





Scooby: Improved Multi-Party Homomorphic Secret Sharing Based on FHE

Ilaria Chillotti¹ , Emmanuela Orsini² , Peter Scholl³ , Nigel Paul Smart^{1,2} , and Barry
Van Leeuwen² 

¹ Zama, Paris, France,

² imec-COSIC, KU Leuven, Leuven, Belgium,

³ U. Aarhus, Aarhus, Denmark.

`ilaria.chillotti@zama.ai,`
`emmanuela.orsini@kuleuven.be,`
`peter.scholl@cs.au.dk,`
`nigel.smart@kuleuven.be,`
`barry.vanleeuwen@kuleuven.be`

Abstract. We present new constructions of multi-party homomorphic secret sharing (HSS) based on a new primitive that we call *homomorphic encryption with decryption to shares* (HEDS). Our first construction, which we call **Scooby**, is based on many popular fully homomorphic encryption (FHE) schemes with a linear decryption property. **Scooby** achieves an n -party HSS for general circuits with complexity $O(|F| + \log n)$, as opposed to $O(n^2 \cdot |F|)$ for the prior best construction based on multi-key FHE. **Scooby** can be based on (ring)-LWE with a super-polynomial modulus-to-noise ratio. In our second construction, **Scrappy**, assuming any generic FHE plus HSS for NC1-circuits, we obtain a HEDS scheme which does not require a super-polynomial modulus. While these schemes all require FHE, in another instantiation, **Shaggy**, we show how in some cases it is possible to obtain multi-party HSS without FHE, for a small number of parties and constant-degree polynomials. Finally, we show that our **Scooby** scheme can be adapted to use multi-key fully homomorphic encryption, giving more efficient spooky encryption and setup-free HSS. This latter scheme, **Casper**, if concretely instantiated with a B/FV-style multi-key FHE scheme, for functions F which do not require bootstrapping, gives an HSS complexity of $O(n \cdot |F| + n^2 \cdot \log n)$.

Table of Contents

Scooby: Improved Multi-Party Homomorphic Secret Sharing Based on FHE	1
<i>Ilaria Chillotti^{ID}, Emmanuela Orsini^{ID}, Peter Scholl^{ID}, Nigel Paul Smart^{ID}, and Barry Van Leeuwen^{ID}</i>	
1 Introduction	3
1.1 Our Contribution	4
2 Preliminaries	6
2.1 Homomorphic Secret Sharing	7
2.2 Spooky Encryption	8
3 Homomorphic Encryption with Decryption to Shares (HEDS)	9
3.1 Multi-input HSS from HEDS Encryption	10
4 Linear-Decryption Based FHE	11
4.1 Two-Party Distributed Decryption: Type <i>lsb</i>	12
5 Scooby: Multi-Party HEDS from LD-based FHE	13
5.1 HEDS Key Generation	14
5.2 Security Assumption	14
5.3 From 2-party to n -party HEDS	14
5.4 BGV Parameters Supporting Scooby	19
6 Multi-Party HEDS from Weaker Assumptions	19
6.1 Scrappy: HEDS from Standard FHE + HSS for NC1	19
6.2 Shaggy: Bootstrapping HEDS to More Parties	22
7 Casper: Friendly AFS-spooky Encryption from Multi-Key FHE	25
7.1 Spooky from MK-TFHE	26
7.2 Spooky from MK-BFV	28
A Two Party Distributed Decryption: Type <i>msb</i>	33
A.1 n -party Scooby Decryption: the <i>msb</i> case	34

1 Introduction

One of the more interesting cryptographic constructions to be developed in recent years has been homomorphic secret sharing (HSS). This concept, which can be seen as a distributed analogue of homomorphic encryption, was introduced in [BGI16a], where a two party construction for branching programs was presented based on the decisional Diffie-Hellman assumption. The idea of HSS starts from the concept of a (traditional) secret sharing scheme, where an input x to some function is split into n shares, $(\mathbf{x}_1, \dots, \mathbf{x}_n)$. This sharing, that in this work we always assume to be a full threshold sharing, is created via an algorithm $(\mathbf{x}_1, \dots, \mathbf{x}_n) \leftarrow \text{Share}^{\text{HSS}}(x)$. An HSS scheme has two additional algorithms, the first $y_i \leftarrow \text{Eval}^{\text{HSS}}(F; \mathbf{x}_i)$ takes a function description F and a share \mathbf{x}_i and produces a corresponding output share y_i . The second $\text{Rec}^{\text{HSS}}(y_1, \dots, y_n)$ takes the output shares and reconstructs the result $F(x)$. To avoid trivial solutions one requires that the length of the y_j 's should be *compact*, i.e. it only depends on the output length of the function F and the security parameter. An important class of HSS schemes are those with *additive reconstruction*, where the function Rec^{HSS} simply computes $y_1 + \dots + y_n$. We refer to these as *additive HSS schemes*. It is such additive HSS schemes that we focus on in this work.

Motivation for HSS. The main application of HSS is towards secure two-party or multi-party computation with succinct communication. Indeed, the breakthrough work of [BGI16a] showed that for a large class of circuits, it's possible to achieve secure computation with sublinear communication in the circuit size under DDH, which was previously only known using fully homomorphic encryption. Since then, HSS has proven useful in various other applications, and is closely related to pseudorandom correlation generators [BCG⁺19] and pseudorandom correlation functions [BCG⁺20], which allow generation of correlated randomness with a minimal amount of interaction. HSS for simple classes of functions, particularly the case of distributed point functions [GI14], has also proven useful for applications including private information retrieval [BGI16b] and secure RAM computation [Ds17]. On a more theoretical side, HSS has also been used to build 2-round secure computation and nearly optimal worst-case to average-case reductions [BGI⁺18].

Additive reconstruction is an important feature of HSS in many secure computation settings, where it may be desirable for the output shares to be re-used in another secure computation based on secret sharing. This is the case, for instance, when using HSS to generate preprocessing material for multi-party computation protocols in the dishonest majority setting [BCG⁺19]. It can also be a useful feature in scenarios where a client reconstructing the output is constrained to perform only lightweight computations.

Current State of HSS and Related Primitives. Related to HSS is the dual concept of function secret sharing (FSS) [BGI15, BGI16b]. In FSS, the shared data is a secret function F (from some publicly known class of functions), such that the parties can locally obtain secret shares of $F(x)$, for any public input x . For general function classes such as polynomially-sized circuits, function secret sharing and homomorphic secret sharing are equivalent.

Obtaining efficient n -party HSS and FSS is complex for general functions. The most efficient known scheme is that based on an LWE-construction from spooky encryption. Spooky encryption, introduced by Dodis et al. [DHRW16], is a rather complex construction based on a multi-key variant of FHE [CM15, MW16], and for our purposes we are only interested in *additive-function-sharing* spooky encryption (or AFS-spooky encryption). Spooky encryption is a semantically secure public-key encryption scheme consisting of the usual three algorithms $(\text{KeyGen}^{\text{Spooky}}, \text{Enc}_{\text{pk}}^{\text{Spooky}}, \text{Dec}_{\text{sk}}^{\text{Spooky}})$

as well as an additional algorithm $\text{Eval}_{\text{pk}_1, \dots, \text{pk}_n}^{\text{Spooky}}(F, \text{ct}_1, \dots, \text{ct}_n)$. The $\text{Eval}^{\text{Spooky}}$ algorithm, given a function F on n arguments from a given class, and n ciphertexts ct_i , encrypting x_i under pk_i , produces n new ciphertexts $\text{ct}'_1, \dots, \text{ct}'_n$ such that, computing $y_i \leftarrow \text{Dec}_{\text{sk}_i}^{\text{Spooky}}(\text{ct}_i)$, we have that $y_1 + \dots + y_n = F(x_1, \dots, x_n)$.

In [DHRW16], it is shown that it is possible to build FSS from AFS-spooky encryption. Roughly, to *share* an input function F the dealer first generates n AFS-spooky key pairs $(\text{pk}_i, \text{sk}_i) \leftarrow \text{KeyGen}^{\text{Spooky}}(1^\lambda)$. The dealer also generates an n -out-of- n description of the function F , i.e. functions $F_i(x)$ such that $F(x) = F_1(x) + \dots + F_n(x)$. Finally, the function secret sharing of the input function F is defined to be the tuple $F_i = (\text{sk}_i, \text{pk}_1, \dots, \text{pk}_n, \text{Enc}_{\text{pk}_1}^{\text{Spooky}}(F_1), \dots, \text{Enc}_{\text{pk}_n}^{\text{Spooky}}(F_n))$.

