins of R .

Equipped with this notion, the traceability experiment is presented in Fig. 1 and the notion of traceability for secret-sharing is defined below.

Experiment $\text{ExpTrace}_{\mathcal{A}, \text{TTSS}, \epsilon, \delta}(\lambda)$
1 : $\mathcal{I} \leftarrow \emptyset$
2 : $(n, t, \mathcal{I}, \text{state}) \leftarrow \mathcal{A}(1^\lambda)$ // \mathcal{A} chooses the parties $\mathcal{I} = \{i_1, \dots, i_f\} \subseteq [n]$ to corrupt
3 : $s^* \leftarrow \text{SCR}\mathcal{T}_\lambda$
4 : $\rho \leftarrow \{0, 1\}^\kappa$
5 : for $i = 1, \dots, n$: $(\text{sh}_i, \text{tk}_i, \text{vk}_i) \leftarrow \text{Share}(1^\lambda, s^*, n, t, \rho)$
6 : $\text{tk} := (\text{tk}_1, \dots, \text{tk}_n)$, $\text{vk} := (\text{vk}_1, \dots, \text{vk}_n)$
7 : $R \leftarrow \mathcal{A}(\text{state}, \text{sh}_{i_1}, \dots, \text{sh}_{i_f})$ where $\mathcal{I} = \{i_1, \dots, i_f\}$
8 : if R is not $(n, t, f, s^*, \rho, \epsilon)$ -good then return 0
9 : $(\mathcal{I}', \pi) \leftarrow \text{Trace}^{R(\cdot)}(\text{tk}, 1^{1/\delta(\lambda)})$
10 : if $\mathcal{I} = \mathcal{I}'$ AND $\text{Verify}(\text{vk}, \mathcal{I}', \pi) = 1$ then return 0, else return 1

Fig. 1. The tracing experiment for a traceable secret-sharing scheme TTSS and an adversary \mathcal{A} . The set $\text{SCR}\mathcal{T} = \{\text{SCR}\mathcal{T}_\lambda\}_{\lambda \in \mathbb{N}}$ denotes the secret space of TTSS.

Definition 5 (Traceability). Let $\text{TTSS} = (\text{Share}, \text{Rec}, \text{Trace}, \text{Verify})$ be a traceable secret-sharing scheme with secret space $\text{SCR}\mathcal{T} = \{\text{SCR}\mathcal{T}_\lambda\}_{\lambda \in \mathbb{N}}$, and let $\epsilon = \epsilon(\lambda)$ and $\delta = \delta(\lambda)$ be functions of

the security parameter. We say that TTSS is **traceable**, if for every probabilistic polynomial time adversary \mathcal{A} , the following function is negligible in λ :

$$\text{Adv}_{\mathcal{A}, \text{TTSS}, \epsilon, \delta}^{\text{trac}}(\lambda) := \left| \epsilon \cdot \Pr [\mathbf{ExpTrace}_{\mathcal{A}, \text{TTSS}, \epsilon, \delta}(\lambda) = 1] - \frac{1}{|\mathcal{SCRT}_\lambda|} \right|.$$

Some remarks are in order:

1. **The advantage term:** The advantage is defined to be

$$\epsilon \cdot \Pr [\mathbf{ExpTrace}_{\mathcal{A}, \text{TTSS}, \epsilon, \delta}(\lambda) = 1] - 1/|\mathcal{SCRT}_\lambda|$$

(instead of simply $\Pr [\mathbf{ExpTrace}_{\mathcal{A}, \text{TTSS}, \epsilon, \delta}(\lambda) = 1]$) to discount for the following trivial attack, which succeeds in “breaking” traceability with probability essentially $1/(\epsilon \cdot |\mathcal{SCRT}_\lambda|)$. Fix a set $\mathcal{S} \subseteq \mathcal{SCRT}_\lambda$. The adversary then outputs a reconstruction box $R_{\mathcal{S}}$ that ignores the shares it receives as input, and outputs a uniformly random secret from \mathcal{S} , i.e., $s \leftarrow_{\$} \mathcal{S}$. Observe that this reconstruction box is essentially untraceable, since it is independent of the shares of the corrupted parties. It is also ϵ -good for any secret $s \in \mathcal{S}$ for $\epsilon = 1/|\mathcal{S}|$. Moreover, the secret s^* chosen by the challenger in $\mathbf{ExpTrace}_{\mathcal{A}, \text{TTSS}, \epsilon, \delta}(\lambda)$ is in \mathcal{S} with probability

$$\alpha = \frac{|\mathcal{S}|}{|\mathcal{SCRT}_\lambda|} = \frac{1}{\epsilon \cdot |\mathcal{SCRT}_\lambda|}.$$

Note that if $|\mathcal{SCRT}_\lambda|$ is super-polynomial in λ , then ϵ or α (or both) must be negligible.

2. **Adaptive adversaries:** For simplicity, our definition requires that the adversary chooses the set of corrupted parties all at once, non-adaptively. However, observe that in symmetric secret sharing schemes, as we consider here, all shares come from the same distribution, and hence querying for shares adaptively would give no additional power to the adversary.
3. **The output of R :** Our definition requires that a good reconstruction R outputs the reconstructed secret s^* , and not some function thereof. For example, if R outputs $H(s^*)$ for some hash function H , our definition makes no guarantees as to the ability to trace R back to the corrupted subset. Looking ahead, our tracing algorithms will indeed rely on R outputting the secret s^* and not, say, $H(s^*)$. This might pose an issue if $H(s^*)$ is also very sensitive information; for instance, if $H(s^*)$ is used by Alice, the secret owner, as her Bitcoin secret key. We argue, however, that this can be justified. The secret owner is the one that gets to decide what the secret is. Therefore, Alice can set the secret s^* to be her secret key, and not a hash of it, in which case a reconstruction R that outputs $H(s^*)$ is of little use.
4. **Private tracing:** In the tracing experiment, the adversary is not given access to the tracing key tk and the verification key vk . This is because, for our constructions, knowledge of these keys can allow an adversary to evade tracing. We leave the problem of constructing secret sharing schemes with public tracing as an avenue for future work.

Non-imputability. Lastly, we require that a traceable secret-sharing scheme satisfy the notion of non-imputability, stating that honest parties cannot be wrongly blamed. In more detail, this means that even a malicious tracer cannot produce an accepting proof for the culpability of an honest party. This is formally captured by the security experiment in Fig. 2. There, the adversary gets the tracing and verification keys tk and vk , and may obtain the secret shares of any party of its choosing, as long as it is not the honest party it tries to blame.

Definition 6 (Non-imputability). We say that a traceable secret-sharing scheme $\text{TTSS} = (\text{Share}, \text{Rec}, \text{Trace}, \text{Verify})$ satisfies **non-imputability**, if for every probabilistic polynomial time adversary \mathcal{A} , the following function is negligible in λ :

$$\text{Adv}_{\mathcal{A}, \text{TTSS}}^{\text{ni}}(\lambda) := \Pr [\mathbf{ExpNI}_{\mathcal{A}, \text{TTSS}}(\lambda) = 1]$$

Experiment $\mathbf{ExpNI}_{\mathcal{A}, \text{TTSS}}(\lambda)$
1 : $\mathcal{J} \leftarrow \emptyset$
2 : $(n, t, i^*, s, \text{state}) \leftarrow \mathcal{A}(1^\lambda)$
3 : $\rho \leftarrow_{\$} \{0, 1\}^\kappa$
4 : for $i = 1, \dots, n$: $(\text{sh}_i, \text{tk}_i, \text{vk}_i) \leftarrow_{\$} \text{Share}(1^\lambda, s^*, n, t, \rho)$
5 : $\text{tk} := (\text{tk}_1, \dots, \text{tk}_n), \text{vk} := (\text{vk}_1, \dots, \text{vk}_n)$
6 : $(\mathcal{I}^*, \pi) \leftarrow_{\$} \mathcal{A}(\text{state}, \{\text{sh}_i\}_{i \in [n] \setminus \{i^*\}}, \text{tk}, \text{vk})$ <div style="padding-left: 20px;">// output a subset \mathcal{I}^* and a proof for its culpability</div>
7 : if $i^* \in \mathcal{I}^*$ AND $\text{Verify}(\text{vk}, \mathcal{I}^*, \pi) = 1$ then return 1, else return 0

Fig. 2. The non-imputability experiment for a traceable secret-sharing scheme TTSS and an adversary \mathcal{A} .

Comparison with the definition of Goyal, Song, and Srinivasan [42]. As mentioned above, our definition takes after that of Goyal et al. [42]. There are, however, a few differences that make our definition incomparable to theirs. On the one hand, our definition gives the tracer much less information to work with, and hence provides stronger security guarantees in this sense. Concretely, in the work of Goyal et al. the leakage of secrets has more structure to it: each corrupted party i submits a function $f(\text{sh}_i)$ of its secret share for an adversarially-chosen f .⁵ The tracer then gets $\{f(\text{sh}_i)\}_{i \in \mathcal{J}}$, where \mathcal{J} is the set of corrupted parties of size at least t , together with a description of the function f . It also gets access to a reconstruction box R , that takes in at least t share encodings $\{f(\text{sh}_i)\}$ as input. On its own, the reconstruction box R is independent of the secret shares. Observe that, indeed, given this information, it is easy to construct a reconstruction box R' in our model, with respect to any subset $\mathcal{I} \subset \mathcal{J}$ of corrupted parties. To do that, one simply hardcodes $\{f(\text{sh}_i)\}_{i \in \mathcal{I}}$ and the function f into R' . Then, on input secret shares $(\text{sh}'_1, \dots, \text{sh}'_{t-|\mathcal{I}|})$, the reconstruction box R' , computes $R(\{f(\text{sh}_i)\}_{i \in \mathcal{I}}, \{f(\text{sh}'_j)\}_{j \in [t-|\mathcal{I}]})$.

On the other hand, in our setting, to trace a reconstruction box R back to a corrupted party, we require that R fully recovers the underlying secret. Goyal et al., however, only require that the reconstruction box R is able to extract some non-trivial information regarding the secret in order to trace it back to a corrupted party (this is formalized by R accomplishing some distinguishing task; see [42] for the formal details). These two requirements are equivalent in the random oracle model, assuming that the user always hashes the secret with a random oracle prior to using it. This way, a reconstruction box that outputs partial information about the hash of the secret, must query the random oracle at the secret. By observing R 's queries to the random oracle, one can extract the

⁵ In the syntax of Goyal et al. each party i may submit a different, party-specific, function f_i of its share. In our setting, however, any information that pertains to the identity of the party who owns the share is considered part of the share. Hence, using just one function f for all shares is without loss of generality.

full secret s thereby obtaining a reconstruction box in our model. Even without the random oracle model, the user can apply a randomness extractor to the secret before using it, thus making a box R that only returns a few bits of information about the secret of no value to the adversary.

On a technical level, Goyal et al. consider a distributed and interactive secret-sharing protocol between the dealer and the n shareholders. The dealer’s view of this interaction serves as both the tracing key tk and the verification key vk . In order to instantiate this sharing step, they use a generic protocol for secure 2-party computation (instantiated n times over, once between the dealer and each of the shareholders). To simplify the presentation, we chose to consider a static and centralized secret sharing algorithm Share , but we note that one can always transform it into a distributed protocol using generic multiparty computation as well.

2.3 A Useful Fact About Good Reconstruction Boxes

In the rest of this section, we introduce a different traceability notion, which we call *universal traceability*. This notion is easier to work with when proving traceability of traceable secret sharing schemes, and — as we will shortly see — any scheme that is universally traceable is also traceable in some formal sense. As we will argue, universal traceability also makes sense as a notion of traceability in its own right.

Looking ahead, universal traceability says that a *universally good* reconstruction box can be traced back to the subset of parties who manufactured it. Very informally, we call a reconstruction box universally good if it correctly reconstructs a random secret from a random sharing of it with high probability.

We begin by defining universally good reconstruction boxes. Our definition makes use of the following notation. For a subset $\mathcal{C} \subseteq \{0, 1\}^\kappa$ of correlation strings, we denote by $\text{sh} \leftarrow_{\$} \mathcal{SH}_{\lambda, n, t}(\mathcal{C})$ the process of sampling a random share for a random secret, conditioned on the correlation string being in \mathcal{C} ; that is, the process: $\rho \leftarrow_{\$} \mathcal{C}$, $s \leftarrow_{\$} \text{SCRT}_{\lambda}$, $\text{sh} \leftarrow_{\$} \text{Share}(1^\lambda, s, n, t, \rho)$. When λ, n and t are clear from context, we may write $\text{sh} \leftarrow_{\$} \mathcal{SH}(\mathcal{C})$.

Definition 7 (Universally-good reconstruction boxes). *Let TTSS be a traceable secret-sharing scheme. Let $\lambda \in \mathbb{N}$, let $n, t, f \in \mathbb{N}$ such that $0 < f < t \leq n$, and let $\kappa = \kappa(\lambda, n, t)$. For $\epsilon \in [0, 1]$, f shares $\mathbf{sh} = (\text{sh}_1, \dots, \text{sh}_f)$, and a subset $\mathcal{C} \subseteq \{0, 1\}^\kappa$, we say that a reconstruction box R is $(n, t, \mathbf{sh}, \mathcal{C}, \epsilon)$ -good if*

$$\Pr [R(\text{sh}'_1, \dots, \text{sh}'_{t-f}) = \text{Rec}(\text{sh}_1, \dots, \text{sh}_f, \text{sh}'_1, \dots, \text{sh}'_{t-f})] \geq \epsilon,$$

where the probability is taken over $\text{sh}'_i \leftarrow_{\$} \mathcal{SH}_{\lambda, n, t}(\mathcal{C})$ for $i = 1, \dots, t - f$ and the random coins of R .

We now use the notion of universally good reconstruction boxes to define universal traceability. Our definition makes use of the following conventions. Let $\Gamma = \{\mathcal{C}_1, \mathcal{C}_2, \dots\}$ be a partition of the space $\{0, 1\}^\kappa$ of correlation strings. For a correlation string $\rho \in \{0, 1\}^\kappa$, we denote by $\Gamma(\rho)$ the unique subset \mathcal{C}_i that contains ρ .

Using this notation, the universal traceability experiment is presented in Fig. 3. It is almost identical to the tracing security experiment from Fig. 1 other than the fact that the reconstruction box R is required to be universally good with respect to $\Gamma(\rho)$, where ρ is the correlation string used to generate the secret shares in the experiment (rather than just good with respect to the secret s^*). For simplicity, we only parameterize the experiment by ϵ , and not δ , by setting $\epsilon = \delta$.

Experiment $\mathbf{ExpUniTrace}_{\mathcal{A}, \text{TTSS}, \Gamma, \epsilon}(\lambda)$
1: $\mathcal{J} \leftarrow \emptyset$
2: $(n, t, \mathcal{I}, \text{state}) \leftarrow \mathcal{A}(1^\lambda)$
3: $s^* \leftarrow_{\$} \text{SCRT}_\lambda$
4: $\rho \leftarrow_{\$} \{0, 1\}^\kappa$
5: for $i = 1, \dots, n$: $(\text{sh}_i, \text{tk}_i, \text{vk}_i) \leftarrow_{\$} \text{Share}(1^\lambda, s^*, n, t, \rho)$
6: $\text{tk} := (\text{tk}_1, \dots, \text{tk}_n)$, $\text{vk} := (\text{vk}_1, \dots, \text{vk}_n)$
7: $R \leftarrow_{\$} \mathcal{A}(\text{state}, \text{sh}_{i_1}, \dots, \text{sh}_{i_f})$ where $\mathcal{I} = (i_1, \dots, i_f)$
8: $\mathbf{sh} := (\text{sh}_{i_1}, \dots, \text{sh}_{i_f})$
9: if R is not $(n, t, \mathbf{sh}, \Gamma(\rho), \epsilon)$ -universally-good then return 0
10: $(\mathcal{I}', \pi) \leftarrow_{\$} \text{Trace}^{R(\cdot)}(\text{tk}, 1^{1/\epsilon(\lambda)})$
11: if $\mathcal{I} = \mathcal{I}'$ AND $\text{Verify}(\text{vk}, \mathcal{I}', \pi) = 1$ then return 0, else return 1

Fig. 3. The universal tracing experiment for a traceable secret-sharing scheme TTSS and an adversary \mathcal{A} . Changes from the tracing security experiment in Fig. 1 are marked in blue.

Definition 8 (Universal traceability). Let $\text{TTSS} = (\text{Share}, \text{Rec}, \text{Trace}, \text{Verify})$ be a traceable secret-sharing scheme with secret space $\text{SCRT} = \{\text{SCRT}_\lambda\}_{\lambda \in \mathbb{N}}$, let $\epsilon = \epsilon(\lambda)$ be a function of the security parameter, and let Γ be a partition of the space $\{0, 1\}^\kappa$ of correlation strings. We say that TTSS satisfies universal traceability, if for every probabilistic polynomial time adversary \mathcal{A} , the following function is negligible in λ :

$$\text{Adv}_{\mathcal{A}, \text{TTSS}, \Gamma, \epsilon}^{\text{uni-trac}}(\lambda) := \left| \Pr [\mathbf{ExpUniTrace}_{\mathcal{A}, \text{TTSS}, \Gamma, \epsilon}(\lambda) = 1] - \frac{1}{|\text{SCRT}_\lambda|} \right|.$$

If Γ is the trivial partition (that is, $\Gamma = \{\{0, 1\}^\kappa\}$), we may omit it from the notation. Observe that the trivial reconstruction box R that outputs a uniformly random secret $s^* \leftarrow_{\$} \text{SCRT}_\lambda$, independently of its inputs, is ϵ -universally good for $\epsilon = 1/|\text{SCRT}_\lambda|$ and is essentially untraceable. This is why we subtract a $1/|\text{SCRT}_\lambda|$ term from the probability that the experiment outputs 1 when defining the advantage of \mathcal{A} above.