To define the FSS evaluation we create a function C_x which takes as input the n additive shares of a function F , and evaluates it on the input x , which is hard-coded into C_x . By applying

$$\text{Eval}_{\text{pk}_1, \dots, \text{pk}_n}^{\text{Spooky}}\left(C_x, \text{Enc}_{\text{pk}_1}^{\text{Spooky}}(F_1), \dots, \text{Enc}_{\text{pk}_n}^{\text{Spooky}}(F_n)\right),$$

we obtain ciphertexts $\text{ct}'_1, \dots, \text{ct}'_n$, where ct'_i can be decrypted (using sk_i) to obtain y_i such that $y_1 + \dots + y_n = F(x)$.

In [BGI⁺18], Boyle et al. showed how the FSS construction from spooky encryption can be modified to enable an additive HSS scheme. The $\text{Share}^{\text{HSS}}(x)$ operation additively shares x into $x = x_1 + \dots + x_n$, generates n spooky key pairs $(\text{pk}_1, \text{sk}_1) \leftarrow \text{KeyGen}^{\text{Spooky}}(1^\lambda)$, and then encrypts x_i via $\text{ct}_i \leftarrow \text{Enc}_{\text{pk}_i}^{\text{Spooky}}(x_i)$. The share values \mathbf{x}_i output by $\text{Share}^{\text{HSS}}(x)$ being $\mathbf{x}_i = (\{\text{pk}_i\}_{i=1}^n, \{\text{ct}_i\}_{i=1}^n, \text{sk}_i)$. The $\text{Eval}^{\text{HSS}}(F, \mathbf{x}_i)$ function executes $\text{Eval}_{\text{pk}_1, \dots, \text{pk}_n}^{\text{Spooky}}$ on the function F and the ciphertext $(\text{ct}_1, \dots, \text{ct}_n)$ so as to obtain n ciphertexts $\text{ct}'_1, \dots, \text{ct}'_n$. The output of $\text{Eval}^{\text{HSS}}(F, \mathbf{x}_i)$ then being $\text{Dec}_{\text{sk}_i}^{\text{Spooky}}(\text{ct}'_i)$.

Thus, there is a strong connection between HSS, FSS and spooky constructions, and, as mentioned above, the prior most efficient n -party HSS and FSS constructions for circuits arise from AFS-spooky based on LWE (and a circular security assumption). The best current construction for AFS-spooky encryption of Dodis et al. [DHRW16] has a complexity of $O(n^2 \cdot |F|)$. In particular, each gate of the underlying arithmetic circuit F requires a bootstrapping operation which in the multi-key FHE setting has complexity $O(n^2)$.

1.1 Our Contribution

We present new constructions of homomorphic secret sharing in the multi-party setting, supporting up to $n - 1$ out of n corruptions. Our constructions improve upon the only previous general construction, based on AFS-spooky encryption [DHRW16], either by being more efficient, or in some cases, relying on different assumptions.

HSS from Homomorphic Encryption with Decryption Shares. We present our constructions as a new primitive called *homomorphic encryption with decryption to shares (HEDS)*, which can be seen as a homomorphic encryption scheme with a special decryption algorithm that (non-interactively) outputs an n -party secret share of the encrypted message. HEDS is closely related to both spooky encryption and homomorphic secret sharing (HSS): the major difference compared to spooky is that HEDS needs to set up private decryption keys under a common public key with either a trusted setup algorithm or a secure multiparty computation protocol, while the difference with HSS is that the homomorphic evaluation algorithm is public. As is the case for spooky, HEDS immediately implies additive HSS for the same class of functions.

Scooby Construction: HEDS from Linear Decryption FHE. We show that HEDS can be built using any FHE scheme with a special decryption property, which we call *linear decryption based fully homomorphic encryption* (LD-based FHE) schemes. Examples of such LD-based FHE schemes are LWE-based constructions like BGV [BGV12], BFV [FV12], GSW [GSW13] and TFHE [CGGI16, CGGI20]. Notice this special property of almost all FHE schemes, where the decryption function is a linear function of the secret key, has been exploited previously, including for HSS in the two-party setting [DHRW16, BKS19] and other applications [BDGM19, GH19].

Any of these schemes can be used to instantiate our Scooby construction, giving additive HSS for circuits. Recall in AFS-spooky the key generation is run independently by the n -parties, in our variation the keys are instead generated by a trusted third party.⁴

Since this construction only requires single-key FHE and not multi-key FHE, we obtain n -party HSS that is simpler and more efficient than the AFS-spooky-based construction. In particular, the computational complexity grows as $O(|F| + \log n)$, whereas AFS-spooky has complexity $O(n^2 \cdot |F|)$ for n parties. In addition, when instantiated with BGV we show that the standard parameter sets for bootstrapping are sufficient for our construction.

At a high level, at the core of Scooby is a well-known 2-party distributed decryption procedure, which non-interactively decrypts an LWE-based ciphertext into two shares, assuming the ciphertext modulus has super-polynomial size. This trick has been used previously, including in the construction of AFS-spooky. Our main contribution is to bootstrap this 2-party non-interactive algorithm into an n -party non-interactive algorithm. We do this by placing the n parties on the leaves of a binary tree, and then homomorphically evaluating the two party protocol at each internal node of the tree. Each party only needs to evaluate the 2-party protocol at each node on the path from the root to its leaf. Each homomorphic evaluation at the internal nodes is exactly equivalent to a bootstrapping operation, namely a homomorphic evaluation of the decryption circuit for some key. Thus, decryption into shares costs $O(\log n)$ operations per party.

Removing the Super-polynomial Modulus. The problem with Scooby, as well as all LWE-based additive HSS schemes, is that we require a super-polynomial modulus-to-noise ratio in the underlying LD-based FHE scheme. This is a stronger form of LWE assumption that usually requires larger parameters to compensate. We give a variant of the construction where we only need standard FHE, together with an HSS scheme for NC1 circuits. Using recent constructions of HSS [OSY21, RS21, ADOS22] based on either Paillier encryption or class groups, we obtain the first additive HSS schemes for circuits that do not require LWE with a super-polynomial modulus. The complexity of the HSS is also $O(|F| + \log n)$. However, it is likely to be less efficient in practice than Scooby, and is also not secure against a quantum adversary. We call this construction **Scrappy**.

Avoiding FHE Entirely. We also show that in certain cases, we can obtain multi-party HSS without using any form of FHE whatsoever. We do this through a variant of the previous construction, where we bootstrap a HEDS scheme to handle more parties by homomorphically evaluating its own decryption circuit. This transformation is more challenging to apply without resorting to FHE, and we are only able to obtain a 4-party HEDS scheme for constant-degree polynomials,

⁴ In some sense the “spooky” behaviour exhibited by spooky encryption cannot really be explained, whereas our “spooky” behaviour can be explained by the setup procedure. This setup procedure in some sense acts like the janitor in Scooby-Do, who has set up the spooky goings-on.

Construction	Assumptions	Setup	Complexity
DHRW [DHRW16] (AFS-Spooky)	LWE with super-polynomial modulus	Uniform CRS	$O(n^2 \cdot F)$
Scooby: §5 (HEDS)	LD-based FHE with super-polynomial modulus	Trusted	$O(F + \log n)$
Scrappy: §6.1 (HEDS)	Generic FHE + 2-party HSS for NC1	Trusted	$O(F + \log n)$
Shaggy: §6.2 (HEDS)	2-party HSS for NC1	Trusted	$O(F)$ ($n = 4$, constant-deg F)
Casper: §7 (AFS-Spooky)	Specific MK-FHE with super-polynomial modulus	Uniform CRS	$O(n \cdot F + n^2 \cdot \log n)$ or $O(n^2 \cdot F + n^2 \cdot \log n)$

Table 1. Summary of n -party HSS Constructions. All FHE-based constructions allow arbitrary functions F , and assume circular security to avoid blow-up in the key sizes (this assumption can be removed by relaxing to bounded-depth circuits). The asymptotic complexities ignore potential factors in λ that are independent of n and F .

based on Paillier encryption. Nevertheless, as far as we are aware, this is the first instance of > 2 -party, dishonest majority HSS for constant-degree polynomials, without relying on FHE. We call this construction **Shaggy**.

We summarize our results in Table 1.

Spooky from HEDS. In addition, in Section 7, we show how our Scooby scheme can be adapted to give a true AFS-spooky encryption, i.e. with no trusted setup and independent keys, if we base our construction on *specific* multi-key FHE (MK-FHE). This instantiation can have a simpler complexity than that given in [DHRW16], in particular, assuming the function F can be evaluated without bootstrapping, our complexity is $O(n \cdot |F| + n^2 \cdot \log n)$. If F requires a bootstrapping for all the operations, it is $O(n^2 \cdot |F| + n^2 \cdot \log n)$. We call this construction **Casper**.

We give two variants of Casper, one based on the TFHE scheme [CCS19], and one based on the BFV scheme [CDKS19]. We note that being MK-FHE schemes, the construction will be less efficient than our Scooby scheme, which works over most (practical) FHE schemes. It is interesting to note that the spooky construction from [DHRW16] also goes via MK-FHE. In particular, they make use of the MK-FHE scheme of [CM15, MW16]. The route though is more complex than our tree-based construction, leading to an increased complexity.