We argue that Definition 8 is a very reasonable one. Intuitively, the definition says that if R is good with respect to the equivalence class $\Gamma(\rho)$ of the *real correlation string* ρ (i.e., the one used to sample the real secret shares sampled by the challenger), then it should be traced back to the corrupted subset. Intuitively, since \mathcal{A} only sees at most $t - 1$ shares, and does not have any information about the real correlation string ρ , it is very reasonable to consider a reconstruction box R that almost always fails on $\Gamma(\rho)$ to be a bad reconstruction box. To make this intuition precise, we prove below that for a certain natural class of traceable secret sharing schemes (as the ones we will construct), if a traceable secret sharing scheme satisfies universal traceability (Definition 8), then it also satisfies standard traceability (Definition 5) with related parameters.

Looking ahead, we will use universal traceability with respect to two extremes of the partition Γ . In Section 3, we will use the trivial partition $\Gamma = \{\{0, 1\}^\kappa\}$; meaning, a universally good reconstruction box should reconstruct random sharings of a random secret. In Section 4 we will use the complete partition into singletons $\Gamma = \{\{\rho\} : \rho \in \{0, 1\}^\kappa\}$; meaning, a universally good reconstruction box should reconstruct random sharings (of random secrets), generated with the *real reconstruction string* ρ sampled by the challenger.

Bidirectional secret sharing schemes. To relate the notion of universal traceability to that of standard traceability, it will be convenient to define a subclass of threshold secret sharing schemes, which we call *bidirectional*. Let Γ be a partition of the space of correlation strings. In such schemes, sampling a secret s^* uniformly at random from the secret space, and then generating t shares $\text{sh}_1, \dots, \text{sh}_t$ of it conditioned on the correlation string being in some $\mathcal{C} \in \Gamma$, is statistically close to sampling t shares $\text{sh}_1, \dots, \text{sh}_t \leftarrow_{\$} \mathcal{SH}(\mathcal{C})$, and then computing $s^* \leftarrow \text{Rec}(\text{sh}_1, \dots, \text{sh}_t)$. Both of our constructions will satisfy this property with respect to the different partitions discussed above.

Definition 9 (Bidirectional secret sharing). Let TTSS be a traceable secret-sharing scheme with secret space $\mathcal{SCR}\mathcal{T} = \{\mathcal{SCR}\mathcal{T}_\lambda\}_{\lambda \in \mathbb{N}}$, and let $\epsilon = \epsilon(\lambda) \in [0, 1]$, $n = n(\lambda)$ and $t = t(\lambda)$ be functions of the security parameter $\lambda \in \mathbb{N}$. Let Γ be a partition of the space $\{0, 1\}^k$ of correlation strings for TTSS. We say that TTSS is (Γ, ϵ) -bidirectional if for every $\lambda \in \mathbb{N}$ and every $\mathcal{C} \in \Gamma$,

$$\text{SD}((s, \text{sh}_1, \dots, \text{sh}_t), (s', \text{sh}'_1, \dots, \text{sh}'_t)) \leq \epsilon(\lambda),$$

where:

1. $s \leftarrow_{\$} \mathcal{SCR}\mathcal{T}_\lambda$, $\rho \leftarrow_{\$} \mathcal{C}$, and $(\text{sh}_i, \text{tk}_i, \text{vk}_i) \leftarrow_{\$} \text{Share}(1^\lambda, s, n, t, \rho)$ for $i = 1, \dots, t$.
2. $\text{sh}'_i \leftarrow_{\$} \mathcal{SH}_{\lambda, n, t}(\mathcal{C})$ for $i = 1, \dots, t$ and $s' \leftarrow \text{Rec}(\text{sh}'_1, \dots, \text{sh}'_t)$.

From universal traceability to traceability. As promised, we will now relate standard traceability (Definition 5) to universal traceability (Definition 8) for bidirectional schemes.

Lemma 1. Let TTSS = (Share, Rec, Trace, Verify) be a traceable secret-sharing scheme that satisfies ν_1 -correctness and (Γ, ν_2) -bidirectionality for negligible functions ν_1, ν_2 of the security parameter $\lambda \in \mathbb{N}$ and a partition Γ of $\{0, 1\}^k$. Let $\epsilon = \epsilon(\lambda)$ and $\delta = \delta(\lambda)$ be functions of the security parameter such that $\delta \leq \min\{2\epsilon, 1/2\}$, and let \mathcal{A} be an adversary. Assume that $\text{Adv}_{\mathcal{A}, \text{TTSS}, \epsilon, \delta}^{\text{trac}}(\lambda) \geq 2\delta$ and $\Pr[\text{ExpTrace}_{\mathcal{A}, \text{TTSS}, \epsilon, \delta}(\lambda) = 1] \geq \frac{1}{\epsilon|\mathcal{SCR}\mathcal{T}_\lambda|}$. Then, there exists a negligible function $\nu = \nu(\cdot)$ such that for every $\lambda \in \mathbb{N}$ it holds that

$$\text{Adv}_{\mathcal{A}, \text{TTSS}, \Gamma, \delta'}^{\text{uni-trac}}(\lambda) \geq \text{Adv}_{\mathcal{A}, \text{TTSS}, \epsilon, \delta}^{\text{trac}}(\lambda)/2 - \nu(\lambda)$$

where $\delta' \geq \delta - \nu(\lambda)$.

The proof of the lemma can be found in Appendix B. The crux of the argument is this: since the adversary sees at most $t - 1$ shares, it has essentially no information about the secret s^* . Hence, if it outputs a reconstruction box that is ϵ -good for the real secret s^* with probability α , then it is good for $\alpha/2$ fraction of the secrets with probability $\alpha/2$. In particular, with probability $\alpha/2$, it is δ' -universally-good for $\delta' \approx \alpha/2$. This means that if α is non-negligible, then the adversary outputs a δ' -universally-good reconstruction box with non-negligible probability, and for a non-negligible δ' . In particular, setting $\alpha \geq 2\delta$ gives $\delta' \approx \delta$.

3 A Scheme Based on Shamir Secret Sharing

In this section, we present our traceable secret-sharing scheme based on Shamir secret sharing. Recall that Shamir secret sharing is defined over a finite field \mathbb{F} . To share a secret $s \in \mathbb{F}$ among n parties, with a threshold of t , the dealer samples a uniformly random polynomial q of degree $t - 1$ over \mathbb{F} , conditioned on its free coefficient being s . Each party $i \in [n]$ is uniquely associated with

some field element $x_i \in \mathbb{F}$, and its share is $\text{sh}_i = (x_i, q(x_i)) \in \mathbb{F}$.⁶ Any t evaluations of q can be used to reconstruct the polynomial, and hence s , e.g. via Lagrange interpolation. In contrast, any tuple of $t - 1$ shares is uniformly random over \mathbb{F}^{t-1} , and hence reveals nothing about the secret.

We begin with an informal overview of our traceable Shamir secret sharing. Then, we will present a basic scheme that does not satisfy non-imputability, but already captures the main ideas behind our tracing procedure. Finally, we will show how to amend this scheme so that it also satisfies non-imputability.

3.1 Overview of our scheme

Random evaluation points. Typically in Shamir secret sharing, party i is deterministically associated with its evaluation point x_i . The simplest example is when the field \mathbb{F} is \mathbb{F}_p , the finite field of integers modulo p . Then, a natural choice is $x_i = i$, when i is interpreted as an \mathbb{F}_p element.

Our first observation is that this approach cannot admit efficient tracing according to our definition. The reason is that an adversary may corrupt a random subset \mathcal{I} of the parties and construct a reconstruction box R with the secret shares of \mathcal{I} embedded in it. Let $f = |\mathcal{I}|$. Then, R will output the secret s if it gets exactly $t - f$ shares associated with parties outside of \mathcal{I} ; the box R can link a share $(x_i, q(x_i))$ given to it as input to party i , because x_i is a deterministic function of the party’s index i . Whenever R is given more than $t - f$ inputs, or if one of its input shares “belongs” to a party in \mathcal{I} , the box R outputs \perp and refuses to output anything meaningful. In this scenario, regardless of the tracing strategy, R outputs anything other than \perp with probability at most $\approx Q \binom{n-(t-f)}{f} / \binom{n}{f}$, where Q is the number of queries that Trace makes to R . To see why that is, consider a different experiment, in which the tracer always gets \perp in response from R , but we still say that it wins if it ever queries R at a subset that does not intersect the corrupted subset \mathcal{I} . The probability that the tracer wins in this modified experiment is the same as in the original experiment since its view is identical in both experiments as long as it has not yet won. In this modified experiment, the queries of the tracer are independent of the corrupted subset \mathcal{I} , and so we can think of \mathcal{I} as being chosen uniformly at random after the queries are determined. In this case, the probability that \mathcal{I} does not intersect a certain subset of size $t - f$ is $\binom{n-(t-f)}{f} / \binom{n}{f}$, and the overall probability that the tracer wins is bounded by $Q \cdot \binom{n-(t-f)}{f} / \binom{n}{f}$. If, for example, $t = c_1 n$ and $f = c_2 t$ for constants $c_1, c_2 \in (0, 1)$ and Q is polynomial in n , making R output anything other than \perp would require $2^{\Omega(n)}$ queries in expectation.

To avoid this problem, we consider a variant of Shamir secret sharing in which each x_i is sampled uniformly at random from the field \mathbb{F} . If the field is large enough, this has a very small impact on the correctness of the scheme. Rejection sampling may also be used to reduce, or even eliminate, the small correctness error that results from the possibility of assigning the same x to two different parties. The benefit is that now a share $(x_i, q(x_i))$ cannot be linked to party i without knowing the dealer’s randomness. As we will now see, this allows us to overcome the impossibility argument sketched above.

Tracing via polynomial interpolation and factorization. To understand the basic idea behind our tracing procedure, it is instructive to consider the task that a pirate reconstruction box R for Shamir secret sharing needs to accomplish. For simplicity, assume that the number of corruptions

⁶ If x_i is a deterministic function of some identity information associated with the party, then x_i need not be explicitly included as part of the share. Looking ahead, we will choose the x_i s randomly, and hence they will need to be included in the share.

is $f = t - 1$. In this case, R gets an additional share as input and outputs the secret s . This secret is the result of reconstructing the value $q(0)$ using the $t - 1$ evaluations of q hardcoded in R and the additional evaluation given to it as input. Suppose that the shares hardcoded in R are $(x_i, y_i = q(x_i))$ for $i = 1, \dots, t - 1$. When feeding it an additional share of the form (x_t, y_t) , R needs to output

$$s = \sum_{i \in [t]} \left(\prod_{k \in [t] \setminus \{i\}} \frac{x_k}{x_k - x_i} \right) \cdot y_i.$$

Otherwise, it is not a good reconstruction box. This can be re-written as

$$s = \sum_{i \in [t-1]} \left(\prod_{k \in [t] \setminus \{i\}} \frac{x_k}{x_k - x_i} \right) \cdot y_i + \left(\prod_{k \in [t-1]} \frac{x_k}{x_k - x_t} \right) \cdot y_t. \quad (1)$$

If we now feed R with the share $(x_t, y_t + 1)$, where x_t and y_t are as before, we obtain some s' satisfying

$$s' = \sum_{i \in [t-1]} \left(\prod_{k \in [t] \setminus \{i\}} \frac{x_k}{x_k - x_i} \right) \cdot y_i + \left(\prod_{k \in [t-1]} \frac{x_k}{x_k - x_t} \right) \cdot (y_t + 1). \quad (2)$$

Subtracting Eq. (1) from Eq. (2) and rearranging, we obtain

$$\prod_{k \in [t-1]} \frac{x_k - x_t}{x_k} = (s' - s)^{-1}. \quad (3)$$

We now consider the univariate polynomial $h(X) = \prod_{k \in [t-1]} \frac{x_k - X}{x_k}$ in the indeterminate X . Observe that the roots of the polynomial h are exactly the x_i values of the corrupted parties. Moreover, we can interpret Eq. (3) as an evaluation of h at the point x_t we fed to R . Repeating the above with additional $t - 1$ fresh x_t values, would give us t evaluations of h . Since h is of degree $t - 1$, this is enough to interpolate h and factor it to find its roots via polynomial factorization [5, 61, 47]. If the true x_i values of all the parties are given to the tracer as part of tk , the tracer can now trace the roots of h back to the corrupted parties.

Sharing and reconstruction. In what follows, we formally describe the **Share** and **Rec** algorithms for our scheme, which we denote **TS** (for “Traceable Shamir”). This is still a simplified version of our scheme that does not provide non-imputability, which we will handle later in this section. The scheme is parameterized by a finite field $\mathbb{F} = \mathbb{F}(\lambda, n, t)$.

We begin with the sharing algorithm. The algorithm receives the security parameter 1^λ , the secret $s \in \mathbb{F}$, integers n and t , and a correlation string ρ that is interpreted as $t - 1$ coefficients in \mathbb{F} .

Share($1^\lambda, s, n, t, \rho$) :

1. Parse ρ as $(a_1, \dots, a_{t-1}) \in \mathbb{F}^{t-1}$.
// observe that $q(X) = a_{t-1}X^{t-1} + \dots + a_1X + s$ is a polynomial of degree $t - 1$ such that $q(0) = s$
2. Sample $x_i \leftarrow_{\$} \mathbb{F}$.
3. Compute $y_i \leftarrow s + \sum_{j=1}^{t-1} a_j x_i^j$.
// observe that $y_i = q(x_i)$

4. Set $\text{sh}_i \leftarrow (x_i, y_i)$, $\text{tk}_i \leftarrow x_i$ and $\text{vk}_i \leftarrow x_i$.
5. Output $(\text{sh}_i, \text{tk}_i, \text{vk}_i)$.

The reconstruction algorithm is the standard Shamir reconstruction algorithm, taking into account the random choices of the x_i s. For simplicity of presentation, we present with shares numbered $1, \dots, t$, and we assume that $x_i \neq x_j$ for every $1 \leq i < j \leq t$.

Rec($\text{sh}_1, \dots, \text{sh}_t$):

1. Parse sh_i as (x_i, y_i) for $i = 1, \dots, t$.
2. Compute $s \leftarrow \sum_{i=1}^t \left(\prod_{k \in [t] \setminus \{i\}} \frac{x_k}{x_k - x_i} \right) \cdot y_i$.
3. Output s .

3.2 Tracing Imperfect Reconstruction Boxes

We now describe our tracing algorithm. Observe, that the informal overview from the previous section inherently assumed that the reconstruction box R is always correct, and in particular, that all evaluations of the polynomial $h(X)$ are correct. To work for imperfect reconstruction boxes, that output the correct reconstructed secret only with a certain non-negligible probability, we need an additional idea. This is because standard polynomial interpolation with erroneous evaluations might fail or give us the wrong polynomial. To resolve this issue, we observe that this problem of interpolating a polynomial of bounded degree from a set of evaluation points with errors is equivalent to the list decoding problem for Reed Solomon codes [66, 43]. Specifically, given a list of evaluations $\{(x'_j, z_j)\}$, the list-decoding algorithms output a list of all polynomials of a certain degree that agree with a pre-determined fraction of the evaluation points. We formally define list decoding of Reed-Solomon codes in Definition 10.

Definition 10 (The RS list decoding problem). *Let \mathbb{F} be a finite field, and let $k, N, C \in \mathbb{N}$ such that $C \leq N < |\mathbb{F}|$. The list decoding problem is defined as follows. Given k, C and N pairs of elements $\{x_i, y_i\}_{i \in [N]} \subseteq \mathbb{F}^2$, output a list H of all univariate polynomials of degree at most k that agree with at least C pairs $\{x_i, y_i\}$. Specifically,*

$$H = \{h \in \mathbb{F}[X] : \deg(h) \leq k \wedge |\{j \in [N] : h(x_j) = y_j\}| \geq C\}.$$

Our scheme uses list decoding for Reed-Solomon codes in a black-box way. For concreteness, we consider the seminal Guruswami-Sudan list decoding algorithm [43] in Theorem 1, but any algorithm that solves the Reed-Solomon list decoding problem defined in Definition 10 in polynomial time may be used instead.

Theorem 1 (Guruswami-Sudan [43]). *Let \mathbb{F} be a finite field, and let $k, N, C \in \mathbb{N}$ such that $C \leq N < |\mathbb{F}|$ and $C \geq \sqrt{kN}$. Then, there exists an algorithm \mathcal{D}_{GS} that solves the list decoding problem as defined in Definition 10 in time polynomial in N, k and $\log |\mathbb{F}|$.*

By the fact that the decoding algorithm runs in polynomial time, the list H that it outputs is also of polynomial length (see also [37]). Denote by $\tau(N, k, \log |\mathbb{F}|)$ the polynomial bounding the list length.⁷

⁷ To solve the list decoding problem with probability 1, the Guruswami-Sudan algorithm runs in expected polynomial time. If we insist that it runs in strict polynomial time, this incurs a negligible error probability. To avoid over-cluttering notation, we will ignore this negligible error in our analysis.

The tracing algorithm. Our full-fledged tracing algorithm obtains a list of evaluations of the polynomial $h(X)$ as described in the previous section. But then, instead of exact polynomial interpolation, it runs a list decoding algorithm to get a list of candidate polynomials. We factor each polynomial in the list one by one. For each polynomial, we check if all its roots belong to the list of true x_i values of all parties. With high probability, only one of the polynomials in the list will satisfy this condition. The roots of *this* polynomial will hence identify our traitors. We formally describe the tracing algorithm in Figure 4. It is parameterized by N and C (the parameters for Guruswami-Sudan decoding), and we will discuss how to set them later in this section.

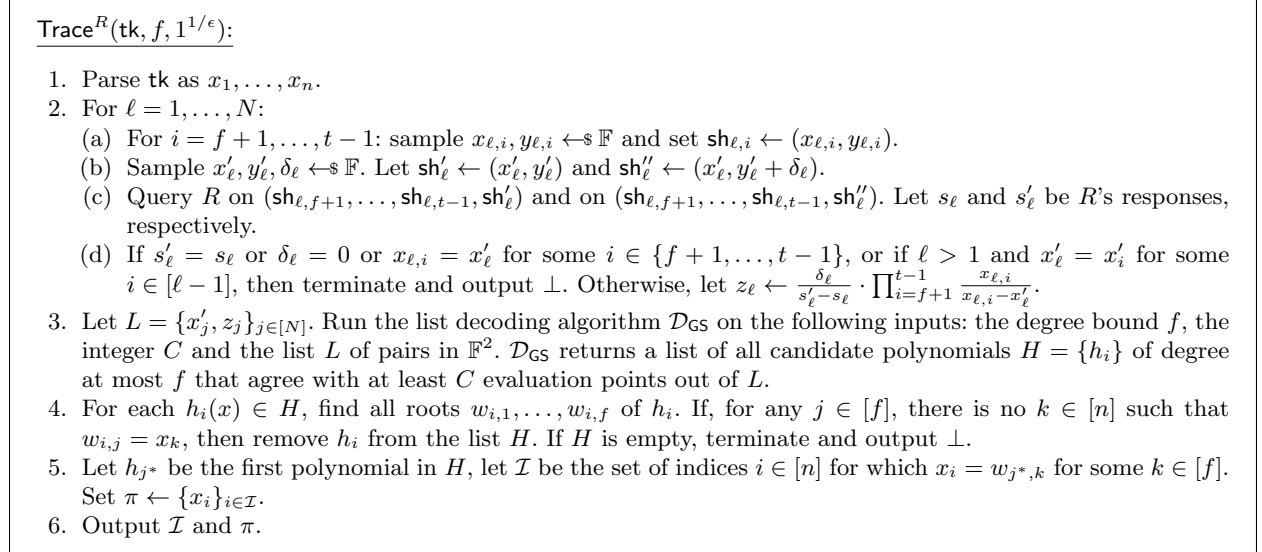


Fig. 4. The tracing algorithm for the Traceable Shamir Secret Sharing scheme TS.

Theorem 2. For every adversary \mathcal{A} , for every $\lambda \in \mathbb{N}$, for every $N, C \in \mathbb{N}$ and $\epsilon \in \left[\sqrt{\frac{2(N+f+t^2)}{p}}, 1 \right]$, such that $\epsilon > \sqrt{2C/N}$ and $\sqrt{fN} \leq C \leq N < p$, it holds that

$$\text{Adv}_{\mathcal{A}, \text{TS}, \epsilon}^{\text{uni-trac}}(\lambda) \leq e^{-\frac{\epsilon^2 N}{2} \cdot (1 - \frac{1}{r})^2} + \frac{f \cdot n \cdot \tau}{p}$$

where $p = |\mathbb{F}|$, $r = \frac{\epsilon^2 N}{2C}$, $n = n(\lambda)$ and $f = f(\lambda)$ are upper bounds on the number of parties and corruptions, respectively, and $\tau = \tau(N, f, \log p)$ is the polynomial upper bounding the length of the output of the Guruswami-Sudan algorithm.

The parameters N and C can be set so that the advantage of the adversary is exponentially small, while meeting the list decoding constraints. Specifically for the Guruswami-Sudan algorithm, C can be set to be $\lceil \sqrt{Nf} \rceil$ and N can be set to $\lceil 16f\lambda/\epsilon^4 \rceil$, where λ is the security paramter. Observe that in this case $r \geq 2$, and hence the term $e^{-\frac{\epsilon^2 N}{2} \cdot (1 - 1/r)^2}$ is at most $e^{-\frac{\epsilon^2 N}{8}}$. Moreover, since $N \geq 8\lambda/\epsilon^2$, we get that the advantage of the adversary is bounded by $e^{-\lambda} + f \cdot n \cdot \tau/p$, which is negligible in λ whenever p is super-polynomial in λ .

Proof (of Theorem 2). Let $\lambda \in \mathbb{N}$, let $\epsilon \in [\sqrt{2(N+f+t^2)/p}, 1]$, and let \mathcal{A} be an adversary taking part in the $\mathbf{ExpUniTrace}_{\mathcal{A}, \text{TS}, \epsilon}(\lambda)$ security experiment. Let \mathbf{G} denote the event in which \mathcal{A} outputs a reconstruction box R that is $(n, t, \mathbf{sh}, \epsilon)$ -good, where n and t are chosen by \mathcal{A} at the beginning of the experiment, and $\mathbf{sh} = (\mathbf{sh}_{i_1}, \dots, \mathbf{sh}_{i_f})$ are the shares given to it by the challenger. Conditioned on $\neg \mathbf{G}$, the output of the experiment is 0 with probability 1, and so we condition the rest of the analysis on \mathbf{G} . Let us denote $\mathbf{sh}_{i_j} = (x_j^*, y_j^*)$ for all $j \in [f]$, and let \mathcal{I} be the set of parties corrupted by \mathcal{A} , i.e. $\{x_j^*\}_{j \in [\mathcal{I}]} = \{x_i\}_{i \in \mathcal{I}}$.

For $\ell = 1, \dots, N$, let $\mathbf{E}_{\delta, \ell}$ be the event $\delta_\ell = 0$. For $\ell \in [N]$ and $j \in \{f+1, \dots, t-1\}$, let $\mathbf{E}_{x, j, \ell, 1}$ be the event that $x_{\ell, j} = x'_\ell$, let $\mathbf{E}_{x, j, \ell, \mathcal{I}}$ be the event that $x_{\ell, j} = x_k^*$ for some $k \in [f]$ and let $\mathbf{E}_{x, j, \ell, 2}$ be the event that $x_{\ell, j} = x_{\ell, i}$ for some $i \in \{f+1, \dots, j-1\}$. For simplicity of notation, let us use $\mathbf{E}_{x, j, \ell}$ to denote the event $\mathbf{E}_{x, j, \ell, 1} \vee \mathbf{E}_{x, j, \ell, \mathcal{I}} \vee \mathbf{E}_{x, j, \ell, 2}$ for all $j \in \{f+1, \dots, t-1\}$ and $\ell \in [N]$. Let \mathbf{E}_ℓ be the event that $x'_\ell = x'_j$ for some $0 < j < \ell$ and let $\mathbf{E}_{\ell, \mathcal{I}}$ be the event that $x'_\ell = x_j^*$ for some $j \in [f]$. Lastly, let $\mathbf{E}_{1, \ell}$ and $\mathbf{E}_{2, \ell}$ be the events that $R(\mathbf{sh}_{\ell, f+1}, \dots, \mathbf{sh}_{\ell, t-1}, \mathbf{sh}'_\ell) = \text{Rec}(\mathbf{sh}_{i_1}, \dots, \mathbf{sh}_{i_f}, \mathbf{sh}_{\ell, f+1}, \dots, \mathbf{sh}_{\ell, t-1}, \mathbf{sh}'_\ell)$ and $R(\mathbf{sh}_{\ell, f+1}, \dots, \mathbf{sh}_{\ell, t-1}, \mathbf{sh}''_\ell) = \text{Rec}(\mathbf{sh}_{i_1}, \dots, \mathbf{sh}_{i_f}, \mathbf{sh}_{\ell, f+1}, \dots, \mathbf{sh}_{\ell, t-1}, \mathbf{sh}''_\ell)$ respectively.

Let \mathcal{J}^* denote the set $\{x_j^*\}_{j \in [f]}$, let $\mathcal{J}_\ell = \{x_{\ell, j}\}_{j \in [f+1, \dots, t-1]}$ and let \mathcal{J}_ℓ^* denote the set $\mathcal{J}^* \cup \mathcal{J}_\ell \cup \{x'_\ell\}$. Then, in the event $\neg \mathbf{E}_{\delta, \ell} \wedge_{j \in [f+1, \dots, t-1]} (\neg \mathbf{E}_{x, j, \ell}) \wedge \neg \mathbf{E}_\ell \wedge \neg \mathbf{E}_{\ell, \mathcal{I}} \wedge \mathbf{E}_{1, \ell} \wedge \mathbf{E}_{2, \ell}$, we have that $|\mathcal{J}_\ell^*| = t$ and

$$s_\ell = \sum_{x_j^* \in \mathcal{J}^*} \left(\prod_{x \in \mathcal{J}_\ell^* \setminus \{x_j^*\}} \frac{x}{x-x_j^*} \right) y_j^* + \sum_{x_{\ell, j} \in \mathcal{J}_\ell} \left(\prod_{x \in \mathcal{J}_\ell^* \setminus \{x_{\ell, j}\}} \frac{x}{x-x_{\ell, j}} \right) y_{\ell, j} \\ + \left(\prod_{x \in \mathcal{J}_\ell^* \setminus \{x'_\ell\}} \frac{x}{x-x'_\ell} \right) y'_\ell$$

and

$$s'_\ell = \sum_{x_j^* \in \mathcal{J}^*} \left(\prod_{x \in \mathcal{J}_\ell^* \setminus \{x_j^*\}} \frac{x}{x-x_j^*} \right) y_j^* + \sum_{x_{\ell, j} \in \mathcal{J}_\ell} \left(\prod_{x \in \mathcal{J}_\ell^* \setminus \{x_{\ell, j}\}} \frac{x}{x-x_{\ell, j}} \right) y_{\ell, j} \\ + \left(\prod_{x \in \mathcal{J}_\ell^* \setminus \{x'_\ell\}} \frac{x}{x-x'_\ell} \right) (y'_\ell + \delta_\ell)$$

This gives us the following:

$$s'_\ell - s_\ell = \delta_\ell \cdot \prod_{x_{\ell, j} \in \mathcal{J}_\ell} \frac{x_{\ell, j}}{x_{\ell, j} - x'_\ell} \cdot \prod_{x_j^* \in \mathcal{J}^*} \frac{x_j^*}{x_j^* - x'_\ell}.$$

Let us define a polynomial $h^*(X) = \prod_{x_j^* \in \mathcal{J}^*} \frac{x_j^* - X}{x_j^*}$. Then, the above equation implies that z_ℓ , as computed by Trace in the ℓ th iteration, is the correct evaluation of $h^*(X)$ at $X = x'_\ell$. Hence, for each $\ell \in [N]$, we get a correct evaluation of $h^*(X)$ at a unique point x'_ℓ , in the event $\neg \mathbf{E}_{\delta, \ell} \wedge_{j \in [f+1, \dots, t-1]} \neg \mathbf{E}_{x, j, \ell} \wedge \neg \mathbf{E}_\ell \wedge \neg \mathbf{E}_{\ell, \mathcal{I}} \wedge \mathbf{E}_{1, \ell} \wedge \mathbf{E}_{2, \ell}$. We will now bound the probability of this event.

We have that $\Pr[\mathbf{E}_{\delta, \ell}] = 1/p$ for all $\ell \in [N]$, since δ_ℓ is uniformly randomly sampled from \mathbb{F} for each ℓ . Next, since $x_{\ell, j}$ is sampled uniformly at random from \mathbb{F} , we have that $\Pr[\mathbf{E}_{x, j, \ell, 1}] = 1/p$, $\Pr[\mathbf{E}_{x, j, \ell, \mathcal{I}}] \leq f/p$ and $\Pr[\mathbf{E}_{x, j, \ell, 2}] \leq (t-f-1)/p$ for all $\ell \in [N]$ and all $j \in \{f+1, \dots, t-1\}$. Hence, we have that $\Pr[\mathbf{E}_{x, j, \ell}] \leq t/p$. Similarly, $\Pr[\mathbf{E}_\ell \vee \mathbf{E}_{\ell, \mathcal{I}}] \leq (N+f)/p$ for all $\ell \in [N]$. Lastly, since \mathbf{R} is $(n, t, \mathbf{sh}, \epsilon)$ -good, we have that, $\Pr[\mathbf{E}_{1, \ell}]$ and $\Pr[\mathbf{E}_{2, \ell}]$ are both at least ϵ . Let us denote the tuple of random variables $(\{(x_{\ell, j}, y_{\ell, j})\}_{j \in [f+1, t-1]}, x'_\ell)$ as $W \in \mathbb{F}^{2(t-f)-1}$. Then,

$$\begin{aligned}
\Pr[\mathbf{E}_{1,\ell} \wedge \mathbf{E}_{2,\ell}] &= \sum_{w \in \mathbb{F}^{2(t-f)-1}} \Pr[W = w] \cdot \Pr_{y'_\ell, \delta_\ell} [\mathbf{E}_{1,\ell} \wedge \mathbf{E}_{2,\ell} | W = w] \\
&= \sum_w \Pr[W = w] \cdot \Pr_{y'_\ell} [\mathbf{E}_{1,\ell} | W = w] \cdot \Pr_{y'_\ell, \delta_\ell} [\mathbf{E}_{2,\ell} | W = w] \tag{4}
\end{aligned}$$

$$\begin{aligned}
&= \mathbb{E}_w \left[\Pr_{y'_\ell} [R(w, y'_\ell) = \text{Rec}(\mathbf{sh}, w, y'_\ell)]^2 \right] \\
&\geq \left(\mathbb{E}_w \left[\Pr_{y'_\ell} [R(w, y'_\ell) = \text{Rec}(\mathbf{sh}, w, y'_\ell)] \right] \right)^2 \tag{5} \\
&= \left(\Pr_{w, y'_\ell} [R(w, y'_\ell) = \text{Rec}(\mathbf{sh}, w, y'_\ell)] \right)^2 \\
&\geq \epsilon^2
\end{aligned}$$

Eq. (4) follows from the fact that the events $\mathbf{E}_{1,\ell}$ and $\mathbf{E}_{2,\ell}$ are independent once they are conditioned on the value of W . Eq. (5) follows from Jensen's inequality.

By combining the above and applying the union bound, we get:

$$\begin{aligned}
&\Pr \left[\neg \mathbf{E}_{\delta,\ell} \bigwedge_{j \in [t-1] \setminus [f]} \neg \mathbf{E}_{x,j,\ell} \wedge \neg \mathbf{E}_\ell \wedge \neg \mathbf{E}_{\ell,\mathcal{I}} \wedge \mathbf{E}_{1,\ell} \wedge \mathbf{E}_{2,\ell} \right] \\
&\geq \Pr[\mathbf{E}_{1,\ell} \wedge \mathbf{E}_{2,\ell}] - \Pr \left[\mathbf{E}_{\delta,\ell} \vee \mathbf{E}_\ell \vee \mathbf{E}_{\ell,\mathcal{I}} \bigvee_{j \in [t-1] \setminus [f]} \mathbf{E}_{x,j,\ell} \right] \\
&\geq \epsilon^2 - 1/p - (N+f)/p - (t-f-1)t/p \\
&\geq \epsilon^2 - (N+f+t^2)/p \\
&\geq \epsilon^2/2
\end{aligned}$$

The last equation follows from our assumption that $\epsilon^2 \geq 2(N+f+t^2)/p$.

Let us define an indicator random variable $Z_\ell = \mathbb{1}[\neg \mathbf{E}_{\delta,\ell} \wedge \bigwedge_{j \in [t-1] \setminus [f]} \neg \mathbf{E}_{x,j,\ell} \wedge \neg \mathbf{E}_\ell \wedge \neg \mathbf{E}_{\ell,\mathcal{I}} \wedge \mathbf{E}_{1,\ell} \wedge \mathbf{E}_{2,\ell}]$, which is 1 if and only if we get a correct evaluation of $h^*(X)$ at a unique point x'_ℓ . Since all the shares are sampled independently for each $\ell \in [N]$, we get that $\Pr[Z_\ell = 1] = E[Z_\ell] \geq \epsilon^2/2$.

Let $Z = \sum_{\ell \in [N]} Z_\ell$. By the Chernoff bound, we have that for every $\eta > 0$,

$$\Pr[Z \leq (1-\eta) \cdot \epsilon^2 N/2] \leq e^{-\frac{\epsilon^2 N \eta^2}{4}}$$

Since $N = (2 \cdot r \cdot C)/\epsilon^2$, we can set $\eta = 1 - 1/r$, to get that $Z > C$ with probability at least $1 - e^{-\frac{\epsilon^2 N}{2} \cdot (1 - \frac{1}{r})^2}$. Hence, with this probability, the list decoding algorithm \mathcal{D}_{GS} will be able to output all polynomials that agree with at least C evaluations out of $\{x'_\ell, z_\ell\}_{\ell \in [N]}$, including the polynomial $h^*(X)$. Next, since all the roots of $h^*(X)$ correspond to x_i values of the parties in \mathcal{I} , h^* will not be eliminated from H in Step 4.

Additionally, we claim that with high probability, there will be no other polynomial in the list H after Step 4. To prove this, let us first define an event $\mathbf{E}_{h,i}$ for all honest parties $i \in [n] \setminus \mathcal{I}$, denoting

that x_i is a root of some polynomial in the list H . Next, observe that the shares of honest parties, i.e. $\{x_i, y_i\}_{i \notin \mathcal{I}}$ are statistically independent from both the view of \mathcal{A} and the shares sampled by the Trace algorithm. Hence, the probability that x_i is a root of any polynomial h_j in H is bounded by f/p , since the degree of h_j is bounded by f for all $h_j \in H$. Since H has upto τ polynomials, we get that $\Pr[\mathbf{E}_{h,i}] \leq f \cdot \tau/p$.