2 Preliminaries

For a set S , we denote by $a \leftarrow S$ the process of drawing a from S with a uniform distribution on the set S . If D is a probability distribution, we denote by $a \leftarrow D$ the process of drawing a with the given probability distribution. For a probabilistic algorithm A , we denote by $a \leftarrow A$ the process of assigning a the output of algorithm A ; with the underlying probability distribution being determined by the random coins of A .

All reductions modulo an integer p will be assumed to be centred, i.e. in the interval $(-p/2, \dots, p/2)$.

We let $R = \mathbb{Z}[X]/(X^N + 1)$ and R_p denote the localisation of R at p , i.e. $(\mathbb{Z}/p\mathbb{Z})[X]/(X^N + 1)$. For a real interval I we let R_I denote the restriction of the set R to have coefficients in the support of I . Thus as sets (but not as rings) we have $R_q = R_{(-q/2, \dots, q/2)}$.

2.1 Homomorphic Secret Sharing

The following definition of public-key HSS is adapted from [BKS19]. Note that, as we are only interested in schemes with additive reconstruction, we can disregard the decoding algorithm, $\text{Dec}_{\text{sk}}^{\text{HSS}}$, that is given in the more general definition of HSS [BGI⁺18]. Concretely, in additive HSS the decoding algorithm simply adds up all the shares.

Definition 2.1 (Additive Public Key Homomorphic Secret Sharing). *An n -party, public-key homomorphic secret sharing (HSS) scheme for a class of functions \mathcal{F} over a ring \mathcal{R} with input space $\mathcal{I} \subseteq R$ consists of PPT algorithms $(\text{KeyGen}^{\text{HSS}}, \text{Share}_{\text{pk}}^{\text{HSS}}, \text{Eval}_{\text{pk}}^{\text{HSS}})$ with the following syntax:*

- $\text{KeyGen}^{\text{HSS}}(1^\lambda, n) \rightarrow (\text{pk}, (\text{ek}_1, \dots, \text{ek}_n))$: Given a security parameter 1^λ , the setup algorithm outputs a public key pk and n evaluation keys $(\text{ek}_1, \dots, \text{ek}_n)$.
- $\text{Share}_{\text{pk}}^{\text{HSS}}(\text{pk}, x) \rightarrow (\mathbf{x}_1, \dots, \mathbf{x}_n)$: Given public key pk and private input value $x \in \mathcal{I}$, the share algorithm outputs shares $(\mathbf{x}_1, \dots, \mathbf{x}_n)$.
- $\text{Eval}_{\text{pk}}^{\text{HSS}}(F; \mathbf{x}_i, \text{ek}_i) \rightarrow y_i$: On input a function $F \in \mathcal{F}$, the parties share \mathbf{x}_i , and it's evaluation key ek_i , the homomorphic evaluation algorithm outputs $y_i \in R$, which is party i 's share of an output $y \in R$.

This definition is in the multi-input setting, meaning that it supports a compact evaluation of a function F on shares of inputs $x^{(1)}, \dots, x^{(\rho)}$ given by ρ parties that are usually referred to as *clients*. More concretely, each client inputs $x^{(i)}$ to the Share algorithm which returns shares $\mathbf{x}_j^{(i)}, j \in [n]$, to n parties (the *servers*). Each server can then locally run Eval on input $(\mathbf{x}_j^{(1)}, \dots, \mathbf{x}_j^{(\rho)})$ obtaining a share y_j such that $F(x^{(1)}, \dots, x^{(\rho)}) = \sum_{j \in [n]} y_j$. Note that the $\text{KeyGen}^{\text{HSS}}$ algorithm cannot be run by any single party, so can be seen as a form of correlated randomness generated by a trusted dealer. We describe the required security properties for the algorithms $(\text{KeyGen}, \text{Share}, \text{Eval})$ according to this more general formulation.

Definition 2.2 (HSS (Statistical) Correctness). *We say that an n -party public-key HSS scheme $(\text{KeyGen}^{\text{HSS}}, \text{Share}_{\text{pk}}^{\text{HSS}}, \text{Eval}_{\text{pk}}^{\text{HSS}})$ is correct for a class of functions \mathcal{F} if, for all security parameters $\lambda \in \mathbb{N}$, for all functions $F \in \mathcal{F}$, for all $x^{(1)}, \dots, x^{(\rho)} \in \mathcal{I}$ (where \mathcal{I} is the input space of F), for all $(\text{pk}, \text{ek}_1, \dots, \text{ek}_n) \leftarrow \text{KeyGen}^{\text{HSS}}(1^\lambda)$ and for all $(\mathbf{x}_1^{(i)}, \dots, \mathbf{x}_n^{(i)}) \leftarrow \text{Share}_{\text{pk}}^{\text{HSS}}(\text{pk}, x^{(i)}), i \in [\rho]$, we have*

$$\Pr \left[y_1 + \dots + y_n = F(x^{(1)}, \dots, x^{(\rho)}) \right] \geq 1 - \text{negl}(\lambda),$$

where

$$y_j \leftarrow \text{Eval}_{\text{pk}}^{\text{HSS}}(F; (\mathbf{x}_j^{(1)}, \dots, \mathbf{x}_j^{(\rho)}), \text{ek}_j), j \in [n],$$

where the probability is taken over the random coins of $\text{KeyGen}^{\text{HSS}}, \text{Share}_{\text{pk}}^{\text{HSS}}$ and $\text{Eval}_{\text{pk}}^{\text{HSS}}$.

Definition 2.3 (HSS Security). *Let I be the set of corrupt servers. For each $j \in I$ and non-uniform adversary \mathcal{A} (of size polynomial in the security parameter λ), it holds that*

$$\left| \Pr[\text{Exp}_{\mathcal{A}, j}^{\text{HSS}, \text{sec}}(\lambda) = 1] \right| \leq \frac{1}{2} + \text{negl}(\lambda),$$

where $\text{Exp}_{\mathcal{A}, j}^{\text{HSS}, \text{sec}}(\lambda)$ is the experiment defined in Figure 1.

Security Experiment $\text{Exp}_{\mathcal{A},j}^{\text{HSS,sec}}(\lambda)$

Let $I \subset [n]$ be the set of corrupt servers.

1. $(\text{pk}, (\text{ek}_1, \dots, \text{ek}_n)) \leftarrow \text{KeyGen}^{\text{HSS}}(1^\lambda)$.
2. $(x_0, x_1, \text{state}) \leftarrow \mathcal{A}(1^\lambda)$.
3. $b \leftarrow \{0, 1\}$.
4. $(\mathbf{x}_{b,1}, \dots, \mathbf{x}_{b,n}) \leftarrow \text{Share}_{\text{pk}}^{\text{HSS}}(\text{pk}, x_b)$.
5. $b' \leftarrow \mathcal{A}(\text{state}, \text{pk}, \{\text{ek}_j, \mathbf{x}_{b,j}\}_{j \in I})$.
6. Return $b' = b$.

Fig. 1. Security Experiment $\text{Exp}_{\mathcal{A},j}^{\text{HSS,sec}}(\lambda)$

Remark 2.1 (Private-key HSS). HSS can also be defined in the single-input, private key setting, which is weaker than the public-key flavour above. Here, there is no `KeyGen` algorithm, and `Share` is run only once on all inputs together, so can be seen as a trusted dealer algorithm that distributes the shares.

2.2 Spooky Encryption

“Spooky” encryption is a type of public key encryption scheme which exhibits a form of limited malleability, so called “spooky action at a distance” [DHRW16]. The particular form of spooky encryption we will use is so called *additive-function-sharing* spooky encryption (or AFS-spooky encryption). We present a definition which works for any finite ring \mathcal{R} , and arithmetic circuit C , and not just for the case of \mathbb{F}_2 as originally presented.

Definition 2.4 (AFS-spooky Encryption). *An AFS-spooky encryption scheme, over a finite field \mathbb{F}_p , is a public-key encryption scheme given by a tuple of four algorithms $(\text{KeyGen}^{\text{Spooky}}, \text{Enc}_{\text{pk}}^{\text{Spooky}}, \text{Dec}_{\text{sk}}^{\text{Spooky}}, \text{Eval}_{\text{pk}_1, \dots, \text{pk}_n}^{\text{Spooky}})$ with the following syntax:*

- $\text{KeyGen}^{\text{Spooky}}(1^\lambda)$: *This is a probabilistic polynomial time algorithm which on input of a security parameter λ outputs a public/private key pair (pk, sk) .*
- $\text{Enc}_{\text{pk}}^{\text{Spooky}}(m)$: *This probabilistic polynomial time algorithm takes a message $m \in \mathcal{R}$ and generates a ciphertext ct encrypting that message under the public key pk .*
- $\text{Dec}_{\text{sk}}^{\text{Spooky}}(\text{ct})$: *Given a ciphertext ct encrypted under the public key associated to sk , this algorithm produces the underlying plaintext.*
- $\text{Eval}_{\text{pk}_1, \dots, \text{pk}_n}^{\text{Spooky}}(C, \text{ct}_1, \dots, \text{ct}_n)$: *Given an arithmetic circuit description $C : \mathcal{R}^n \rightarrow \mathcal{R}$, n public keys $\text{pk}_1, \dots, \text{pk}_n$, and n of ciphertexts $\text{ct}_1, \dots, \text{ct}_n$, this produces n ciphertexts $\text{ct}'_1, \dots, \text{ct}'_n$.*

An AFS-spooky encryption scheme must be correct, as an encryption scheme, i.e. we must have

$$\forall (\text{pk}, \text{sk}) \leftarrow \text{KeyGen}^{\text{Spooky}}(1^\lambda), \forall m \in \mathcal{R} : \text{Dec}_{\text{sk}}^{\text{Spooky}}(\text{Enc}_{\text{pk}}^{\text{Spooky}}(m)) = m.$$

It must also be IND-CPA as an encryption scheme and satisfy the following form of limited malleability called AFS-spooky correctness.