Then, by applying a union bound over all honest x_i , we get that

$$\Pr \left[\bigvee_{i \in [n] \setminus \mathcal{I}} \mathbf{E}_i \right] \leq \frac{f \cdot (n - f) \cdot \tau}{p}$$

Lastly, observe that \mathcal{A} can win the game only if (a) $Z < C$ so that the list decoding algorithm fails to find h^* or (b) if any of the events $\mathbf{E}_{h,i}$ occur, causing an honest party to be blamed. So we get that,

$$\text{Adv}_{\mathcal{A}, \text{TS}, \epsilon}^{\text{uni-trac}}(\lambda) \leq e^{-\frac{\epsilon^2 N}{2} \cdot (1 - \frac{1}{r})^2} + \frac{f \cdot (n - f) \cdot \tau}{p}$$

This proves the theorem. □

Learning f . If the number f of corruptions is not known in advance by the tracing algorithm, it can learn it by simply trying $f = t - 1, t - 2, \dots$ until it reaches a value of f that works; that is, a value of f for which the above tracing algorithm indeed finds exactly f corrupted x_i values. Suppose that the real number of corruptions is f^* . For each value $f > f^*$ that Trace tries, outputting a subset \mathcal{I} of size f means outputting at least one honest party. By the analysis in the proof of Theorem 2, the probability that it outputs such a subset is at most $f \cdot (n - f) \cdot \tau/p \leq n^2 \tau/p$. Moreover, when Trace tries $f = f^*$, then Theorem 2 tells us that it will fail to output the correct subset with probability at most $e^{-\lambda} + n^2 \tau/p$ (for the choices of parameters discussed above). Hence, by a union bound, the probability that Trace correctly traces R back the corrupted subset is at least $1 - (e^{-\lambda} + 2n^3 \tau/p)$.

3.3 Adding Non-Imputability

We now present the changes needed to make our Traceable Shamir scheme TS satisfy non-imputability against a malicious tracer. This yields our full-fledged traceable secret sharing scheme based on Shamir secret sharing. In the simplified scheme presented in the previous sections, the fact that tk contains x_1, \dots, x_n explicitly, allows a malicious tracer to falsely accuse any party i by including their x_i as part of the proof π . To remedy this situation, we need to a-priori hide the x_i s, but in a way that still allows the tracer (and, later on, the verifier) to link an x_i value, once it has been extracted from a reconstruction box R , back to party i . Thus, instead of including x_1, \dots, x_n in the clear in tk and vk , we include $F(x_1), \dots, F(x_n)$ where F is a one-way function. The proof would still consist of the x_i s of the corrupted parties. This way, intuitively speaking, falsely accusing an honest party i amounts to inverting $F(x_i)$ for a randomly chosen x_i . Specifically, the changes to our scheme are as follows:

1. The share algorithm $\text{Share}(1^\lambda, s, n, t, \rho)$ now also computes $u_i \leftarrow F(x_i)$. The tracing key component tk_i and verification key component vk_i are set to be $\text{tk}_i = \text{vk}_i = u_i$ (instead of x_i).

2. The tracing algorithm `Trace` computes the list $H = \{h_j\}_j$ of polynomials as before. For each j , `Trace` now checks if h_j is the correct polynomial by factoring it to find its roots, w_1, \dots, w_f . It then computes $u'_i \leftarrow F(w_i)$ for every $i \in [f]$, and looks for u'_i in `tk`. If for some index $i \in [f]$, u'_i does not appear in the tracing key `tk`, then `Trace` eliminates h_j from the list H . Suppose that at the end of this process, there is at least one polynomial left in H . Let h^* be the first polynomial in H and let w_1^*, \dots, w_f^* be its roots. For each $i \in [f]$, `Trace` finds the index $k \in [n]$ for which $F(w_i^*) = u_k$, and adds k to the subset \mathcal{I} of corrupted parties. The proof π is set to be (w_1^*, \dots, w_f^*) .
3. To verify a proof $\pi = \{w_1, \dots, w_f\}$ against a subset $\mathcal{I} = \{i_1, \dots, i_f\}$, the verification algorithm `Verify` checks that $F(w_j) = u_{i_j}$ for $j \in [f]$, where $\text{vk} = (u_1, \dots, u_n)$.

We denote the resulting scheme by NITS (for “Non-Imputable Traceable Shamir”). The following theorem establishes its non-imputability.

Theorem 3. *For every adversary \mathcal{A} for the non-imputability of NITS, there exists an algorithm \mathcal{B} such that for every $\lambda \in \mathbb{N}$, it holds that*

$$\text{Adv}_{\mathcal{A}, \text{NITS}}^{\text{ni}}(\lambda) = \Pr [F(\mathcal{B}(F(x^*))) = F(x^*)],$$

where $x^* \leftarrow_{\$} \mathbb{F}$ and the probability is also over the random coins of \mathcal{B} .

Proof. Let \mathcal{A} be as in the statement of the theorem and consider the following adversary \mathcal{B} trying to invert F . The inverter \mathcal{B} gets $u^* = F(x^*)$ as input for a uniformly random $x^* \leftarrow_{\$} \mathbb{F}_p$. It then simulates the non-imputability experiment $\text{ExpNI}_{\mathcal{A}, \text{NITS}}(\lambda)$ to \mathcal{A} :

1. \mathcal{B} invokes $\mathcal{A}(\lambda)$ and gets n, t, i^* and the secret s from \mathcal{A} .
2. \mathcal{B} samples a correlation string $\rho = (a_1, \dots, a_{t-1}) \leftarrow_{\$} \mathbb{F}^{t-1}$. It defines the polynomial $q(X) = s + \sum_{i=1}^{t-1} a_i X^i$.
3. For each $i \in [n] \setminus \{i^*\}$, \mathcal{B} samples $x_i \leftarrow_{\$} \mathbb{F}$.
4. \mathcal{B} then generates the secret shares, tracing key, and verification key to pass to \mathcal{A} :
 - (a) To generate secret shares for every $i \in [n] \setminus \{i^*\}$, \mathcal{B} computes $y_i \leftarrow q(x_i)$ and sets $\text{sh}_i = (x_i, y_i)$. Note that \mathcal{B} does not need to pass to \mathcal{A} the secret share of party i^* .
 - (b) To generate the tracing key `tk` and verification key:
 - For every $i \in [n] \setminus \{i^*\}$, \mathcal{B} computes $u_i \leftarrow F(x_i)$ and sets $\text{tk}_i = \text{vk}_i = u_i$.
 - For party i^* , \mathcal{B} sets $\text{tk}_{i^*} = \text{vk}_{i^*} = u^*$.
 The tracing key is $\text{tk} = (\text{tk}_1, \dots, \text{tk}_n)$ and the verification key is $\text{vk} = (\text{vk}_1, \dots, \text{vk}_n)$.
5. \mathcal{A} then outputs a subset \mathcal{I}^* and a proof π . If $i^* \in \mathcal{I}^*$ and π is accepting, then it must include a field element x' such that $F(x') = u^*$. \mathcal{B} thus outputs this x' . If \mathcal{A} does not output \mathcal{I}^* and π satisfying these conditions, \mathcal{B} aborts.

Observe that \mathcal{B} perfectly simulates $\text{ExpNI}_{\mathcal{A}, \text{NITS}}(\lambda)$ to \mathcal{A} . Moreover, \mathcal{B} outputs a preimage of u^* if and only if \mathcal{A} outputs a subset that contains i^* and an accepting proof. Hence,

$$\text{Adv}_{\mathcal{A}, \text{NITS}}^{\text{ni}}(\lambda) = \Pr [F(\mathcal{B}(F(x^*))) = F(x^*)],$$

and the theorem follows. □

Sampling random evaluation points. The evaluation points x_i for each party i can be sampled from a subset of the field \mathbb{F} , as long as it is large enough to accommodate tracing security and non-imputability. Specifically, the Trace algorithm produces a set of “candidate polynomials” via the list-decoding algorithm. To argue that none of these polynomials have the x value of an uncorrupted party as a root, we use the fact that the values are chosen from a super-polynomial-sized set. Similarly, for non-imputability, we need the one-way function evaluation $F(x_i)$ to hide x_i , meaning that the evaluation points must be drawn from a super-polynomial-sized set.

4 A Scheme Based on Blakley Secret Sharing

In this section, we show how to trace leakage in a variant of Blakley’s seminal secret sharing scheme [6]. We begin by putting forth an extension of Blakley’s original scheme. We then present a very efficient tracing algorithm for this extended scheme.

4.1 An Extended Blakley Scheme

We begin by recalling Blakley’s original scheme [6], and then present our extended scheme.

Blakley’s original scheme and its limitations. In Blakley’s scheme, the secret is encoded as one coordinate of a point in a field \mathbb{F}_p^t , where p is an appropriately chosen modulus, and t is the threshold. Concretely, the dealer chooses a random point $\mathbf{x} \in \mathbb{F}_p^t$ conditioned on its first coordinate x_1 being equal to the secret $s \in \mathbb{F}_p$. The secret of party i is then a uniformly random hyperplane $H_i \subset \mathbb{F}_p^t$ that passes through \mathbf{x} . For a large enough modulus p , a collection of t dealt hyperplanes intersect in a single point with high probability. Since by design, all hyperplanes pass through \mathbf{x} , this unique intersection point must be \mathbf{x} .

As for secrecy, observe that the intersection of any k hyperplanes is a subspace of dimension at least $t - k$. Hence, if $k < t$, this subspace is of dimension at least 1. Suppose for simplicity that it is of dimension exactly 1; that is, it is a line ℓ in \mathbb{F}_p^t . If ℓ is not perpendicular to the axis of the first coordinate (which happens with very small probability), then all points on ℓ have different first coordinates. Intuitively, this means that the line reveals no information about what is the first coordinate of the secret point \mathbf{x} , and hence no information about the secret $s \in \mathbb{F}_p$. We will make everything precise when we present our extension of Blakley’s scheme.

For our tracing procedure, however, we will need to encode the secret as a full point in \mathbb{F}_p^t . A naive extension of Blakley’s scheme might set the secret $\mathbf{s} \in \mathbb{F}_p^t$ as the intersection point of all hyperplanes. This naive extension is, however, completely insecure! Observe that even one hyperplane – that is, a single secret share – reveals much information about the secret, as it contains only a $1/p$ -fraction of the points in \mathbb{F}_p^t . More generally speaking, it is not hard to see that if we set the secret to be the first k coordinates of the intersection point $\mathbf{x} \in \mathbb{F}_p^t$, then already $t - k + 1$ hyperplanes reveal information about the secret.

Our extended Blakley scheme. Instead of the naive approach above, we consider a more delicate way of “smearing” the secret across multiple entries of the intersection point \mathbf{x} , by introducing more randomness to the secret shares. Suppose that we want to encode the secret as a full point in \mathbb{F}_p^t . The idea is to randomly split \mathbf{s} to t shares $\mathbf{r}_1, \dots, \mathbf{r}_t$ via standard additive secret sharing; that is, $\mathbf{r}_1, \dots, \mathbf{r}_t$ are uniformly random in \mathbb{F}_p^t subject to $\mathbf{r}_1 + \dots + \mathbf{r}_t = \mathbf{s}$. Then, we share each of these \mathbf{r}_i s using the naive extension of Blakley’s scheme described above. Meaning, the i th secret share now includes t random hyperplanes $H_i^{(1)}, \dots, H_i^{(t)}$ that pass through $\mathbf{r}_1, \dots, \mathbf{r}_t$, respectively. Per the

security of this scheme, envision the t instances of “naive Blakley” as taking place in t orthogonal subspaces of $\mathbb{F}_p^{t^2}$. Then, any subset of $t - 1$ parties holds $t \cdot (t - 1)$ hyperplanes in $\mathbb{F}_p^{t^2}$. This allows us to hide exactly t elements of \mathbb{F}_p . Intuitively, this is since each “candidate secret” $\mathbf{s}^* \in \mathbb{F}_p^t$ adds t affine equations over $\mathbb{F}_p^{t^2}$, and so the affine system of equations induced by this secret and the $t \cdot (t - 1)$ hyperplanes contains exactly t^2 equations. Hence, it is not overdetermined, making \mathbf{s}^* plausible. The formal analysis requires much more care. We now present the scheme in detail.

The scheme, which we denote by Blakley^+ , is parameterized by a function $p = p(n, t, \lambda)$, determining the size of the finite field. The secret sharing algorithm takes in n, t , a secret vector $\mathbf{s} \in \mathbb{F}_p^t$, and a correlation string ρ that is interpreted as the randomness for the additive secret sharing of \mathbf{s} .

Share($1^\lambda, n, t, \mathbf{s}, \rho$) :

1. Parse ρ as $\mathbf{r}_2, \dots, \mathbf{r}_t \in \mathbb{F}_p^t$.
2. Set $\mathbf{r}_1 \leftarrow \mathbf{s} - (\mathbf{r}_2 + \dots + \mathbf{r}_t)$.
3. For $j = 1, \dots, t$: sample $\mathbf{a}_j \leftarrow \$_\mathbb{F}_p^t$ and set $b_j \leftarrow \langle \mathbf{a}_j, \mathbf{r}_j \rangle$.
4. Set the share to be $\text{sh} = \{(\mathbf{a}_j, b_j)\}_{j \in [t]}$.
5. Output sh .

The reconstruction algorithm takes in t shares and outputs a secret $\mathbf{s} \in \mathbb{F}_p^k$. For simplicity of presentation, we number the shares it takes in by $1, \dots, t$, but the algorithm is defined exactly the same for every subset of t shares.

Rec($\text{sh}_1, \dots, \text{sh}_t$):

1. Parse sh_i as $\{(\mathbf{a}_{i,j}, b_{i,j})\}_{j \in [t]}$ for every $i \in [t]$.
2. For $j = 1, \dots, t$: Consider the system of affine equations $\{\langle \mathbf{a}_{i,j}, \mathbf{R}_j \rangle = b_{i,j}\}_{i \in [t]}$ in the indeterminates $\mathbf{R}_j = (R_{j,1}, \dots, R_{j,t})$. This is a system with t equations in t variables. If the equations are linearly independent, compute its unique solution \mathbf{r}_j . Otherwise, if there is more than one solution to the system, terminate and output \perp .
3. Compute $\mathbf{s} \leftarrow \mathbf{r}_1 + \dots + \mathbf{r}_t$.
4. Output \mathbf{s} .

Correctness. Let $n \in \mathbb{N}$ and $t \leq n$, and let $\mathbf{s} \in \mathbb{F}_p^t$ be the secret. Let $\text{sh}_1, \dots, \text{sh}_t$ be any t shares out of the n shares generated by **Share**. We show that with high probability, $\text{Rec}(\text{sh}_1, \dots, \text{sh}_t) = \mathbf{s}$. Fix $j \in [t]$, and consider the system of equations $\{\langle \mathbf{a}_{i,j}, \mathbf{R}_j \rangle = b_{i,j}\}_{i \in [t]}$. If $\{\mathbf{a}_{i,j}\}_{i \in [t]}$ are linearly independent, then this system has a unique solution that must be equal to the share \mathbf{r}_j sampled by **Share**. The probability that $\{\mathbf{a}_{i,j}\}_{i \in [t]}$ are linearly independent is

$$\frac{\prod_{\ell=0}^{t-1} (p^t - p^\ell)}{p^{t^2}} = \prod_{\ell=1}^t (1 - p^{-\ell}) \geq 1 - \sum_{\ell=1}^t p^{-\ell} > 1 - \frac{1}{p-1}.$$

By union bound, the probability that **Rec** computes $\mathbf{r}_1, \dots, \mathbf{r}_k$ correctly is at least $1 - t/(p-1)$. Moreover, whenever $\mathbf{r}_1, \dots, \mathbf{r}_t$ are computed correctly, so is the secret \mathbf{s} . All in all, we get that the correctness error of our scheme can be bounded by $t/(p-1)$, which is negligible whenever t is polynomial and p is super-polynomial.

Succinctly representing hyperplanes. Consider the share dealt to party i in our scheme. It consists of t pairs $\{(\mathbf{a}_{i,j}, b_{i,j})\}_{j \in [t]}$, and hence its size is $t \cdot (t+1) \cdot \log p$. However, observe that for each $j \in [t]$, $\mathbf{a}_{i,j}$ a uniformly random vector in \mathbb{F}_p^t , independent of $\{\mathbf{a}_{i,j'}\}_{j' \neq j}$. Hence, instead of representing the $\mathbf{a}_{i,j}$ s explicitly, we can replace them by a seed σ_i to a pseudorandom generator G . Now, to reconstruct the secret, one first expands σ_i to the $\mathbf{a}_{i,j}$ s by computing $(\mathbf{a}_{i,1}, \dots, \mathbf{a}_{i,t}) \leftarrow G(\sigma_i)$ and then proceeds as before. This reduces the share size to $\lambda + t \cdot \log p$ (since the $b_{i,j}$ s still need to be represented explicitly), at the expense of having only computational secrecy. Another option is to succinctly represent $(\mathbf{a}_{i,1}, \dots, \mathbf{a}_{i,t})$ using ϵ -biased sets [53, 1, 67]. This approach can reduce the share size to $\log(t^2 \log p) + O(\alpha) + k \cdot \log p$, while degrading statistical security by an additive $n \cdot 2^{-\alpha}$ factor. Finally, as observed by Brikell [16], all the $\mathbf{a}_{i,j}$ s (or the succinct representations thereof) can be made public without affecting security.⁸ If this is done, then the secret information each party needs to store is only of size $t \log p$, which exactly matches the size of the secret, making our Extended Blakley scheme ideal (meaning, secret shares are the same size as the secret).⁹

Secrecy. We now analyze the security of our extended scheme. We prove the following theorem.

Theorem 4. *For every $p = p(n, t, \lambda)$, the scheme Blakley⁺ described above is ϵ -secret for $\epsilon = (t+1)/(p-1)$.*

Proof. We prove that with high probability over the shares generation process, the shares of parties $1, \dots, t-1$ reveal no information about the secret $\mathbf{s} \in \mathbb{F}_p^k$. The proof is identical for every other subset $\mathcal{J} \subset [n]$ of size t .

Let $\{(\mathbf{a}_{i,j}, b_{i,j})\}_{j \in [t]}$ denote the share of party i for $i = 1, \dots, t-1$. For $i \in [t-1]$ and $j \in [t]$, let $\boldsymbol{\mu}_{i,j}^T = (0^{t \times (j-1)}, \mathbf{a}_{i,j}^T, 0^{t \times (t-j)}) \in \mathbb{F}_p^{t^2}$. That is, $\boldsymbol{\mu}_{i,j}$ is a vector of length t^2 composed of t blocks of length t each; the j th block is $\mathbf{a}_{i,j}$ and all other blocks are all 0s. For $\ell \in [t]$, let $\boldsymbol{\psi}_\ell^T = (\mathbf{e}_\ell^T, \mathbf{e}_\ell^T, \dots, \mathbf{e}_\ell^T) \in \mathbb{F}_p^{t^2}$, where $\mathbf{e}_\ell \in \mathbb{F}_p^t$ is the ℓ th standard basis vector for \mathbb{F}_p^t . That is, $\boldsymbol{\psi}_\ell$ is a vector of length t^2 , where each entry $m \in [t^2]$ is 1 if $m = \ell \pmod t$, and is 0 otherwise. Let $A_j = \{\boldsymbol{\mu}_{i,j}\}_{i \in [t-1]}$ for $j \in [t]$, let $A = \bigcup_{j \in [t]} A_j$, let $B = \{\boldsymbol{\psi}_\ell\}_{\ell \in [t]}$, and let $C = A \cup B$. Note that C contains at most t^2 vectors in $\mathbb{F}_p^{t^2}$.

The proof proceeds in two steps. First, we argue that if C spans $\mathbb{F}_p^{t^2}$ then every secret $\mathbf{s} \in \mathbb{F}_p^t$ has a unique collection of additive shares $\mathbf{r}_1, \dots, \mathbf{r}_t \in \mathbb{F}_p^t$ such that $\mathbf{s} = \mathbf{r}_1 + \dots + \mathbf{r}_t$, and these additive shares form a solution to the set of linear equations induced by the $t-1$ secret shares. We deduce that in this case, the shares reveal no information about the secret \mathbf{s} . Secondly, we bound the probability that C spans $\mathbb{F}_p^{t^2}$.

Suppose C spans $\mathbb{F}_p^{t^2}$, and let $\mathbf{s} = (s_1, \dots, s_t) \in \mathbb{F}_p^t$ be some secret. Consider the following system of (affine) linear equations

$$\begin{cases} \langle \mathbf{a}_{i,j}, \mathbf{R}_j \rangle = b_{i,j} \text{ for } j \in [t] \text{ and } i \in [t-1] \\ R_{1,\ell} + \dots + R_{t,\ell} = s_\ell \text{ for } \ell \in [t] \end{cases}$$

where $\mathbf{R}_j = (R_{j,1}, \dots, R_{j,t})$. This is a system with t^2 equations in t^2 variables $\{R_{j,\ell}\}_{j \in [t], \ell \in [t]}$. Observe that this system can be expressed in matrix form as

$$M\mathbf{z} = \mathbf{b},$$

⁸ Brikell [16] made this observation with respect to Blakley's original scheme, but it equally applies to our scheme.

⁹ In our full-fledged scheme with non-immutability, it is not possible to publish all $\mathbf{a}_{i,j}$ s in the clear to obtain an ideal secret sharing scheme. However, for secrets in $\{0, 1\}^\lambda$, our derandomization approach gives shares of size 2λ , coming close to it.

where M is a $t^2 \times t^2$ matrix whose rows are the vectors in the set C , $\mathbf{z} = (\mathbf{R}_1, \dots, \mathbf{R}_t) \in \mathbb{F}_p^{t^2}$ is the vector of indeterminates, and $\mathbf{b} \in \mathbb{F}_p^{t^2}$ is the solution vector that depends on $\{b_{i,j}\}_{i \in [t-1], j \in [t]}$ and the secret \mathbf{s} . Since we assumed C spans $\mathbb{F}_p^{t^2}$ and M is a square matrix, it follows that M is full rank. Hence, the system has a unique solution. Since this holds for every secret \mathbf{s} , it follows that conditioned on C spanning $\mathbb{F}_p^{t^2}$, the distribution over the $t-1$ shares is identical for any two shares $\mathbf{s}, \mathbf{s}' \in \mathbb{F}_p^t$.

Lemma 2. *The set C is linearly independent with probability at least $1 - (t+1)/(p-1)$, where the probability is over the choice of vectors in A .*

Proof (Lemma 2). We prove the lemma in two steps. First, we lower bound the probability that each A_j is linearly independent. Then, we upper bound the probability that there is a linear combination of the vectors in C that assigns non-zero coefficients to any of the vectors in B and evaluates to $\mathbf{0}$.

First, observe that for each $j \in [t]$, A_j is linearly independent if and only if the set of vectors $\{\mathbf{a}_{1,j}, \dots, \mathbf{a}_{t-1,j}\}$ is linearly independent over \mathbb{F}_p^t . This probability is at least $1 - 1/(p-1)$, as proven in our correctness analysis. By a union bound, the probability that each A_j is linearly independent is at least $1 - t/(p-1)$. Observe that conditioned on each A_j being linearly independent, any linear combination of the vectors in C that assigns zero coefficient to all the vectors in B **cannot** give the zero vector $\mathbf{0} \in \mathbb{F}_p^{t^2}$.

Now consider a linear combination of the vectors in C that assigns a non-zero coefficient to at least one vector in B . Such a linear combination can be written as

$$\mathbf{v} = \sum_{j \in [t]} \alpha_j \boldsymbol{\psi}_j + \sum_{j \in [t]} \sum_{i \in [t-1]} \beta_{i,j} \boldsymbol{\mu}_{i,j},$$

where at least one α_j is non-zero. By the structure of $\{\boldsymbol{\psi}_1, \dots, \boldsymbol{\psi}_t\}$, it holds that

$$\sum_{j \in [t]} \alpha_j \boldsymbol{\psi}_j = (\boldsymbol{\eta}, \boldsymbol{\eta}, \dots, \boldsymbol{\eta}),$$

for some non-zero vector $\boldsymbol{\eta} \in \mathbb{F}_p^t$. Now, if $\mathbf{v} = \mathbf{0}$, this means that $\boldsymbol{\eta}$ is in the linear span of $\{\mathbf{a}_{1,j}, \dots, \mathbf{a}_{t-1,j}\}$ for **every** index $j \in [t]$. In other words $\boldsymbol{\eta}$ is in the intersection of the subspaces spanned by $\{\mathbf{a}_{1,j}, \dots, \mathbf{a}_{t-1,j}\}$ for $j \in [t]$. That is, if we denote

$$V = \bigcap_{j \in [t]} \text{Span}(\{\mathbf{a}_{1,j}, \dots, \mathbf{a}_{t-1,j}\}),$$

then it holds that $\boldsymbol{\eta} \in V$. If $V = \{\mathbf{0}\}$, then we have arrived at a contradiction, and it must be the case that $\boldsymbol{\eta} = \mathbf{0}$. Note that for every $j \in [t]$, the dimension of $\text{Span}(\{\mathbf{a}_{1,j}, \dots, \mathbf{a}_{t-1,j}\})$ is at most $t-1$. Hence, the orthogonal subspace is of dimension at least 1; let w_j be a uniformly random vector in this subspace. Observe that if w_1, \dots, w_t span \mathbb{F}_p^t then V is trivial. Since w_1, \dots, w_t are independent and uniform in \mathbb{F}_p^t , the same analysis as before shows that they span \mathbb{F}_p^t with probability at least $1 - 1/(p-1)$.

We have shown that:

1. If each A_j is linearly independent, there can be no linear combination of C that yields $\mathbf{0}$ and assigns all zero coefficients to B . The probability that at each A_j is linearly independent is $1 - t/(p-1)$.

2. If the subspace V is trivial, there can be no linear combination of C that yields $\mathbf{0}$ and assigns a non-zero coefficient to a vector in B . This occurs with probability at least $1 - 1/(p - 1)$.

By a union bound, we obtain that the probability that C is linearly independent is at least $1 - (t + 1)/(p - 1)$. \square

We proved that C is linearly independent with probability at least $1 - (t + 1)/(p - 1)$, and that conditioned on C being linearly independent, the distributions over any tuple of $t - 1$ shares are identical for any two secrets \mathbf{s} and \mathbf{s}' . Hence, the statistical distance between the two distributions is at most $(t + 1)/(p - 1)$, concluding the proof of the theorem. \square

4.2 The Basic Tracing Procedure

We now present our tracing algorithm for the extended Blakeley scheme described above. We begin by presenting a tracing algorithm that relies on the full knowledge of all secret shares $\text{sh}_1, \dots, \text{sh}_n$ dealt to the parties. In the following section, we show how to (slightly) modify the tracing key and the tracing algorithm so that tracing does not require explicit knowledge of the secret shares, and the scheme can thus provide non-imputability.

We begin with an informal overview of the construction. In Blakeley⁺, the secret is a point $\mathbf{s} \in \mathbb{F}_p^t$. To derive the n shares, the Share algorithm samples $\mathbf{r}_1, \dots, \mathbf{r}_t$ that sum up to \mathbf{s} , and for $i \in [n]$, the i th share consists of t random hyperplanes $H_i^{(1)}, \dots, H_i^{(t)}$, passing through $\mathbf{r}_1, \dots, \mathbf{r}_t$, respectively. The tracing key tk will consist of $\mathbf{r}_2, \dots, \mathbf{r}_t$, and in addition, the first hyperplane of each party, $H_1^{(1)}, \dots, H_n^{(1)}$. The verification key vk consists of $H_i^{(1)}, \dots, H_i^{(t)}$. For the sake of simplicity, suppose in this informal overview that we know that $t - 1$ secret shares are hardcoded to the reconstruction box R , which takes in one additional share as input. To do that, we sample a random share to feed into R , consistently with $\mathbf{r}_1, \dots, \mathbf{r}_t$. That is, for every $j \in [t]$, we sample a hyperplane $H^{(j)}$ as follows:

1. For $j = 2, \dots, t$: Sample a hyperplane $H^{(j)}$ uniformly at random, conditioned on passing through \mathbf{r}_j .
2. For $j = 1$, sample $H^{(1)}$ uniformly at random.

The tracing algorithm then feeds the share $\text{sh} = (H^{(1)}, H^{(2)}, \dots, H^{(t)})$ to R , which returns some secret $\mathbf{s}' \in \mathbb{F}_p^t$. Suppose for simplicity that R is perfectly correct. Then, by construction, we know that for each $j \geq 2$, $H^{(j)}$ intersects with the j th hyperplanes of the shares embedded in R at \mathbf{r}_j . This means that $H^{(1)}$ intersects the first hyperplanes of the shares embedded in R at $\mathbf{r}'_1 = \mathbf{s}' - (\mathbf{r}_2 + \dots + \mathbf{r}_t)$. The tracing algorithm thus outputs all parties $i \in [n]$ for which \mathbf{r}'_1 is on their first hyperplane $H_i^{(1)}$. Since R is correct, for every party i whose share is hardcoded to R , \mathbf{r}'_1 will indeed be on $H_i^{(1)}$, and so all corrupted parties are caught by our tracing algorithm. On the other hand, if party i is uncorrupted, then \mathbf{r}'_1 will be statistically independent from $H_i^{(1)}$, and hence will lie on it with very small probability, so no honest parties are wrongly accused.

We now present the tracing algorithm in detail. For simplicity of presentation, we assume that the number of corruptions $f < t$ hardcoded into the reconstruction box R is given as input to Trace. This assumption may be removed in the same manner as in Section 3.

Trace^R(tk, f, 1^{1/ε}):

1. Parse tk as $(\mathbf{r}_2, \dots, \mathbf{r}_k, (\mathbf{a}_{1,1}, b_{1,1}), \dots, (\mathbf{a}_{n,1}, b_{n,1}))$.
2. For $\ell = 1, \dots, 2\lambda/\epsilon$:
 - (a) Sample $t - f$ shares to feed to R . For $i = 1, \dots, t - f$:
 - i. For $j = 2, \dots, t$: sample $\mathbf{a}'_{i,j} \leftarrow_{\$} \mathbb{F}_p^t$ and set $b'_{i,j} \leftarrow \langle \mathbf{a}'_{i,j}, \mathbf{r}_j \rangle$.
// sample a random hyperplane that passes through \mathbf{r}_j
 - ii. For $j = 1$: $\mathbf{a}'_{i,1} \leftarrow_{\$} \mathbb{F}_p^t$ and $b'_{i,1} \leftarrow_{\$} \mathbb{F}_p$.
// sample a uniformly random hyperplane
 - iii. Set $\text{sh}'_i = \left\{ (\mathbf{a}'_{i,j}, b'_{i,j}) \right\}_{j \in [t]}$.
 - (b) Query R on $(\text{sh}'_1, \dots, \text{sh}'_{t-f})$ and get back a secret $\mathbf{s} \in \mathbb{F}_p^t$.
 - (c) Compute $\mathbf{r}_1 \leftarrow \mathbf{s} - (\mathbf{r}_2 + \dots + \mathbf{r}_k)$.
 - (d) Compute the subset $\mathcal{I} \subseteq [n]$ as the set of indices $i \in [n]$ for which $\langle \mathbf{a}_{i,1}, \mathbf{r}_1 \rangle = b_{i,1}$. Set $\pi \leftarrow \{(\mathbf{a}_{i,1}, b_{i,1})\}_{i \in \mathcal{I}}$.
 - (e) If $|\mathcal{I}| = f$, output \mathcal{I} and π and terminate.
3. If reached, output \emptyset .

The verification algorithm `Verify` takes in a verification key $\text{vk} = (\mathbf{a}_{1,1}, b_{1,1}), \dots, (\mathbf{a}_{n,1}, b_{n,1})$, a subset \mathcal{I} and a proof $\pi = \{(\mathbf{a}'_{i,1}, b'_{i,1})\}_{i \in \mathcal{I}}$ and outputs 1 if and only if \mathcal{I} and π are consistent with vk ; that is $(\mathbf{a}'_{i,1}, b'_{i,1}) = (\mathbf{a}_{i,1}, b_{i,1})$ for every $i \in \mathcal{I}$. This makes it trivial for the tracer to frame innocent parties, but as mentioned, we will show how to address this fact in the following section.

Let `TB` = (`Share`, `Rec`, `Trace`, `Verify`) (`TB` for “Traceable Blakley”) be the traceable secret sharing scheme described above. That is, `Share` and `Rec` are as defined in Section 4.1, when `Share` additionally outputs `tk` and `vk` as described above. The algorithms `Trace` and `Verify` are as described in this section.

The traceability of `TB` is established by the following theorem. It proves that `TB` is universally traceable with respect to the partition of the correlation strings space into singletons; that is $\Gamma_{\text{sngltn}} = \{\{(\mathbf{r}_2, \dots, \mathbf{r}_t)\} : \mathbf{r}_2, \dots, \mathbf{r}_t \in \mathbb{F}_p^t\}$. Standard traceability then follows by Lemma 1.

Theorem 5. *For every adversary \mathcal{A} and for every $\lambda \in \mathbb{N}$, and every $\epsilon \in [0, 1]$ it holds that*

$$\text{Adv}_{\mathcal{A}, \text{TB}, \Gamma_{\text{sngltn}}, \epsilon}^{\text{uni-trac}}(\lambda) \leq e^{-\lambda} + \frac{2n\lambda}{\epsilon p},$$

where $n = n(\lambda)$ is an upper bound on the number of parties, and assuming $p \geq 2n/\epsilon$.

Proof. Let $\lambda \in \mathbb{N}$, let $\epsilon \in [0, 1]$, and let \mathcal{A} be an adversary taking part in the $\text{ExpUniTrace}_{\mathcal{A}, \text{TB}, \Gamma_{\text{sngltn}}, \epsilon}(\lambda)$ security experiment. Let G denote the event in which \mathcal{A} outputs a reconstruction box R that is $(n, t, \mathbf{sh}, \Gamma_{\text{sngltn}}, \epsilon)$ -good, where n and t are chosen by \mathcal{A} at the beginning of the experiment, and $\mathbf{sh} = (\text{sh}_{i_1}, \dots, \text{sh}_{i_f})$ are the shares given to it by the challenger. Conditioned on $\neg \text{G}$, the output of the experiment is 0 with probability 1, and so we condition the rest of the analysis on G .

Note that in each iteration of Step 2, the tracing algorithm feeds R with uniformly random shares $\text{sh}'_1, \dots, \text{sh}'_{t-f}$ conditioned on the correlation string being $\mathbf{r}_2, \dots, \mathbf{r}_t$. Hence, since R is $(n, t, \mathbf{sh}, \Gamma_{\text{sngltn}}, \epsilon)$ -good, in each iteration it holds that

$$R(\text{sh}'_1, \dots, \text{sh}'_{t-f}) = \text{Rec}(\text{sh}_{i_1}, \dots, \text{sh}_{i_f}, \text{sh}'_1, \dots, \text{sh}'_{t-f}) \quad (6)$$

with probability at least ϵ . By definition of the reconstruction algorithm, whenever Eq. (6) holds, we know that

$$R(\text{sh}'_1, \dots, \text{sh}'_{t-f}) = \mathbf{r}_1 + \dots + \mathbf{r}_t, \quad (7)$$

where for each $j \in [t]$, \mathbf{r}_j is the single solution to the system of equations

$$\{\langle \mathbf{a}_{i_k, j}, \mathbf{R}_j \rangle = b_{i_k, j}\}_{k \in [f]} \cup \{\langle \mathbf{a}'_{i, j}, \mathbf{R}_j \rangle = b'_{i, j}\}_{i \in [t-f]}$$

in the variables $\mathbf{R}_j = (R_{j,1}, \dots, R_{j,t})$. Recall that Trace computes \mathbf{r}'_1 as

$$\mathbf{r}'_1 = R(\text{sh}'_1, \dots, \text{sh}'_{t-f}) - (\mathbf{r}_2 + \dots + \mathbf{r}_t). \quad (8)$$

Hence, whenever Eq. (7) holds, it also holds that $\mathbf{r}'_1 = \mathbf{r}_1$. Hence, \mathbf{r}'_1 satisfies $\langle \mathbf{a}_{i_k, 1}, \mathbf{r}'_1 \rangle = b_{i_k, 1}$ for $k = 1, \dots, f$. This means that all corrupted parties – that is, parties i_1, \dots, i_f – are included in the set \mathcal{I}' computed by Trace.