Definition 2.5 (AFS-spooky Correctness). *There exists a negligible function ν such that for all $\lambda \in \mathbb{N}$, every arithmetic circuit C computing a n -argument function $f : \mathcal{R}^n \rightarrow \mathcal{R}$, and all*

inputs x_1, \dots, x_n of C , we have

$$\Pr \left[\sum_{i \in [n]} y_i = C(x_1, \dots, x_n) : \begin{array}{l} \forall i \in [n], (\mathbf{pk}_i, \mathbf{sk}_i) \leftarrow \text{KeyGen}^{\text{Spooky}}(1^\lambda), \\ \forall i \in [n], \text{ct}_i \leftarrow \text{Enc}_{\mathbf{pk}}^{\text{Spooky}}(x_i), \\ (\text{ct}'_1, \dots, \text{ct}'_n) \leftarrow \text{Eval}_{\mathbf{pk}_1, \dots, \mathbf{pk}_n}^{\text{Spooky}}(C, \text{ct}_1, \dots, \text{ct}_n), \\ \forall i \in [n], y_i \leftarrow \text{Dec}_{\mathbf{sk}_i}^{\text{Spooky}}(\text{ct}'_i) \end{array} \right] \geq 1 - \nu(\lambda)$$

In [DHRW16], it is shown how to construct an AFS-spooky encryption scheme in the CRS model using an LWE-based multi-key FHE [CM15, MW16] and assuming a circular security assumption. The common reference string (output by a separate generation algorithm), necessary in the multi-key FHE construction, is assumed as input to the key generation algorithm, and correctness and security hold for all outputs of the common reference string generator.

In their work, Dodis et al. [DHRW16] show that AFS-spooky encryption implies FSS for general circuit; in [BGI⁺18], Boyle et al. show that AFS-spooky also enables HSS for multiple inputs; in fact, it implies HSS without any setup, where the key generation algorithm is simply run locally by each client providing input.

3 Homomorphic Encryption with Decryption to Shares (HEDS)

In this section we formally introduce the notion of a scheme which implements Homomorphic Encryption with Decryption to Shares (HEDS) and relate it with other concepts described in previous sections. Loosely speaking, a HEDS encryption scheme is similar to public-key HSS, except with a public evaluation algorithm that outputs a ciphertext, more akin to evaluation in homomorphic encryption. The ciphertext is then converted into shares in the decryption algorithm, which uses one party's private key. In addition, similarly to HSS, but unlike in spooky encryption, the parties need to engage in a protocol, or assume a trusted third party, to set up the associated public and secret keys. Thus the action from the outside seems spooky, but this can be explained away as an effect of the setup protocol.

We start by giving the definition of HEDS, and then we show that it enables both homomorphic and function secret sharing.

Definition 3.1 (HEDS Encryption). A HEDS encryption scheme for a class of functions $\mathcal{F} : \mathcal{R}^* \rightarrow \mathcal{R}$, over a ring \mathcal{R} , is given by a tuple of PPT algorithms $(\text{Setup}^{\text{HEDS}}, \text{Enc}_{\mathbf{pk}}^{\text{HEDS}}, \text{Dec}_{\mathbf{sk}}^{\text{HEDS}}, \text{Eval}_{\mathbf{pk}}^{\text{HEDS}})$, with the following syntax:

- $\text{Setup}^{\text{HEDS}}(1^\lambda, n)$: This randomized algorithm takes as input a security parameter λ , a number of parties n . It outputs the tuple $(\mathbf{pk}, \mathbf{sk}_1, \dots, \mathbf{sk}_n)$.
- $\text{Enc}_{\mathbf{pk}}^{\text{HEDS}}(m)$: This takes as input the public key and a message $m \in \mathcal{R}$, and outputs a ciphertext ct .
- $\text{Dec}_{\mathbf{sk}_i}^{\text{HEDS}}(\text{ct})$: Given a ciphertext ct encrypted under the public key this outputs a value y_i for each $i \in [n]$.
- $\text{Eval}_{\mathbf{pk}}^{\text{HEDS}}(C, (\text{ct}_1, \dots, \text{ct}_\rho))$: On input of the public key \mathbf{pk} , a set of n ciphertexts, and an arithmetic circuit description $C : \mathcal{R}^\rho \rightarrow \mathcal{R}$ of a function from the specified class, this produces a ciphertext ct .

The algorithms $(\text{Setup}^{\text{HEDS}}, \text{Enc}_{\mathbf{pk}}^{\text{HEDS}}, \text{Dec}_{\mathbf{sk}}^{\text{HEDS}}, \text{Eval}_{\mathbf{pk}}^{\text{HEDS}})$ should satisfy the following correctness and security requirements.

Definition 3.2 (HEDS Correctness). *There exists a negligible function ν such that for all $\lambda \in \mathbb{N}$, every arithmetic circuit C computing a ρ -argument function $f : \mathcal{R}^\rho \rightarrow \mathcal{R}$ in \mathcal{F} , and all inputs x_1, \dots, x_ρ of C , we have*

$$\Pr \left[\sum_{i \in [n]} y_i = C(x_1, \dots, x_\rho) : \begin{array}{l} (\text{pk}, \text{sk}_1, \dots, \text{sk}_n) \leftarrow \text{Setup}^{\text{HEDS}}(1^\lambda, n), \\ \forall i \in [\rho], \text{ct}_i \leftarrow \text{Enc}_{\text{pk}}^{\text{HEDS}}(x_i), \\ \text{ct} \leftarrow \text{Eval}_{\text{pk}}^{\text{HEDS}}(C, (\text{ct}_1, \dots, \text{ct}_\rho)), \\ \forall i \in [n], y_i \leftarrow \text{Dec}_{\text{sk}_i}^{\text{HEDS}}(\text{ct}) \end{array} \right] \geq 1 - \nu(\lambda).$$

Definition 3.3 (HEDS Security). *For all subsets $A \subset [n]$ of size $< n$, and all probabilistic polynomial time adversaries $(\mathcal{A}_1, \mathcal{A}_2)$ we have*

$$\Pr \left[b = b' : \begin{array}{l} (\text{pk}, \text{sk}_1, \dots, \text{sk}_n) \leftarrow \text{Setup}^{\text{HEDS}}(1^\lambda, n), b \in \{0, 1\}, \\ (m_0, m_1, \text{state}) \leftarrow \mathcal{A}_1(\text{pk}, \{\text{sk}_i\}_{i \in A}), \\ \text{ct} \leftarrow \text{Enc}_{\text{pk}}^{\text{HEDS}}(m_b), \\ b' \leftarrow \mathcal{A}_2(\text{ct}, \text{state}) \end{array} \right] \leq \text{negl}(\lambda),$$

i.e. the encryption scheme is IND-CPA, even when up to $n-1$ secret keys are given to the adversary.

Compactness. Just as with fully homomorphic encryption, we say that HEDS is *compact* if the share decryption algorithm is independent of the evaluated function.

3.1 Multi-input HSS from HEDS Encryption

Here we relate HEDS encryption and HSS showing that HEDS encryption implies HSS with multiple inputs. Let \mathcal{P} be a set of n servers and \mathcal{C} be a set of m clients. Let C be a circuit representing a function $F : \mathcal{R}^m \rightarrow \mathcal{R}$ in a class function \mathcal{F} . To build an HSS-scheme, we need to define three algorithms $\text{KeyGen}^{\text{HSS}}$, $\text{Share}^{\text{HSS}}$, Eval^{HSS} as in Definition 2.1. Let $(\text{Setup}^{\text{HEDS}}, \text{Enc}_{\text{pk}}^{\text{HEDS}}, \text{Dec}_{\text{sk}}^{\text{HEDS}}, \text{Eval}_{\text{pk}}^{\text{HEDS}})$ be a HEDS encryption scheme for \mathcal{F} , as defined in the previous section, we proceed as follows.

- $\text{KeyGen}^{\text{HSS}}(1^\lambda, n)$:
 1. Run $(\text{pk}, \text{sk}_1, \dots, \text{sk}_n) \leftarrow \text{Setup}^{\text{HEDS}}(1^\lambda, n)$
 2. For each $i \in [n]$, set $\text{ek}_i := \text{sk}_i$
 3. Return pk and $(\text{ek}_1, \dots, \text{ek}_n)$
- $\text{Share}_{\text{pk}}^{\text{HSS}}(x^{(j)})$: Each client $P_j \in \mathcal{P}$, on input $x^{(j)}$ performs the following steps. We recall that the goal is to obtain shares $(\mathbf{x}_1, \dots, \mathbf{x}_n)$ of (x^1, \dots, x^m) .
 1. For $i \in [n]$ and $j \in [m]$, generate $x_i^{(j)}$ such that $x^{(j)} = x_1^{(j)} + \dots + x_n^{(j)}$.
 2. For each $x_i^{(j)}$, compute $\text{ct}_i^{(j)} = \text{Enc}_{\text{pk}}^{\text{HEDS}}(x_i^{(j)})$.
 3. Set $\mathbf{x}_i = \{\text{ct}_i^{(j)}\}_{j \in [m]}$, for $i \in [n]$.
- $\text{Eval}_{\text{pk}}^{\text{HSS}}(F; \mathbf{x}_i, \text{ek}_i)$: Given a function $F : \mathcal{R}^m \rightarrow \mathcal{R}$, each server $i \in [n]$ computes circuit description C of F and proceeds as follows.
 1. Compute $\text{ct}_i = \text{Eval}_{\text{pk}}^{\text{HEDS}}(C, (\text{ct}_i^{(1)}, \dots, \text{ct}_i^{(m)}))$
 2. Compute $y_i = \text{Dec}_{\text{ek}_i}^{\text{HEDS}}(\text{ct}_i)$

By Definition 3.2, we know that the evaluation algorithm outputs to the servers the shares y_1, \dots, y_n such that $\sum_{i \in [n]} y_i = y = F(x^{(1)}, \dots, x^{(m)})$.

Proposition 3.1. *Assuming the existence of a HEDS encryption scheme for a class of functions \mathcal{F} , there exists a public-key multi-input HSS scheme for \mathcal{F} .*

Proof. Correctness follows by inspection of the scheme described above and by correctness of the underlying HEDS construction. Security also follows from the security of HEDS. \square

In the other direction, we observe that a public-key HSS scheme implies HEDS for the same class of functions, however, the resulting HEDS scheme may not be compact. This is because the HSS evaluation algorithm will have to be carried out in the HEDS decryption step, since HSS uses a private key for evaluation.

4 Linear-Decryption Based FHE

Our main constructions are based on a form of FHE which comes from LWE-style systems. We abstract much of the details of the specific construction away in what follows, for example the specific key generation and encryption algorithms. This allows us to capture schemes as diverse as BGV [BGV12], BFV [FV12], GSW [GSW13] and TFHE [CGGI16, CGGI20]. These schemes all have the same form of decryption equation, namely one based on a linear inner product combination of the ciphertext with the secret key, modulo the ciphertext modulus. The result of this inner product is then processed to produce the plaintext (which is an element of R_p for some prime p) in one of two distinct ways, depending on whether the message is embedded at the top of the range modulo q (as in FV), or the bottom of the range modulo q (as in BGV). We refer to these two types of decryption as FHE as being of type `msb` and type `lsb` respectively. We call the whole class of such FHE systems Linear Decryption based, or LD-based FHE. Similar definitions have been considered previously [BKS19, GH19, BDGM19].

Let `sec` denote some statistical security parameter and λ denote a computational security parameter. We define such a scheme as follows, the precise encryption and evaluation algorithms are not important for our discussion.

Definition 4.1 (LD-based FHE). *An LD-based FHE scheme is given by a tuple of algorithms $(\text{KeyGen}^{\text{FHE}}, \text{Enc}_{\text{pk}}^{\text{FHE}}, \text{Dec}_{\text{sk}}^{\text{FHE}}, \text{Eval}_{\text{pk}}^{\text{FHE}})$, as follows:*

- $\text{KeyGen}^{\text{FHE}}(1^\lambda, p)$: *This randomized algorithm takes as input the security parameter λ and a plaintext modulus space p . It outputs a tuple $(q, N, B, d, \Delta, S, \text{pk}, \text{sk})$. The value q will correspond to the ciphertext modulus⁵, the value N will be the LWE-ring dimension (which for convenience we assume is a power of two), the value B will be a “noise bound”, the value d is one less than the dimension of the ciphertext space, the value Δ is set to be $\lfloor q/p \rfloor$, the value S is a bound on the secret key size S , and pk (resp. sk) will be the public (resp. private) keys.*

The private key $\text{sk} = (s_1, \dots, s_d)$ is assumed to be a random element in R_q^d sampled such that $\|\text{sk}\|_\infty \leq S$. Note, this is not necessarily sampled uniformly at random subject to this constraint. All subsequent algorithms are assumed to take the tuple (N, q, d, B, Δ) implicitly as input parameters.

⁵ In practice there may be many ciphertext moduli depending on which level a ciphertext is sitting at, at a high level this can be ignored. Although it can be important in practice

- $\text{Enc}_{\text{pk}}^{\text{FHE}}(m, \text{type})$: On input of $m \in R_p$ this will output a ciphertext $\text{ct} \in R_q^{d+1}$ such that

$$\text{ct} \cdot (1, -\text{sk}) = \begin{cases} m + p \cdot \epsilon \pmod{q} & \text{If type} = \text{lsb}, \\ \Delta \cdot m + \epsilon \pmod{q} & \text{If type} = \text{msb}. \end{cases}$$

A ciphertext such that $\|\epsilon\|_\infty \leq B$ will be called valid. The encryption algorithm produces such a valid ciphertext. The precise algorithm use for encryption will depend on the public key, and the specific scheme. All that concerns us is the form of the ciphertext.

- $\text{Eval}_{\text{pk}}^{\text{FHE}}(F(x_1, \dots, x_\ell), \{\text{ct}_1, \dots, \text{ct}_\ell\})$: On input of ℓ valid ciphertexts ct_i and an arithmetic function $F(x_1, \dots, x_\ell)$ this function will homomorphically evaluate the function F over the ciphertexts, producing a valid ciphertext as output.
- $\text{Dec}_{\text{sk}}^{\text{FHE}}(\text{ct})$: On input of a valid ciphertext and a secret key this will compute the message as

$$m = \begin{cases} (\text{ct} \cdot (1, -\text{sk}) \pmod{q}) \pmod{p} & \text{If type} = \text{lsb}, \\ \lfloor (\text{ct} \cdot (1, -\text{sk}) \pmod{q}) \cdot p/q \rfloor & \text{If type} = \text{msb}. \end{cases}$$

The correctness requirement simply says that Eval^{FHE} , when given ℓ valid ciphertexts, outputs a valid encryption of the correct result. The security requirement is the standard notion of IND-CPA security.

For example: In the case of the BGV scheme [BGV12] from ring-LWE we will have that $\text{ct} = (c_0, c_1)$, so that decryption is given by $\text{ct} \cdot (1, -\text{sk}) = c_0 - s_1 \cdot c_1$, and, hence, for this scheme we have $n = 1$ and $\text{sk} = s_1$. The BFV scheme [FV12] has the same structure, the main difference being that BFV uses the `msb` decryption, while BGV uses `lsb`.

In the case of Ring-GSW, a ciphertext is in $R_q^{(d+1) \times (d+1)\ell}$, with $d = 1$ and $\text{sk} = s_1$. In practice it is composed by $2 \cdot \ell$ FV-like ciphertext (i.e., with the message encrypted in the `msb`). To decrypt a Ring-GSW ciphertext, we only decrypt one of these ciphertexts: the others contain redundant information. Another way of seeing a Ring-GSW ciphertext, is with a very sparse secret key $\text{sk} = (\text{sk}_1, \dots, \text{sk}_{2\ell})$, where all the keys corresponding to the FV-like ciphertext that we are not going to decrypt are set to zero. The TFHE scheme [CGGI16, CGGI20] uses a combination of FV-like ciphertexts (with message encrypted in the `msb`, called LWE and RLWE ciphertexts) and Ring-GSW ones.

Parameters for Decryption to Shares. For such LD-based FHE schemes we have a special form of non-interactive two party distributed decryption, which we shall now outline in the `lsb` and the `msb` cases. We will require the parameters are selected so that

$$q > 2 \cdot p \cdot (B + 1) \cdot 2^{\text{sec}}, \tag{1}$$

where `sec` is the statistical security parameter. This two-party distributed decryption, which is essentially the same technique as in [DHRW16, BKS19], will form the basis of our first multi-party HEDS construction in Section 5.