On the other hand, for every party $i \in [n] \setminus \{i_1, \dots, i_f\}$, it holds that its share $\{\langle \mathbf{a}_{i, j}, b_{i, j} \rangle\}_{j \in [t]}$ is statistically independent from both the view of \mathcal{A} and the shares $\text{sh}'_1, \dots, \text{sh}'_{t-f}$ sampled by Trace. Hence, it is independent of \mathbf{r}_1 and $\Pr[\langle \mathbf{a}_{i, 1}, \mathbf{r}'_1 \rangle = b_{i, 1}] = 1/p$. By a union bound, the probability that any honest party is included in \mathcal{I}' is at most t/p .

Overall, we have that in each iteration, the probability that Trace outputs the correct subset $\mathcal{I}' = \mathcal{I}$ is at least $\epsilon - t/p \geq \epsilon/2$. Hence, the probability that it computes the correct subset in at least one iteration is at least

$$1 - (1 - \epsilon/2)^\ell \geq 1 - e^{-\epsilon\ell/2} = 1 - e^{-\lambda}. \quad (9)$$

Moreover, the probability that it outputs the wrong subset $\mathcal{I}' \neq \mathcal{I}$ in a given iteration is at most n/p , and hence the probability that it outputs the wrong subset in any of the $\ell = 2\lambda/\epsilon$ iterations is at most $2n\lambda/(\epsilon p)$.

Overall, we get that the advantage of \mathcal{A} is bounded by $e^{-\lambda} + 2n\lambda/(\epsilon p)$, concluding the proof. \square

4.3 Adding Non-Imputability

In the scheme as presented above, knowledge of the tracing key tk trivially lets one falsely accuse an innocent party of contributing its share to a pirate reconstruction box R . This violates the non-imputability property (Definition 6). The reason is that tk contains the first hyperplane $(\mathbf{a}_{i,1}, b_{i,1})$ of each party in the clear. In what follows, we show that this is not necessary. We observe, that full knowledge of $(\mathbf{a}_{i,1}, b_{i,1})$ is not required for tracing; all that is needed is the ability to determine whether a certain point – namely, \mathbf{r}_1 reconstructed by Trace – lies on the hyperplane induced by $(\mathbf{a}_{i,1}, b_{i,1})$. To this end, it is sufficient to include in the tracing key $f((\mathbf{a}_{i,1})_1), \dots, f((\mathbf{a}_{i,1})_t)$ and $f(b_{i,1})$, where $\mathbf{a}_{i,1} = ((\mathbf{a}_{i,1})_1, \dots, (\mathbf{a}_{i,1})_t)$ and f is a *homomorphic* one-way function. The tracing key tk and verification key vk then include these f -evaluations, where tk additionally includes $\mathbf{r}_2, \dots, \mathbf{r}_t$ as before. The homomorphism allows Trace to check whether $\langle \mathbf{a}_{i,1}, \mathbf{r}_1 \rangle = b_{i,1}$ (“under the hood” of f), and Verify can do the same. Intuitively, if party i is honest, then the reconstruction box R is statistically independent of $\mathbf{a}_{i,1}$. Hence, falsely accusing party i involves finding some non-trivial linear relation among the elements of $f((\mathbf{a}_{i,1})_1), \dots, f((\mathbf{a}_{i,1})_t)$ and $f(b_{i,1})$, which is equivalent to inverting the one-way function f .

We now present the changes to the tracing and verification algorithm in detail, and claim that non-imputability holds. For concreteness, we focus on the case where f is the exponentiation

function in a cyclic group, and hence its one-wayness is based on the hardness of the discrete log problem in the group. For asymptotic reasoning, we consider a distribution ensemble over discrete-log hard groups. This is formalized by the existence of a group generation algorithm \mathcal{G} that takes the security parameter as input and outputs a triple (\mathbb{G}, g, p) , where \mathbb{G} is a description of a group of order p generated by g . We also use the following standard notation. For a group element $h \in \mathbb{G}$ and a vector $\mathbf{x} = (x_1, \dots, x_\ell)$ of \mathbb{Z}_p elements, we write $h^{\mathbf{x}}$ to denote the vector $(h^{x_1}, \dots, h^{x_\ell})$. For a vector $\mathbf{h} = (h_1, \dots, h_\ell)$ of group elements and a vector \mathbf{x} as before, we write denote $\mathbf{h}^{\mathbf{x}} := \prod_i h_i^{x_i}$.

The changes to the scheme are as follows:

1. The share algorithm $\text{Share}(1^\lambda, s, n, t)$ now also samples a group $(\mathbb{G}, g, p) \leftarrow \mathcal{G}(1^\lambda)$. It computes $\mathbf{r}_1, \dots, \mathbf{r}_t$ and $\text{sh}_i = \{(\mathbf{a}_{i,j}, b_{i,j})\}_{j \in [t]}$ as before. Share then computes $\mathbf{y}_i \leftarrow g^{\mathbf{a}_{i,1}}$ and $z_i \leftarrow g^{b_{i,1}}$ for each $i \in [n]$. It sets the tracing key to be

$$\text{tk} \leftarrow ((\mathbb{G}, g, p), \mathbf{r}_2, \dots, \mathbf{r}_t, \mathbf{y}_1, \dots, \mathbf{y}_n, z_1, \dots, z_n)$$

and the verification key $\text{vk} \leftarrow ((\mathbb{G}, g, p), \mathbf{y}_1, \dots, \mathbf{y}_n, z_1, \dots, z_n)$.

2. The tracing algorithm Trace computes \mathbf{r}_1 as before, and sets the set \mathcal{I} of corrupted party to be the set of all indices i that satisfy $\mathbf{y}_i^{\mathbf{r}_1} = z_i$. The proof π now consists of the vector \mathbf{r}_1 .
3. To verify a proof $\pi = \mathbf{r}_1$ against a subset \mathcal{I} , the verification algorithm Verify checks that $\mathbf{y}_i^{\mathbf{r}_1} = z_i$ for every $i \in \mathcal{I}$ and that there exists an $i' \in [n] \setminus \mathcal{I}$ for which $\mathbf{y}_{i'}^{\mathbf{r}_1} \neq z_{i'}$. Verify outputs 1 if and only if both of these conditions are met.

We denote the scheme resulting from these changes as NITB (for “Non-Imputable Traceable Blakeley”). The fact that NITB satisfies correctness, secrecy, and traceability follows from the same analysis as in previous subsections. We now prove that NITB additionally satisfies non-imputability. We do so, by reducing it to the problem of finding a non-trivial linear-in-the-exponent relation among uniformly-random group elements in the group \mathbb{G} .

Definition 11. Let \mathcal{G} be a group generation algorithm and let $\ell = \ell(\lambda)$ be a function of the security parameter $\lambda \in \mathbb{N}$. We say that the Discrete-Log Relation (DLR) problem is hard relative to \mathcal{G} if for every probabilistic polynomial-time algorithm \mathcal{A} there exists a negligible function $\nu(\cdot)$ such that

$$\text{Adv}_{\mathcal{A}, \mathcal{G}}^{\text{dlr}}(\lambda) := \Pr \left[\langle \mathbf{a}, \mathbf{r} \rangle = 0 \wedge \mathbf{r} \neq \mathbf{0} : \begin{array}{l} (\mathbb{G}, p, g) \leftarrow \mathcal{G}(1^\lambda) \\ \mathbf{a} \leftarrow \mathbb{Z}_p^\ell \\ \mathbf{r} \leftarrow \mathcal{A}((\mathbb{G}, p, g), g^{\mathbf{a}}) \end{array} \right] < \nu(\lambda)$$

for all $\lambda \in \mathbb{N}$.

A simple reduction shows that the DLR problem in \mathbb{G} is tightly equivalent to the discrete-log problem in \mathbb{G} (see, for example [15]).

Theorem 6. For every adversary \mathcal{A} for the non-imputability of NITB, there exists an algorithm \mathcal{B} for the DLR problem, such that for every $\lambda \in \mathbb{N}$, it holds that

$$\text{Adv}_{\mathcal{A}, \text{NITB}}^{\text{ni}}(\lambda) = \text{Adv}_{\mathcal{B}, \mathcal{G}}^{\text{dlr}}(\lambda).$$

Proof. Let \mathcal{A} be an adversary participating in the non-imputability experiment for NITB. We construct an algorithm \mathcal{B} that breaks the discrete log relation problem with respect to \mathcal{G} . \mathcal{B} gets as input a vector $g^{\mathbf{x}} = (g^{x_1}, \dots, g^{x_t})$ of group elements sampled uniformly at random from \mathbb{G} . It invokes \mathcal{A} on 1^λ to obtain $n \in \mathbb{N}$, $t \leq n$, a party $i^* \in [n]$, and a secret $\mathbf{s} \in \mathbb{F}_p^t$. It then generates the tracing key, verification key, and shares for \mathcal{A} as follows:

1. It samples an additive secret sharing $\mathbf{r}_1, \dots, \mathbf{r}_t$ for \mathbf{s} .
2. For every $i \in [n] \setminus \{i^*\}$, \mathcal{B} samples a secret share $\text{sh}_i = \{(\mathbf{a}_{i,j}, b_{i,j})\}_{j \in [t]}$ for $\mathbf{r}_1, \dots, \mathbf{r}_t$ honestly. That is, for every $j \in [t]$, $\mathbf{a}_{i,j}$ is uniformly-random and $b_{i,j} = \langle \mathbf{a}_{i,j}, \mathbf{r}_j \rangle$. \mathcal{B} also computes $\mathbf{y}_i \leftarrow g^{a_{i,1}}$ and $z_i \leftarrow g^{b_{i,1}}$ for every $i \neq i^*$.
3. For $i = i^*$, \mathcal{B} sets $\mathbf{y}_{i^*} \leftarrow g^x$ and $z_{i^*} \leftarrow \mathbf{y}_{i^*}^{r_1}$.
4. \mathcal{B} gives \mathcal{A} the shares $\{\text{sh}_i\}_{i \in [n] \setminus \{i^*\}}$, as well as the tracing key

$$\text{tk} \leftarrow ((\mathbb{G}, g, p), \mathbf{r}_2, \dots, \mathbf{r}_t, \mathbf{y}_1, \dots, \mathbf{y}_n, z_1, \dots, z_n)$$

and the verification key $\text{vk} \leftarrow ((\mathbb{G}, g, p), \mathbf{y}_1, \dots, \mathbf{y}_n, z_1, \dots, z_n)$.

\mathcal{A} then outputs a proof π implicating a subset that includes party i^* . If π is not an accepting proof for implicating party i^* , \mathcal{B} aborts and outputs \perp .

For \mathcal{A} to win the security experiment, it must be that π is a vector \mathbf{r}^* such that

$$\mathbf{y}_{i^*}^{r^*} = z_{i^*}. \quad (10)$$

Moreover, for π to be accepting, it must be that $\mathbf{r}^* \neq \mathbf{r}_1$, since $\langle \mathbf{a}_{i',1}, \mathbf{r}_1 \rangle = b_{i',1}$ for some $i' \in [n]$. By construction, we know that

$$\mathbf{y}_{i^*}^{r_1} = z_{i^*}. \quad (11)$$

Dividing Eq. (10) by Eq. (11), we obtain that

$$\mathbf{y}_{i^*}^{r^* - r_1} = 1_{\mathbb{G}}.$$

Recalling that $y_{i^*} = g^x$, this implies that $\langle \mathbf{x}, \mathbf{r}^* - \mathbf{r}_1 \rangle = 0$. Hence, if not aborted before, \mathcal{B} outputs the vector $\mathbf{r}^* - \mathbf{r}_1$.

Observe that \mathcal{B} perfectly simulates the non-imputability experiment to \mathcal{A} . Moreover, whenever \mathcal{A} wins in the experiment, \mathcal{B} succeeds in solving the discrete log relation problem. Hence, $\text{Adv}_{\mathcal{A}, \text{NITB}}^{\text{ni}}(\lambda) = \text{Adv}_{\mathcal{B}, \mathcal{G}}^{\text{dlr}}(\lambda)$, completing the proof of the theorem. \square

5 Discussion and Future Directions

This work raises several open questions regarding traceable secret sharing and coding theory.

Tracing other secret sharing schemes. In this work, we presented tracing procedures for two classic secret sharing schemes (or variants thereof) – those of Shamir [64] and of Blakley [6]. An interesting open question is to devise tracing procedures for other existing secret-sharing schemes. One prime candidate that comes to mind is the Asmuth-Bloom secret sharing scheme, based on the Chinese Remainder Theorem [3, 51, 36]. Coming up with such a tracing procedure is not only a very interesting number theoretic problem, it may also pave the way for traceable secret sharing beyond standard threshold access structures. In particular, as was previously observed, the Asmuth-Bloom scheme can be efficiently extended to give secret sharing for the *weighted threshold* access structure, in which each party is associated with a weight and the secret can be reconstructed if and only if the cumulative weight of the reconstruction quorum clears a certain threshold [71, 32].

Subsequently to this work, Hoffmann [44] presented a tracing procedure for a variant of Mignotte's CRT-based secret sharing scheme [51]. This scheme, however, provides only a weak notion of secrecy, since the secret is the entire CRT code's information word (and not just one of the codeword's coordinates, as in the Asmuth-Bloom secret-sharing scheme). For a secret-sharing scheme that is

derived from an erasure code by encoding the secret as the entire information word (like Mignotte’s scheme), tracing seems easier: Given two random evaluations of the reconstruction box R , the corrupted parties are likely to be the ones whose codeword coordinates remain unchanged. Hoffmann’s tracing algorithm can be seen as implementing this idea for the CRT-specific case. In any case, the problem of tracing leaks in Asmuth-Bloom secret-sharing remains an interesting open problem.

Another direction is to consider traceability in verifiable secret sharing (VSS) schemes. Our constructions cannot support verifiability since our tracing procedures inherently rely on querying the reconstruction box with invalid secret shares as input. We leave the problem of devising traceable VSS schemes as an interesting open problem.

Traceable threshold VUFs. Threshold Verifiable unpredictable functions (VUF) have recently found important applications in blockchains, including for randomness beacons [31, 18, 21, 26] and deterministic wallets [25, 55]. Several constructions of threshold VUFs have been suggested over the years (see, for example, [49, 27, 29, 31, 26] and the references therein). One of the most efficient constructions builds on BLS signatures [9, 49] in bilinear groups, where the signing key is shared among n evaluators using Shamir secret sharing.

As in secret sharing, one can think of the worrisome scenario in which evaluators of a threshold VUF leak their secret evaluation keys. Similarly to the secret sharing case, we can model such leakage as an evaluation box E that has $f < t$ shares of the evaluation key hardcoded in it. An important and interesting open question is to come up with a VUF that has an efficient procedure for *tracing* such a box back to the corrupted evaluators. Though it is tempting to try and use our tracer for Shamir secret sharing to solve this task for the BLS-based threshold VUF, this encounters a problem; this tracing procedure is not linear, and hence cannot be performed in the exponent. Our tracing procedure for Blakley’s scheme *is* linear, but reconstruction in this scheme is not. In an upcoming work, we construct a traceable secret sharing scheme that has linear reconstruction and linear or quadratic tracing.¹⁰ Such a traceable secret sharing scheme will immediately yield a traceable threshold VUF from BLS.

Connection to erasure codes. Many works have observed a tight connection between secret sharing and erasure codes (e.g., [50, 19, 23, 24, 2]). It is easy to see that a t -out-of- n threshold secret sharing scheme immediately gives rise to an erasure code with words of length n that can tolerate up to $n - t$ erasures. The encoding of a word w is the concatenation of its n secret shares (with some fixed randomness for the sharing algorithm). Efficient decoding is guaranteed by the correctness of the secret-sharing scheme. The other direction is less clear: to share a secret s , one can encode it into a codeword of $m \geq n$ symbols. The share of each party is then a random index $i \in [m]$ together with the i th symbol in the codeword.¹¹ However, since encoding is deterministic, this naive approach cannot provide secrecy. A general template for fixing this issue is to first randomly embed the secret s in a higher dimension space, and only then encode it. If this embedding is reversible, then correctness is preserved. Indeed, the previous works that drew a connection between secret sharing and erasure codes can be seen as falling within this paradigm. We observe that the seminal secret-sharing schemes tackled in this work, Shamir’s and Blakley’s, can also be viewed as special cases of this general paradigm. In Shamir’s scheme, the secret $s \in \mathbb{F}_p$ is first randomly extended into a degree $t - 1$ polynomial, which is then encoded using the Reed-Solomon code [62]. In Blakley’s

¹⁰ Tracing may be quadratic, since the bilinear map can be used to compute quadratic functions in the exponent. Reconstruction, however, cannot be quadratic, since the output of the VUF needs to be an element of the source group to allow for efficient verification.

¹¹ If the code supports decoding from any n coordinates, indices can be assigned deterministically.

scheme, the secret $s \in \mathbb{F}_p$ is first randomly extended to a point \mathbf{x} in \mathbb{F}_p^t . Then, \mathbf{x} is encoded using a Hadamard-like code over the field \mathbb{F}_p . That is, positions in the codeword that encode \mathbf{x} are indexed by vectors $\mathbf{a} \in \mathbb{F}_p^t$ (a codeword is of length p^t) and the \mathbf{a} -th position of the codeword is the inner product $\langle \mathbf{a}, \mathbf{x} \rangle$ (our extension of Blakley’s scheme from Section 4.1 can be seen in similar terms).