4.1 Two-Party Distributed Decryption: Type `lsb`

Suppose sk is split into two keys sk_1 and sk_2 with $\text{sk} = \text{sk}_1 + \text{sk}_2 \pmod{q}$, with sk_1 held by party P_1 and sk_2 held by party P_2 . Now we can, without interaction, given a valid ciphertext ct encrypting a message m , compute an additive sharing of $m = m_1 + m_2 \pmod{p}$ between P_1 and P_2 as follows. We require that the parties have agreed upon a public random value for each decryption, but later will remove this using a PRF.

2-party DistDec^{lsb}: Let $R \leftarrow \mathbb{Z}_q$ be a public random nonce.

1. P_1 computes $d_1 \leftarrow \text{ct} \cdot (1, -\text{sk}_1) + R \pmod{q}$ and then $m_1 \leftarrow d_1 \pmod{p}$.
2. P_2 computes $d_2 \leftarrow \text{ct} \cdot (0, -\text{sk}_2) - R \pmod{q}$ and then $m_2 \leftarrow d_2 \pmod{p}$.

We prove that this leads to a correct result with overwhelming probability.

Proposition 4.1. *Given an LD-based FHE scheme of type msb (Definition 4.1), where $(q, N, B, d, \Delta, S, \text{pk}, \text{sk}) \leftarrow \text{KeyGen}^{\text{FHE}}(1^\lambda, p)$, with $q > 2 \cdot p \cdot (B + 1) \cdot 2^{\text{sec}}$ and $\text{sk}_1 + \text{sk}_2 = \text{sk}$. Let (ct, m) be a pair of ciphertext/plaintext messages and m_1 and m_2 values obtained with the 2-party distributed decryption procedure described above. Then, it holds that*

$$m = m_1 + m_2 \pmod{p},$$

with probability at least $1 - N \cdot 2^{-\text{sec}}$.

Proof. First we notice that

$$m = ((d_1 + d_2) \pmod{q}) \pmod{p},$$

and that we will always have $m = m_1 + m_2 \pmod{p}$ if the internal reduction modulo q in the decryption equation for m does not need to compensate for a wrap around. However, since we know ct is valid (i.e., $\text{ct} \cdot (1, -\text{sk}) = m + p \cdot \epsilon \pmod{q}$ with $\|\epsilon\|_\infty \leq B$) we also know that the coefficients of $d_1 + d_2 \pmod{q}$ will lie in the range $(-p \cdot (B + 1), \dots, p \cdot (B + 1))$. Thus, the distributed decryption will potentially result in an error if and only if the coefficients of d_1 lie in one of the two ranges $(-q/2, -q/2 + p \cdot (B + 1))$ or $(q/2 - p \cdot (B + 1), q/2)$. Since each party added or subtracted the random R , it holds that d_1 is uniformly distributed in the range $(-q/2, \dots, q/2)$. Therefore, the probability there is a wraparound in a single coefficient is bounded by $2 \cdot p \cdot (B + 1)/q < 2^{-\text{sec}}$. However, we also know that, if there is a wrap around, it will definitely result in an invalid distributed decryption, as the error only consists of the addition of a single value of $q \pmod{p} \neq 0$. Thus, a single coefficient will be correct with probability $1 - 2^{-\text{sec}}$. To obtain a correct decryption we need all coefficients to be correct, which will happen with probability

$$(1 - 2^{-\text{sec}})^N \approx 1 - N \cdot 2^{-\text{sec}}.$$

□

We report details on the two party distributed decryption for the type msb in Appendix A.

5 Scooby: Multi-Party HEDS from LD-based FHE

In this section we detail how to construct a HEDS encryption scheme for the underlying ring R_p , from generic LD-based FHE. We call our construction **Scooby**, as it is similar to a spooky encryption but with a trusted setup. To denote the specific nature of this construction we refer to $\text{Setup}^{\text{Scooby}}$, $\text{Enc}_{\text{pk}}^{\text{Scooby}}$, etc., instead of $\text{Setup}^{\text{HEDS}}$, $\text{Enc}_{\text{pk}}^{\text{HEDS}}$, etc.

At the core of **Scooby** is the 2-party distributed decryption procedure described in the previous section. We show that, assuming an LD-based FHE scheme, this directly yields a 2-party **Scooby**. We then show how to bootstrap the 2-party scheme to the multi-party setting.

5.1 HEDS Key Generation

First, we need to slightly modify the `KeyGen` algorithm for the underlying FHE scheme to take a “special” form that is common to all standard FHE constructions. More concretely, the algorithm $\text{KeyGen}^{\text{FHE}}(1^\lambda, p)$ proceeds as follows, using two sub-procedures `ParamGen()` and `PubKeyGen()`:

1. $\text{params} \leftarrow \text{ParamGen}(1^\lambda, p)$: This algorithm takes as input a security parameter λ , a plaintext modulo p and produces the scheme parameters $\text{params} = (q, N, B, d, \Delta, S)$.
2. $\text{sk} \leftarrow R_q^n$ such that $\|\text{sk}\|_\infty \leq S$.
3. $\text{pk} \leftarrow \text{PubKeyGen}(1^\lambda, \text{sk}, \text{params})$: This algorithm, on input the secret key and scheme parameters, samples and outputs an associated public key pk .

5.2 Security Assumption

In our construction, we generate an FHE public key based on a secret-key $\text{sk} = \text{sk}_0 + \text{sk}_1$, where sk_0, sk_1 are both sampled uniformly with coefficients bounded by the parameter S . For security, we require that the scheme defined by (pk, sk) satisfies the standard IND-CPA security notion, even when the adversary is given one of the original secret keys sk_i . This is formalized as follows.

Definition 5.1 (Bounded secret key IND-CPA security). *Let $\text{FHE} = (\text{KeyGen}^{\text{FHE}}, \text{Enc}_{\text{pk}}^{\text{FHE}}, \text{Dec}_{\text{sk}}^{\text{FHE}}, \text{Eval}_{\text{pk}}^{\text{FHE}})$ be a linear decryption-based FHE scheme, where $\text{KeyGen}^{\text{FHE}}$ is split into two sub-routines `ParamGen`, `PubKeyGen` as above.*

We require that for $(q, N, B, d, \Delta, S) \leftarrow \text{ParamGen}(1^\lambda, p)$, and $\text{sk}_0, \text{sk}_1 \leftarrow R_q^d$ with $\|\text{sk}_i\| \leq S$, $\text{sk} = \text{sk}_0 + \text{sk}_1$ and $\text{pk} \leftarrow \text{PubKeyGen}(\text{sk})$, it holds that for any PPT algorithm \mathcal{A} , for any $\sigma \in \{0, 1\}$, messages m_0, m_1 and bit $b \leftarrow \{0, 1\}$:

$$\Pr[\mathcal{A}(1^\lambda, \text{pk}, \text{sk}_\sigma, \text{Enc}_{\text{pk}}^{\text{FHE}}(m_b)) = b] \leq 1/2 + \text{negl}(\lambda).$$

It is straightforward to verify that, given a linear decryption-based FHE scheme that satisfies the bounded secret-key IND-CPA security, we obtain a 2-party `Scooby` encryption scheme using the prior algorithms for 2-party distributed decryption into shares described in the previous section. Indeed, this 2-party distributed decryption forms the basis of the 2-party spooky construction in [DHRW16] and HSS construction in [BKS19]. However, to obtain an n -party generalization is not immediate. A direct application of the trick used for 2-party to, say, 3-parties results in decryption errors due to unaccounted for wrap-arounds in the reduction modulo q of the local decryption. Coping with these wrap-arounds, without resorting to interaction, thus seems a challenge. A challenge which we solve in the next section.

5.3 From 2-party to n -party HEDS

Here we give the details of our construction `Scooby`, for n -party HEDS. The encryption and evaluation algorithms of `Scooby` are identical to that of the underlying linear decryption FHE scheme, so here we only describe the setup and share decryption procedures. We give two different variants of the construction, depending on whether the FHE scheme encodes the message in the `lsb` or `msb` of the ciphertext. In this section, we focus on a linear decryption FHE scheme that encodes the message in the `lsb` of the ciphertext; in Appendix A.1, we give a variant for the `msb` type.