Adopting this perspective, a fascinating open problem is to devise tracing procedures for other secret sharing schemes that are obtained from erasure codes. Of particular interest are secret-sharing schemes obtained from low-density parity-check codes (LDPC codes) [2, 23]. If such a tracing procedure is linear or quadratic, it might also help make progress towards traceable VUFs. More generally speaking, it may be interesting to study notions of tracing for error-correcting codes, and their potential applications beyond just secret sharing.

In our construction of traceable secret sharing, the codewords were of exponential size, which was necessary due to the non-imputability requirement. However, if we do not care about non-imputability, then the size m of the codeword may be as small as $\approx n^2$, where n is the number of parties. In particular, this means that the problem is meaningful already for explicit LDPC codes with polynomial-length codewords. Observe, that the tracing problem becomes essentially impossible to solve efficiently if $m \ll n^2$, for the same reasons as the lower bound sketched in Section 3.1.¹²

Tracing a reconstruction service. The literature on the related notion of traitor tracing for encryption schemes distinguishes between two types of pirate decoders that need to be traced back to a corrupted party. On the one hand, most works on traitor tracing (e.g., [48, 54, 7, 30, 63, 45, 52, 28, 17, 33, 14, 39, 20, 70, 69, 38]) consider a “decoder box” D_{box} , which is a *stateless* algorithm, whose output distribution is the same across different queries issued to it. On the other hand, some works [46, 60, 65] consider a “decoding service”, which can be thought of as a *stateful* algorithm D_{service} , that can keep state across queries, altering its output distribution from one query to another. This makes tracing harder, since D_{service} may shut down if it notices that the queries issued to it are part of a tracing attempt.

The same distinction can be drawn in the setting of traceable secret sharing. The previous construction of Goyal, Song, and Srinivasan [42] can only trace back stateless reconstruction boxes. Our Blakley-based construction (Section 4) gives a solution to this problem, since tracing requires only one successful query. This is in contrast to our Shamir-based scheme (Section 3), in which tracing requires many different correlated queries.

Acknowledgments. This work was funded by NSF, DARPA, the Simons Foundation, UBRI, and NTT Research. Opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of DARPA.

References

1. N. Alon, O. Goldreich, J. Håstad, and R. Peralta. Simple constructions of almost k -wise independent random variables. In *31st FOCS*, pages 544–553, St. Louis, MO, USA, Oct. 22–24, 1990. IEEE Computer Society Press.
2. B. Applebaum, O. Nir, and B. Pinkas. How to recover a secret with $o(n)$ additions. In H. Handschuh and A. Lysyanskaya, editors, *CRYPTO 2023, Part I*, volume 14081 of *LNCS*, pages 236–262, Santa Barbara, CA, USA, Aug. 20–24, 2023. Springer, Heidelberg, Germany.

¹² If $m \ll n^2$, it becomes hard to even make the reconstruction box R output anything other than \perp , since it requires an exponential number of queries in expectation to query R on a subset of positions such that none of which is corrupted. If $m = \Omega(n^2)$, though, a random subset will not intersect the corrupted subset with high probability.

3. C. Asmuth and J. Bloom. A modular approach to key safeguarding. *IEEE Transactions on Information Theory*, 29(2):208–210, 1983.
4. M. Bellare, W. Dai, and P. Rogaway. Reimagining secret sharing: Creating a safer and more versatile primitive by adding authenticity, correcting errors, and reducing randomness requirements. In *Proceedings in Privacy Enhancing Technologies Symposium (PoPETS)*, volume 4, pages 461–490, 2020.
5. E. R. Berlekamp. Factoring polynomials over large finite fields. *Mathematics of computation*, 24(111):713–735, 1970.
6. G. R. Blakley. Safeguarding cryptographic keys. In *1979 International Workshop on Managing Requirements Knowledge (MARK)*, pages 313–318, 1979.
7. D. Boneh and M. K. Franklin. An efficient public key traitor tracing scheme. In M. J. Wiener, editor, *CRYPTO'99*, volume 1666 of *LNCS*, pages 338–353, Santa Barbara, CA, USA, Aug. 15–19, 1999. Springer, Heidelberg, Germany.
8. D. Boneh, A. Kiayias, and H. W. Montgomery. Robust fingerprinting codes: A near optimal construction. In *Proceedings of the Tenth Annual ACM Workshop on Digital Rights Management, DRM '10*, page 3–12, New York, NY, USA, 2010. Association for Computing Machinery.
9. D. Boneh, B. Lynn, and H. Shacham. Short signatures from the Weil pairing. In C. Boyd, editor, *ASIACRYPT 2001*, volume 2248 of *LNCS*, pages 514–532, Gold Coast, Australia, Dec. 9–13, 2001. Springer, Heidelberg, Germany.
10. D. Boneh and M. Naor. Traitor tracing with constant size ciphertext. In P. Ning, P. F. Syverson, and S. Jha, editors, *ACM CCS 2008*, pages 501–510, Alexandria, Virginia, USA, Oct. 27–31, 2008. ACM Press.
11. D. Boneh, A. Partap, and L. Rotem. Accountability for misbehavior in threshold decryption via threshold traitor tracing. Cryptology ePrint Archive, Paper 2023/1724, 2023. <https://eprint.iacr.org/2023/1724>. To be published in CRYPTO 2024.
12. D. Boneh, A. Sahai, and B. Waters. Fully collusion resistant traitor tracing with short ciphertexts and private keys. In S. Vaudenay, editor, *EUROCRYPT 2006*, volume 4004 of *LNCS*, pages 573–592, St. Petersburg, Russia, May 28 – June 1, 2006. Springer, Heidelberg, Germany.
13. D. Boneh and J. Shaw. Collusion-secure fingerprinting for digital data (extended abstract). In D. Coppersmith, editor, *CRYPTO'95*, volume 963 of *LNCS*, pages 452–465, Santa Barbara, CA, USA, Aug. 27–31, 1995. Springer, Heidelberg, Germany.
14. D. Boneh and M. Zhandry. Multiparty key exchange, efficient traitor tracing, and more from indistinguishability obfuscation. In J. A. Garay and R. Gennaro, editors, *CRYPTO 2014, Part I*, volume 8616 of *LNCS*, pages 480–499, Santa Barbara, CA, USA, Aug. 17–21, 2014. Springer, Heidelberg, Germany.
15. J. Bootle, A. Cerulli, P. Chaidos, J. Groth, and C. Petit. Efficient zero-knowledge arguments for arithmetic circuits in the discrete log setting. In M. Fischlin and J.-S. Coron, editors, *EUROCRYPT 2016, Part II*, volume 9666 of *LNCS*, pages 327–357, Vienna, Austria, May 8–12, 2016. Springer, Heidelberg, Germany.
16. E. F. Brickell. Some ideal secret sharing schemes. In J.-J. Quisquater and J. Vandewalle, editors, *EUROCRYPT'89*, volume 434 of *LNCS*, pages 468–475, Houthalen, Belgium, Apr. 10–13, 1990. Springer, Heidelberg, Germany.
17. H. Chabanne, D. H. Phan, and D. Pointcheval. Public traceability in traitor tracing schemes. In R. Cramer, editor, *EUROCRYPT 2005*, volume 3494 of *LNCS*, pages 542–558, Aarhus, Denmark, May 22–26, 2005. Springer, Heidelberg, Germany.
18. Chainlink vrf: On-chain verifiable randomness. [link](#).
19. H. Chen, R. Cramer, S. Goldwasser, R. de Haan, and V. Vaikuntanathan. Secure computation from random error correcting codes. In M. Naor, editor, *EUROCRYPT 2007*, volume 4515 of *LNCS*, pages 291–310, Barcelona, Spain, May 20–24, 2007. Springer, Heidelberg, Germany.
20. Y. Chen, V. Vaikuntanathan, B. Waters, H. Wee, and D. Wichs. Traitor-tracing from LWE made simple and attribute-based. In A. Beimel and S. Dziembowski, editors, *TCC 2018, Part II*, volume 11240 of *LNCS*, pages 341–369, Panaji, India, Nov. 11–14, 2018. Springer, Heidelberg, Germany.
21. K. Choi, A. Manoj, and J. Bonneau. SoK: Distributed randomness beacons. In *2023 IEEE Symposium on Security and Privacy*, pages 75–92, San Francisco, CA, USA, May 21–25, 2023. IEEE Computer Society Press.
22. B. Chor, A. Fiat, and M. Naor. Tracing traitors. In Y. Desmedt, editor, *CRYPTO'94*, volume 839 of *LNCS*, pages 257–270, Santa Barbara, CA, USA, Aug. 21–25, 1994. Springer, Heidelberg, Germany.
23. R. Cramer, I. B. Damgård, N. Döttling, S. Fehr, and G. Spini. Linear secret sharing schemes from error correcting codes and universal hash functions. In E. Oswald and M. Fischlin, editors, *EUROCRYPT 2015, Part II*, volume 9057 of *LNCS*, pages 313–336, Sofia, Bulgaria, Apr. 26–30, 2015. Springer, Heidelberg, Germany.

24. R. Cramer and C. Xing. Blackbox secret sharing revisited: A coding-theoretic approach with application to expansionless near-threshold schemes. In A. Canteaut and Y. Ishai, editors, *EUROCRYPT 2020, Part I*, volume 12105 of *LNCS*, pages 499–528, Zagreb, Croatia, May 10–14, 2020. Springer, Heidelberg, Germany.
25. P. Das, S. Faust, and J. Loss. A formal treatment of deterministic wallets. In L. Cavallaro, J. Kinder, X. Wang, and J. Katz, editors, *ACM CCS 2019*, pages 651–668, London, UK, Nov. 11–15, 2019. ACM Press.
26. S. Das, B. Pinkas, A. Tomescu, and Z. Xiang. Distributed randomness using weighted vrfs. Cryptology ePrint Archive, Paper 2024/198, 2024. <https://eprint.iacr.org/2024/198>.
27. Y. Dodis. Efficient construction of (distributed) verifiable random functions. In Y. Desmedt, editor, *PKC 2003*, volume 2567 of *LNCS*, pages 1–17, Miami, FL, USA, Jan. 6–8, 2003. Springer, Heidelberg, Germany.
28. Y. Dodis and N. Fazio. Public key trace and revoke scheme secure against adaptive chosen ciphertext attack. In Y. Desmedt, editor, *PKC 2003*, volume 2567 of *LNCS*, pages 100–115, Miami, FL, USA, Jan. 6–8, 2003. Springer, Heidelberg, Germany.
29. Y. Dodis and A. Yampolskiy. A verifiable random function with short proofs and keys. In S. Vaudenay, editor, *PKC 2005*, volume 3386 of *LNCS*, pages 416–431, Les Diablerets, Switzerland, Jan. 23–26, 2005. Springer, Heidelberg, Germany.
30. A. Fiat and T. Tassa. Dynamic traitor training. In M. J. Wiener, editor, *CRYPTO'99*, volume 1666 of *LNCS*, pages 354–371, Santa Barbara, CA, USA, Aug. 15–19, 1999. Springer, Heidelberg, Germany.
31. D. Galindo, J. Liu, M. Ordean, and J.-M. Wong. Fully distributed verifiable random functions and their application to decentralised random beacons. In *2021 IEEE European Symposium on Security and Privacy (EuroS&P)*, pages 88–102, 2021.
32. S. Garg, A. Jain, P. Mukherjee, R. Sinha, M. Wang, and Y. Zhang. Cryptography with weights: MPC, encryption and signatures. In H. Handschuh and A. Lysyanskaya, editors, *CRYPTO 2023, Part I*, volume 14081 of *LNCS*, pages 295–327, Santa Barbara, CA, USA, Aug. 20–24, 2023. Springer, Heidelberg, Germany.
33. S. Garg, A. Kumarasubramanian, A. Sahai, and B. Waters. Building efficient fully collusion-resilient traitor tracing and revocation schemes. In E. Al-Shaer, A. D. Keromytis, and V. Shmatikov, editors, *ACM CCS 2010*, pages 121–130, Chicago, Illinois, USA, Oct. 4–8, 2010. ACM Press.
34. O. Goldreich. *Foundations of Cryptography*. Cambridge University Press, 2001.
35. O. Goldreich and L. A. Levin. A hard-core predicate for all one-way functions. In *21st ACM STOC*, pages 25–32, Seattle, WA, USA, May 15–17, 1989. ACM Press.
36. O. Goldreich, D. Ron, and M. Sudan. Chinese remaindering with errors. In *31st ACM STOC*, pages 225–234, Atlanta, GA, USA, May 1–4, 1999. ACM Press.
37. O. Goldreich, R. Rubinfeld, and M. Sudan. Learning polynomials with queries: The highly noisy case. In *36th FOCS*, pages 294–303, Milwaukee, Wisconsin, Oct. 23–25, 1995. IEEE Computer Society Press.
38. J. Gong, J. Luo, and H. Wee. Traitor tracing with $N^{1/3}$ -size ciphertexts and $O(1)$ -size keys from k -Lin. In C. Hazay and M. Stam, editors, *EUROCRYPT 2023, Part III*, volume 14006 of *LNCS*, pages 637–668, Lyon, France, Apr. 23–27, 2023. Springer, Heidelberg, Germany.
39. R. Goyal, V. Koppula, and B. Waters. Collusion resistant traitor tracing from learning with errors. In I. Diakonikolas, D. Kempe, and M. Henzinger, editors, *50th ACM STOC*, pages 660–670, Los Angeles, CA, USA, June 25–29, 2018. ACM Press.
40. V. Goyal. Reducing trust in the PKG in identity based cryptosystems. In A. Menezes, editor, *CRYPTO 2007*, volume 4622 of *LNCS*, pages 430–447, Santa Barbara, CA, USA, Aug. 19–23, 2007. Springer, Heidelberg, Germany.
41. V. Goyal, S. Lu, A. Sahai, and B. Waters. Black-box accountable authority identity-based encryption. In P. Ning, P. F. Syverson, and S. Jha, editors, *ACM CCS 2008*, pages 427–436, Alexandria, Virginia, USA, Oct. 27–31, 2008. ACM Press.
42. V. Goyal, Y. Song, and A. Srinivasan. Traceable secret sharing and applications. In T. Malkin and C. Peikert, editors, *CRYPTO 2021, Part III*, volume 12827 of *LNCS*, pages 718–747, Virtual Event, Aug. 16–20, 2021. Springer, Heidelberg, Germany.
43. V. Guruswami and M. Sudan. Improved decoding of Reed-Solomon and algebraic-geometric codes. In *39th FOCS*, pages 28–39, Palo Alto, CA, USA, Nov. 8–11, 1998. IEEE Computer Society Press.
44. C. Hoffmann. Traceable secret sharing based on the chinese remainder theorem. Cryptology ePrint Archive, Paper 2024/811, 2024. <https://eprint.iacr.org/2024/811>.
45. A. Kiayias and M. Yung. Self protecting pirates and black-box traitor tracing. In J. Kilian, editor, *CRYPTO 2001*, volume 2139 of *LNCS*, pages 63–79, Santa Barbara, CA, USA, Aug. 19–23, 2001. Springer, Heidelberg, Germany.
46. A. Kiayias and M. Yung. On crafty pirates and foxy tracers. In *Security and Privacy in Digital Rights Management: ACM CCS-8 Workshop DRM 2001 Philadelphia, PA, USA, November 5, 2001 Revised Papers*, pages 22–39. Springer, 2002.

47. D. E. Knuth. *Art of computer programming, volume 2: Seminumerical algorithms*. Addison-Wesley Professional, 2014.
48. K. Kurosawa and Y. Desmedt. Optimum traitor tracing and asymmetric schemes. In K. Nyberg, editor, *EUROCRYPT'98*, volume 1403 of *LNCS*, pages 145–157, Espoo, Finland, May 31 – June 4, 1998. Springer, Heidelberg, Germany.
49. A. Lysyanskaya. Unique signatures and verifiable random functions from the DH-DDH separation. In M. Yung, editor, *CRYPTO 2002*, volume 2442 of *LNCS*, pages 597–612, Santa Barbara, CA, USA, Aug. 18–22, 2002. Springer, Heidelberg, Germany.
50. J. L. Massey. Some applications of coding theory in cryptography. *Codes and Ciphers: Cryptography and Coding IV*, pages 33–47, 1995.
51. M. Mignotte. How to share a secret? In T. Beth, editor, *EUROCRYPT'82*, volume 149 of *LNCS*, pages 371–375, Burg Feuerstein, Germany, Mar. 29 – Apr. 2, 1983. Springer, Heidelberg, Germany.
52. D. Naor, M. Naor, and J. Lotspiech. Revocation and tracing schemes for stateless receivers. In J. Kilian, editor, *CRYPTO 2001*, volume 2139 of *LNCS*, pages 41–62, Santa Barbara, CA, USA, Aug. 19–23, 2001. Springer, Heidelberg, Germany.
53. J. Naor and M. Naor. Small-bias probability spaces: Efficient constructions and applications. In *22nd ACM STOC*, pages 213–223, Baltimore, MD, USA, May 14–16, 1990. ACM Press.
54. M. Naor and B. Pinkas. Threshold traitor tracing. In H. Krawczyk, editor, *CRYPTO'98*, volume 1462 of *LNCS*, pages 502–517, Santa Barbara, CA, USA, Aug. 23–27, 1998. Springer, Heidelberg, Germany.
55. J. Nick, T. Ruffing, Y. Seurin, and P. Wuille. MuSig-DN: Schnorr multi-signatures with verifiably deterministic nonces. In J. Ligatti, X. Ou, J. Katz, and G. Vigna, editors, *ACM CCS 2020*, pages 1717–1731, Virtual Event, USA, Nov. 9–13, 2020. ACM Press.
56. K. Nuida. A general conversion method of fingerprint codes to (more) robust fingerprint codes against bit erasure. In K. Kurosawa, editor, *ICITS 09*, volume 5973 of *LNCS*, pages 194–212, Shizuoka, Japan, Dec. 3–6, 2010. Springer, Heidelberg, Germany.
57. B. Pfitzmann. Trials of traced traitors. In R. Anderson, editor, *Information Hiding*, pages 49–64, Berlin, Heidelberg, 1996. Springer Berlin Heidelberg.
58. B. Pfitzmann and M. Schunter. Asymmetric fingerprinting (extended abstract). In U. M. Maurer, editor, *EUROCRYPT'96*, volume 1070 of *LNCS*, pages 84–95, Saragossa, Spain, May 12–16, 1996. Springer, Heidelberg, Germany.
59. B. Pfitzmann and M. Waidner. Asymmetric fingerprinting for larger collusions. In R. Graveman, P. A. Janson, C. Neuman, and L. Gong, editors, *ACM CCS 97*, pages 151–160, Zurich, Switzerland, Apr. 1–4, 1997. ACM Press.
60. D. H. Phan. Traitor tracing for stateful pirate decoders with constant ciphertext rate. In *International Conference on Cryptology in Vietnam*, pages 354–365. Springer, 2006.
61. M. O. Rabin. Probabilistic algorithms in finite fields. *SIAM Journal on computing*, 9(2):273–280, 1980.
62. I. S. Reed and G. Solomon. Polynomial codes over certain finite fields. *Journal of the society for industrial and applied mathematics*, 8(2):300–304, 1960.
63. R. Safavi-Naini and Y. Wang. Sequential traitor tracing. In M. Bellare, editor, *CRYPTO 2000*, volume 1880 of *LNCS*, pages 316–332, Santa Barbara, CA, USA, Aug. 20–24, 2000. Springer, Heidelberg, Germany.
64. A. Shamir. How to share a secret. *Communications of the Association for Computing Machinery*, 22(11):612–613, Nov. 1979.
65. T. Sirvent. Traitor tracing scheme with constant ciphertext rate against powerful pirates. Cryptology ePrint Archive, Paper 2006/383, 2006. <https://eprint.iacr.org/2006/383>.
66. M. Sudan. Maximum likelihood decoding of reed solomon codes. In *37th FOCS*, pages 164–172, Burlington, Vermont, Oct. 14–16, 1996. IEEE Computer Society Press.
67. A. Ta-Shma. Explicit, almost optimal, epsilon-balanced codes. In H. Hatami, P. McKenzie, and V. King, editors, *49th ACM STOC*, pages 238–251, Montreal, QC, Canada, June 19–23, 2017. ACM Press.
68. G. Tardos. Optimal probabilistic fingerprint codes. *J. ACM*, 55(2), may 2008.
69. H. Wee. Functional encryption for quadratic functions from k -lin, revisited. In R. Pass and K. Pietrzak, editors, *TCC 2020, Part I*, volume 12550 of *LNCS*, pages 210–228, Durham, NC, USA, Nov. 16–19, 2020. Springer, Heidelberg, Germany.
70. M. Zhandry. New techniques for traitor tracing: Size $N^{1/3}$ and more from pairings. In D. Micciancio and T. Ristenpart, editors, *CRYPTO 2020, Part I*, volume 12170 of *LNCS*, pages 652–682, Santa Barbara, CA, USA, Aug. 17–21, 2020. Springer, Heidelberg, Germany.
71. X. Zou, F. Maino, E. Bertino, Y. Sui, K. Wang, and F. Li. A new approach to weighted multi-secret sharing. In *2011 Proceedings of 20th International Conference on Computer Communications and Networks (ICCCN)*, pages 1–6. IEEE, 2011.

A A Simple Tracing Algorithm

In this section we present a simple tracing algorithm that can trace perfect reconstruction boxes for any secret sharing scheme. The main drawback of this scheme, compared to the tracing constructions in the main body, is that in this scheme, the tracing key has to include the secret s .

The basic idea is as follows. Suppose for simplicity that the reconstruction box R has $f = t - 1$ shares $\{\text{sh}_{i_j}\}_{j \in [t-1]}$ hardcoded in it, and it takes one additional share sh as input. By perfect, we mean that R always outputs the result of the Rec algorithm applied to $\{\text{sh}_{i_j}\}_{j \in [t-1]}$ and the input share sh . The Rec algorithm outputs \perp when it is given fewer than t distinct shares as input. For each party $i \in [n]$, the tracer simply runs the box with input sh_i . Observe that if party i is honest, then the box can successfully reconstruct the secret, and since it is perfect, it must output the reconstructed secret. On the other hand, if i is corrupt, then the input share is already contained in the box, and hence the box does not have enough information to reconstruct the secret. The tracer exonerates all the parties i for which the output of the box when given sh_i is the correct secret, and blames all the remaining parties.

We generalize to any number of corrupted parties $1 \leq f < t - 1$ as follows. For each party $i \in [n]$, the tracer samples $t - f - 1$ random shares $\{\text{sh}'_i\}_{i \in [t-f-1]}$ by running the Share algorithm $t - f - 1$ times. If any of these shares is equal to sh_j for some party $j \in [n]$, then the tracer aborts. Otherwise, it runs the reconstruction box with these random shares $\{\text{sh}'_i\}_{i \in [t-f-1]}$ along with sh_i as input. Similar to before, the tracer exonerates party i if the box outputs the correct secret. To bound the tracer's abort probability, we assume that the Share algorithm outputs a uniformly random share from a super-polynomial-sized space¹³.

Figure 5 presents the full tracing algorithm. This tracing procedure is generic; it builds a traceable secret sharing scheme from any secret sharing scheme SS , as long as it is symmetric (as defined in Section 2) and the shares come from a super-polynomial-sized space. We denote this generic traceable secret sharing scheme as GTSS . Correctness and secrecy of GTSS follow directly from that of the underlying secret sharing scheme SS . Theorem 7 proves universal traceability of GTSS . Our analysis relies on the assumption that the Share algorithm outputs a uniformly random share from a super-polynomial-sized space.




Fig. 5. The trivial tracing algorithm for any secret sharing scheme SS with a super-polynomial-sized share space.

¹³ Note that if the share space were polynomial-sized, then tracing would be impossible by the lower bound discussed in Section 3.1.

Theorem 7. For every adversary \mathcal{A} , for every $\lambda \in \mathbb{N}$, it holds that

$$\text{Adv}_{\mathcal{A}, \text{GTSS}}^{\text{uni-trac}}(\lambda) \leq \frac{n \cdot t \cdot (n + t)}{q}$$

where q is the size of the space of secret shares for SS, $n = n(\lambda)$, $t = t(\lambda)$ and $f = f(\lambda)$ are upper bounds on the number of parties, threshold and the number of corruptions, respectively.

Proof. Let \mathcal{A} be an adversary taking part in the $\text{ExpUniTrace}_{\mathcal{A}, \text{GTSS}, 1}(\lambda)$ security experiment. Let G denote the event in which \mathcal{A} outputs a reconstruction box R that is $(n, t, \mathbf{sh}, 1)$ -good, where n and t are chosen by \mathcal{A} at the beginning of the experiment, and $\mathbf{sh} = (\text{sh}_{i_1}, \dots, \text{sh}_{i_f})$ are the shares given to it by the challenger. Conditioned on $\neg \text{G}$, the output of the experiment is 0 with probability 1, and so we condition the rest of the analysis on G . Let \mathcal{I}^* be the set of parties corrupted by \mathcal{A} , i.e. $\{\text{sh}_{i_j}\}_{j \in [f]} = \{\text{sh}_i\}_{i \in \mathcal{I}^*}$.

Let s, ρ denote the secret and the correlation string sampled by the challenger.

For every $\ell \in [n]$, let E_ℓ denote the event that the tracer aborts due to $\text{sh}_{\ell, i} = \text{sh}_{\ell, j}$ for some $i \neq j \in \{f + 1, \dots, t - 1\}$ or $\text{sh}_{\ell, i} = \text{sh}_j$ for some $i \in \{f + 1, \dots, t - 1\}$ and some $j \in [n]$. For any $j \in \{f + 1, \dots, t - 1\}$, the probability that $\text{sh}_{\ell, j} = \text{sh}_{\ell, i}$ for some $f + 1 \leq i < j$ or $\text{sh}_{\ell, j} = \text{sh}_k$ for some $k \in [n]$ is $\leq (n + t - f - 1)/q$, because the Share algorithm outputs a uniformly random share from the share space, which is of size q . Hence, by a union bound, we have that $\Pr[\text{E}_\ell] \leq (t - f - 1)(n + t - f - 1)/q$.

Since R is a perfect universally-good box, we have that, $\forall \ell \in [n]$,

$$R(\text{sh}_{\ell, f+1}, \dots, \text{sh}_{\ell, t}) = \text{Rec}(\text{sh}_{i_1}, \dots, \text{sh}_{i_f}, \text{sh}_{\ell, f+1}, \dots, \text{sh}_{\ell, t}).$$

Now, consider the following two cases:

- $\ell \in \mathcal{I}^*$. In this case, $\text{sh}_{\ell, t}$ is equal to some share in \mathbf{sh} . This means that the tuple $(\text{sh}_{i_1}, \dots, \text{sh}_{i_f}, \text{sh}_{\ell, f+1}, \dots, \text{sh}_{\ell, t})$ has less than t distinct shares, and hence, reconstruction will not succeed and Rec will output \perp .
- $\ell \in [n] \setminus \mathcal{I}^*$. In this case, observe that if the tracer does not abort, then, the tuple $(\text{sh}_{i_1}, \dots, \text{sh}_{i_f}, \text{sh}_{\ell, f+1}, \dots, \text{sh}_{\ell, t})$ contains t distinct secret shares. This is because (i) $\text{sh}_{\ell, t} = \text{sh}_\ell \notin \mathbf{sh}$, (ii) there is no $i \neq j \in \{f + 1, \dots, t - 1\}$ such that $\text{sh}_{\ell, i} = \text{sh}_{\ell, j}$ and (iii) none of the $\text{sh}_{\ell, i}$ can be equal to the shares in the box. Hence, the Rec algorithm will successfully reconstruct the secret, and output s .

Let Abort be the event that the tracer aborts. By the above analysis, in the event $\overline{\text{Abort}}$, if $s_\ell = s$, then party ℓ must be innocent. Equivalently, if $s_\ell \neq s$, then party ℓ must be in \mathcal{I}^* . This means that the tracer successfully traces the box to all the corrupt parties. We now bound the probability of the tracer aborting:

$$\begin{aligned} \Pr[\text{Abort}] &= \Pr \left[\bigvee_{\ell \in [n]} \text{E}_\ell \right] \leq \sum_{\ell \in [n]} \Pr[\text{E}_\ell] \\ &\leq n \cdot \frac{(t - f - 1)(n + t - f - 1)}{q} \leq \frac{n \cdot t \cdot (n + t)}{q} \end{aligned}$$

This proves the theorem. □

If the number of corruptions f is not known in advance, the tracing algorithm can learn it by simply trying $f = t - 1, t - 2, \dots$ until it reaches a value that works. Specifically, suppose that the real number of corruptions is f^* . Then, for any value of $f > f^*$, the box will not be able to reconstruct the secret for any $\ell \in [n]$, since the box will only have upto $f^* + t - f$ distinct shares which is less than t . Hence, if the tracer does not abort, it can detect that $f > f^*$. Additionally, by the analysis in the proof of Theorem 7, when Trace tries $f = f^*$, the algorithm will fail to output the correct subset with probability atmost $nt(n+t)/q$. We get that, by a union bound, the probability that Trace correctly traces R back to the corrupted set is atleast $1 - n^2t(n+t)/q$.

Limitations. The tracing procedure presented in this section has three major drawbacks. First, the tracer knows the secret s as well as the correlation string ρ since they are both part of the tracing key. Second, the procedure only works for perfect boxes, which always output the correct reconstructed secret. Third, it is unclear how to achieve non-imputability since the tracer needs to know the shares sh_i for all n parties to be able to trace. We leave the task of overcoming these limitations as an interesting direction for future work.

B Proof of Lemma 1

Proof (Lemma 1). Let \mathcal{A} , ϵ , δ , and λ be as in the statement of the lemma. Consider the tracing security experiment $\mathbf{ExpTrace}_{\mathcal{A}, \text{TTSS}, \epsilon, \delta}(\lambda)$, and denote by $\alpha := \text{Adv}_{\mathcal{A}, \text{TTSS}, \epsilon, \delta}^{\text{trac}}(\lambda)$. By definition 5 and our assumption that $\Pr[\mathbf{ExpTrace}_{\mathcal{A}, \text{TTSS}, \epsilon, \delta}(\lambda) = 1] \geq \frac{1}{\epsilon|\mathcal{SCR}\mathcal{T}_\lambda|}$, we get that

$$\Pr[\mathbf{ExpTrace}_{\mathcal{A}, \text{TTSS}, \epsilon, \delta}(\lambda) = 1] = \frac{\alpha + \frac{1}{|\mathcal{SCR}\mathcal{T}_\lambda|}}{\epsilon}. \quad (12)$$

Let G denote the event in which the reconstruction box R outputted by \mathcal{A} is $(n, t, f, s^*, \rho, \epsilon)$ -good, i.e., the probability that the experiment does not return 0 in line 8. Let T denote the event in which the tracing algorithm Trace outputs the correct subset \mathcal{I} and an accepting proof in line 9. Then, by definition

$$\Pr[\mathbf{ExpTrace}_{\mathcal{A}, \text{TTSS}, \epsilon, \delta}(\lambda) = 1] = \Pr[\mathsf{G} \wedge \neg\mathsf{T}] \quad (13)$$

where the probability is over the choice of $s^* \leftarrow_{\$} \mathcal{SCR}\mathcal{T}_\lambda$, $\rho \leftarrow_{\$} \{0, 1\}^\kappa$, $(\text{sh}_i, \text{tk}_i, \text{vk}_i) \leftarrow_{\$} \text{Share}(1^\lambda, s^*, n, t, \rho)$, and the random coins of \mathcal{A} and Trace.

We now consider the probability that the reconstruction box R outputs s^* on $t - f$ random well-formed secret shares of s^* and is not traced back to the subset \mathcal{I} of corrupted parties. From Eq. (12) and (13), the definition of a good reconstruction box, and the law of total probability, we have that

$$\Pr[R(\text{sh}'_1, \dots, \text{sh}'_{t-f}) = s^* \wedge \neg\mathsf{T}] \geq \alpha + \frac{1}{|\mathcal{SCR}\mathcal{T}_\lambda|}, \quad (14)$$

where the probability is over the same choices as in Eq. 13 and also over the choice of $(\text{sh}'_i, \text{tk}'_i, \text{vk}'_i) \leftarrow_{\$} \text{Share}(1^\lambda, s^*, n, t, \rho)$ for $i = 1, \dots, t - f$ and the random coins of R .

The correctness guarantee of TTSS implies that

$$\begin{aligned} \Pr[R(\text{sh}'_1, \dots, \text{sh}'_{t-f}) = \text{Rec}(\text{sh}_1, \dots, \text{sh}_f, \text{sh}'_1, \dots, \text{sh}'_{t-f}) \wedge \neg\mathsf{T}] \\ \geq \alpha + \frac{1}{|\mathcal{SCR}\mathcal{T}_\lambda|} - \nu_1, \end{aligned} \quad (15)$$

Now consider a different way of sampling the random variables underlying the probability in Eq. (15). We sample random coins for \mathcal{A} , R and Trace as before. The f shares $\text{sh}_1, \dots, \text{sh}_f$ given to \mathcal{A} are sampled by:

1. Sample a random subset $\mathcal{C} \in \Gamma$ by sampling $\rho \leftarrow_{\$} \{0, 1\}^\kappa$ and setting $\mathcal{C} \leftarrow \Gamma(\rho)$.
2. For each $i \in [f]$ we sample $\text{sh}_i \leftarrow_{\$} \mathcal{SH}(\mathcal{C})$.

We then sample $\text{sh}'_1 \dots, \text{sh}'_{t-f}$ independently according to $\mathcal{SH}(\mathcal{C})$ as well.

By ν_2 -bidirectionality, we obtain that

$$\begin{aligned} \Pr [R(\text{sh}'_1, \dots, \text{sh}'_{t-f}) = \text{Rec}(\text{sh}_1, \dots, \text{sh}_f, \text{sh}'_1, \dots, \text{sh}'_{t-f}) \wedge \neg \top] \\ \geq \alpha + \frac{1}{|\mathcal{SCRT}_\lambda|} - (\nu_1 + \nu_2), \end{aligned} \tag{16}$$

where the distribution is over the modified sampling of $\text{sh}_1, \dots, \text{sh}_f, \text{sh}'_1, \dots, \text{sh}'_{t-f}$ as described above, and over the random coins of \mathcal{A}, R and Trace .

By total probability, with probability at least $\alpha/2 + 1/|\mathcal{SCRT}_\lambda|$ over the choice of \mathcal{C} , the choice of $\text{sh}_1, \dots, \text{sh}_f \leftarrow_{\$} \mathcal{SH}(\mathcal{C})$ and the random coins of \mathcal{A} and Trace it holds that $\neg \top$ occurs and

$$\Pr [R(\text{sh}'_1, \dots, \text{sh}'_{t-f}) = \text{Rec}(\text{sh}_1, \dots, \text{sh}_f, \text{sh}'_1, \dots, \text{sh}'_{t-f})] \geq \frac{\alpha}{2} - (\nu_1 + \nu_2) \tag{17}$$

over the choice of $\text{sh}'_1, \dots, \text{sh}'_{t-f} \leftarrow_{\$} \mathcal{SH}(\mathcal{C})$ and the random coins of R . The lemma then follows from invoking bidirectionality one more time, and sampling the shares and secret as in the original experiment. \square