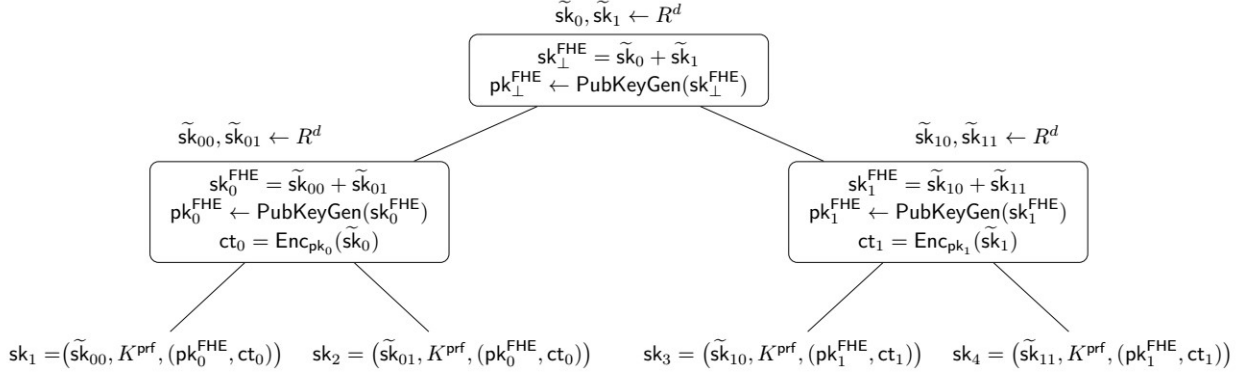


Fig. 2. Scooby setup for $n = 4$

Scooby Setup. Recall that the setup algorithm in HEDS takes as input a security parameter and outputs a global public key pk , as well as secret keys sk_1, \dots, sk_n to each of the n parties. For Scooby, in both the *lsb* and *msb* variants of LD-based FHE scheme, the underlying $\text{SetUp}^{\text{Scooby}}$ algorithm is the same. Note that in the following, the $\text{SetUp}^{\text{Scooby}}$ algorithm should be seen as a trusted setup procedure that is either run by a trusted third party, or executed via an MPC protocol, which can be done, for instance, based on the techniques from [?].

The $\text{SetUp}^{\text{Scooby}}$ algorithm is described in Figure 3. Recall that the main challenge is to setup up some key material which allows n parties to convert an FHE ciphertext into shares of the message, while using the 2-party distributed decryption method from the previous section. We build a binary tree with n leaves, where the original FHE ciphertext lives at the root node. We split the FHE secret key sk^{FHE} into two shares $\tilde{sk}_0, \tilde{sk}_1$, and then generate a fresh FHE key pair for each of the two child nodes, and encrypt each \tilde{sk}_b , for $b \in \{0, 1\}$, under the corresponding public key. This process is repeated with the FHE secret keys generated for the children, and so on throughout the tree. Note that we abuse notation by writing $ct_v = \text{Enc}_{pk_v}^{FHE}(\tilde{sk}_v)$, even though \tilde{sk}_v may not lie in the plaintext space; we implicitly assume here that \tilde{sk}_v is broken up into bits (or possibly larger chunks), so ct_v is actually a vector of ciphertexts encrypting each bit separately.

The idea is that, during the decryption phase, the parties can homomorphically evaluate the 2-party distributed decryption function at each node of the tree, obtaining a share of the message, now encrypted under a child node's public key. The i -th party repeats this for each node on the path to leaf i , where it finally obtains a ciphertext encrypting an n -party sharing of the original message, which it can decrypt.

Given this setup procedure we define $\text{Enc}_{pk}^{\text{Scooby}}$ and $\text{Eval}_{pk}^{\text{Scooby}}$ exactly as is the case in the underlying LD-based FHE scheme. Next, we detail the $\text{Dec}_{sk_i}^{\text{Scooby}}$ procedure in the *lsb* case.

Scooby Decryption. The decryption algorithms for Scooby in the *lsb/msb*-mode are described in Figure 4 and Figure 11, respectively. The decryption algorithm requires $\lceil \log n \rceil - 1$ evaluations of the Eval^{FHE} function for the underlying LD-based FHE scheme, each for a different public key. Note that the circuit used in Eval^{FHE} is almost exactly the decryption circuit, so the complexity of each of these homomorphic operations is the same as a bootstrapping operation in the underlying FHE scheme.

Algorithm Setup^{Scooby}(λ, p, n)

The algorithm takes as input the security parameter λ , plaintext modulus p , and number of parties n . It outputs a public key \mathbf{pk} and secret keys $(\mathbf{sk}_1, \dots, \mathbf{sk}_n)$.

1. Let $\text{params} = (q, N, B, d, \Delta, S) \leftarrow \text{ParamGen}(1^\lambda, p)$.
2. Sample a key $K^{\text{prf}} \leftarrow \{0, 1\}^\lambda$.
3. We construct a complete (but not necessarily full at the last layer) binary tree with n leaves and height $h = \lceil \log(n) \rceil$, and index the levels from 0 up to h . Each node in level i of the tree is labelled with a string of i bits, so the root is the empty string \perp , and the children of node v are $v\|0$ and $v\|1$.
4. Sample $\tilde{\mathbf{sk}}_0, \tilde{\mathbf{sk}}_1 \leftarrow R_q^d$ such that $\|\tilde{\mathbf{sk}}_j\|_\infty \leq S$.
5. Let $\mathbf{sk}_\perp^{\text{FHE}} = \tilde{\mathbf{sk}}_0 + \tilde{\mathbf{sk}}_1$ and sample $\mathbf{pk}_\perp^{\text{FHE}} = \text{PubKeyGen}(1^\lambda, \mathbf{sk}_\perp^{\text{FHE}}, \text{params})$.
6. For each internal node v (excluding the root and leaves) with children $v\|0$ and $v\|1$:
 - (a) Sample $\tilde{\mathbf{sk}}_{v\|0}, \tilde{\mathbf{sk}}_{v\|1} \leftarrow R_q^d$ such that $\|\tilde{\mathbf{sk}}_j\|_\infty \leq S$.
 - (b) Let $\mathbf{sk}_v^{\text{FHE}} = \tilde{\mathbf{sk}}_{v\|0} + \tilde{\mathbf{sk}}_{v\|1}$, sample $\mathbf{pk}_v^{\text{FHE}} = \text{PubKeyGen}(1^\lambda, \mathbf{sk}_v^{\text{FHE}}, \text{params})$.
 - (c) Let $\text{ct}_v = \text{Enc}_{\mathbf{pk}_v^{\text{FHE}}}(\tilde{\mathbf{sk}}_v)$.
7. Let \mathbf{sk}_i contain the leaf secret key $\tilde{\mathbf{sk}}_i$, together with K^{prf} and the public keys and ciphertexts on the path from the root to leaf i .
8. Output $\mathbf{pk} := \mathbf{pk}_\perp^{\text{FHE}}$ and the secret keys $(\mathbf{sk}_1, \dots, \mathbf{sk}_n)$.

Fig. 3. Trusted setup algorithm for the Scooby construction

Algorithm Dec _{\mathbf{sk}_i} ^{Scooby}(ct) (for lsb-based construction)

Let $F : \{0, 1\}^\lambda \times [n] \rightarrow R_q$ be a pseudorandom function.

Dec _{\mathbf{sk}_i} ^{Scooby}(ct):

1. Parse \mathbf{sk}_i as $\tilde{\mathbf{sk}}_i, K^{\text{prf}}$ and $(\mathbf{pk}_v^{\text{FHE}}, \text{ct}_v)$, for every node v from the root to leaf i .
2. Let $\tilde{\mathbf{ct}}_\perp := \text{ct}$.
3. For each internal node v on the path from the root to leaf i (excluding the root and leaf):
 - (a) Write $v = u\|b$, where u is the parent of v (so $b = 0$ if v is a left child and $b = 1$ otherwise).
 - (b) Define the function:

$$f_{\tilde{\mathbf{ct}}_u}^b : \mathbf{sk} \mapsto \left(\tilde{\mathbf{ct}}_u \cdot (b, -\mathbf{sk}) + (-1)^b \cdot F(K^{\text{prf}}, u) \pmod{q} \right) \pmod{p}$$

- (c) Compute $\tilde{\mathbf{ct}}_v := \text{Eval}_{\mathbf{pk}_v^{\text{FHE}}}(f_{\tilde{\mathbf{ct}}_u}^b, \text{ct}_v)$.

4. Write $i = u\|b$, then take the leaf ciphertext $\tilde{\mathbf{ct}}_i$ and output the share

$$m_i = \left(\tilde{\mathbf{ct}}_i \cdot (b, -\tilde{\mathbf{sk}}_i) + (-1)^b \cdot F(K^{\text{prf}}, u) \pmod{q} \right) \pmod{p}$$

Fig. 4. Decryption to shares for lsb-based Scooby

It is also clear that, due to the fact that at each internal branch we are homomorphically evaluating the two-party distributed decryption method from either Section 4.1 (for the lsb case) or Appendix A (for the msb case), the final n messages m_i will sum up to the decryption of the ciphertext ct . The only difference is that instead of adding or subtracting a random nonce R , the parties are using the PRF F to randomize their shares in distributed decryption; thus, the correctness property of the scheme relies on the security of F .

Theorem 5.1. *Let F be a pseudorandom function, and suppose there is an LD-like FHE scheme which satisfies the hardness assumption from Definition 5.1, such that $(q, N, B, d, \Delta, S) \leftarrow \text{ParamGen}(1^\lambda, p)$ with $q > 2 \cdot p \cdot (B + 1) \cdot 2^{\text{sec}}$. Then the Scooby construction in Fig. 3–4 is a secure n -party homomorphic encryption scheme with decryption to shares.*

Proof. Let n be the number of parties. Let $(\text{pk}, \text{sk}_1, \dots, \text{sk}_n) \leftarrow \text{SetUp}^{\text{Scooby}}(\lambda, p, n)$. We need to show that Scooby, with the setup and decryption algorithms described in Figures 3, 4 and 11, satisfies the definition of correctness and security given in Section 3.

CORRECTNESS. Here, we only consider the lsb-based construction, the proof for the msb-based variant is similar.

We assume $(\text{pk}, \text{sk}_1, \dots, \text{sk}_n)$ be the output of the setup algorithm. Each party P_i holds sk_i , corresponding to a leaf i in the binary tree generated in the setup algorithm. We say that a party P_i is *involved in a node \mathcal{N}* of this binary tree if the path from the leaf i to the root contains the node \mathcal{N} .

Let ct be a ciphertext corresponding to public-key pk and encrypting a message m (possibly obtained through $\text{Eval}_{\text{pk}}^{\text{Scooby}}$), we show that $m_i \leftarrow \text{Dec}_{\text{sk}_i}^{\text{Scooby}}(\text{ct}), i \in [n]$, satisfy $\sum_i m_i = m$, except with negligible probability.

During the decryption algorithm, and starting from the root, each party P_i , using its secret key sk_i , iteratively performs homomorphic evaluations of the 2-party distributed decryption function until it reaches the leaf i .

More concretely, in the first step at the root level, we have $\tilde{\text{ct}} := \text{ct}$ and the 2-party decryption functions:

$$\begin{aligned} f_{\tilde{\text{ct}}}^0 : \text{sk} &\mapsto \left(\tilde{\text{ct}} \cdot (0, -\text{sk}) + F(K^{\text{prf}}) \pmod{q} \right) \pmod{p} \\ f_{\tilde{\text{ct}}}^1 : \text{sk} &\mapsto \left(\tilde{\text{ct}} \cdot (1, -\text{sk}) - F(K^{\text{prf}}) \pmod{q} \right) \pmod{p}. \end{aligned}$$

Going a level down in the tree, each party P_i is involved either in the node $v = 0$ or in $v = 1$, hence it will compute either

$$\tilde{\text{ct}}_0 = \text{Eval}_{\text{pk}_0}^{\text{FHE}}(f_{\tilde{\text{ct}}}^0, \text{ct}_0) = \text{Eval}_{\text{pk}_0}^{\text{FHE}}(f_{\tilde{\text{ct}}}^0, \text{Enc}_{\text{pk}_0}^{\text{FHE}}(\tilde{\text{sk}}_0))$$

or

$$\tilde{\text{ct}}_1 = \text{Eval}_{\text{pk}_1}^{\text{FHE}}(f_{\tilde{\text{ct}}}^1, \text{ct}_1) = \text{Eval}_{\text{pk}_1}^{\text{FHE}}(f_{\tilde{\text{ct}}}^1, \text{Enc}_{\text{pk}_1}^{\text{FHE}}(\tilde{\text{sk}}_0))$$

Since $\tilde{\text{ct}}$ is an encryption under the public key pk corresponding to $\text{sk} = \tilde{\text{sk}}_0 + \tilde{\text{sk}}_1$, the two ciphertexts $\tilde{\text{ct}}_0, \tilde{\text{ct}}_1$ are the result of an homomorphic evaluation of the 2-party decryption functions $f_{\tilde{\text{ct}}}^0, f_{\tilde{\text{ct}}}^1$ and represent encryptions of m_0^1 and m_1^1 under pk_0 and pk_1 , respectively, such that $m_0^1 + m_1^1 = m$.

Further following the path towards the leaf, each party will repeat the same procedure for at most one node per level, always obtaining an encryption of a share of m that they are not able to decrypt until they reach the level just above the leaves. In the last step, each P_i considers the leaf ciphertext and, performing a 2-party distributed decryption, outputs a share of m .

The correctness of the procedure directly follows from the fact that all the ciphertexts are randomized using the PRF and by Proposition 4.1. Each party involved in the decryption holds all the necessary encryptions and secret-key material to be able to successfully conclude the algorithm with probability at least $1 - N \cdot 2^{-\text{sec}}$. The latter probability arising as we have N coefficients which need to be decrypted correctly, which happens with probability of $1 - 2^{-\text{sec}}$, via the method in Section 4.1, or Appendix A. Therefore, the overall probability that all the parties are correct is

$$1 - n \cdot N \cdot 2^{-\text{sec}}.$$

SECURITY. We consider the strongest case of $n-1$ corruptions, so the adversary is given all-but-one of the secret keys sk_i . We need to show that the resulting encryption scheme is IND-CPA secure. Let j be the index of the honest party. We consider an experiment where the adversary is given a public key pk together with the secret keys $\{\text{sk}_i\}_{i \neq j}$ and the encryption of a challenge message m . We show that this is indistinguishable from an experiment where the adversary receives an encryption of a random message. We proceed via a hybrid argument, starting with the bottom layer h of the tree.

Hybrid H_h . This is the experiment defined above, where the adversary is given pk , $\{\text{sk}_i\}_{i \neq j}$ and an encryption of m .

Hybrid H_ℓ , for $\ell \in \{h-1, \dots, 1\}$. Let v be the index of the node on layer ℓ of the tree that lies on the path to leaf j . This hybrid is defined as in $H_{\ell+1}$, except here we modify the ciphertext ct_v , which was previously an encryption of the secret key share sk_v under pk_v , and replace it with an encryption of zero under pk_v .

Hybrid H_0 . This hybrid is defined to be the same as H_1 , except we replace the ciphertext ct , which previously encrypted the message m under the root key pk_\perp , with an encryption of zero. This hybrid is now independent of m .

We now argue indistinguishability. For each $\ell \in [h]$, hybrids H_ℓ and $H_{\ell-1}$ are computationally indistinguishable due to the bounded IND-CPA security property of the LD-based FHE scheme (Definition 5.1). We can apply this property because in hybrid $\ell-1$ at node u , it always holds that one of the secret keys $\text{sk}_{u||0}$ or $\text{sk}_{u||1}$ is unknown to the adversary, as required. Furthermore, there is no circular security issue, because the encrypted secret key $\tilde{\text{sk}}_u$ is sampled independently of the keys $\text{sk}_{u||0}$ and $\text{sk}_{u||1}$. \square

Remark 5.1. Note that for correctness to hold it is not sufficient that for a single party the path from the root to the node is correctly split. We need this to happen for *all* parties simultaneously. This means that the obtained probability is in fact $1 - n \cdot N \cdot 2^{-\text{sec}}$ and not, as initially might be believed, $1 - \log(n) \cdot N \cdot 2^{-\text{sec}}$.

A simpler variant relying on circular security. The previous construction avoids relying on a circular security assumption by switching to a freshly sampled FHE key at each node of the tree. We could instead simplify this slightly, with a variant of the construction where only one set of FHE secret keys is used. Here, we would start by sampling an independent secret key $\tilde{\text{sk}}_i$ for each leaf i . The public key associated with node v is then defined as $\text{pk}_v = \text{PubKeyGen}(1^\lambda, \text{sk}_v, \text{params})$, where sk_v is the sum of all the leaf secret keys that are descendants of v . We additionally encrypt sk under pk_v and give this out to the relevant parties. This introduces a circular security assumption, however, it does not seem to offer any significant efficiency benefits except for a slightly simpler setup algorithm.

5.4 BGV Parameters Supporting Scooby

It would appear that at first sight the parameters needed for Scooby are larger than those needed for standard FHE bootstrapping, due to the increase in q required by Equation (1). However, this is not necessarily the case, as we now explain in the case of the BGV encryption scheme.

Standard BGV decryption simply requires the bound $q > 2 \cdot p \cdot (B + 1)$ for valid decryption, so we appear to have boosted the size of q by a factor of 2^{sec} . However, bootstrappable BGV as implemented in (say) HELib [HS21] utilizes an underlying levelled SHE scheme. At level zero, where no further homomorphic operations may take place without bootstrapping, we have a ciphertext modulus q_0 which satisfies $q_0 > 2 \cdot p \cdot (B + 1)$. At level L , i.e., the initial encryption level, we have a ciphertext modulus q_L which satisfies $q_L > 2 \cdot p \cdot (B + 1) \cdot 2^{b_p \cdot L}$, where b_p is the (average) bits-per-level of the chain of ciphertext moduli. On passing from each level from L down to zero, the size of the ciphertext modulus drops by (on average) 2^{b_p} . Note that, when bootstrapping a ciphertext from level zero, we do not end up with a ciphertext at level L , instead we obtain a ciphertext at level U (which denotes the so-called “usable” number of levels).

To see how this affects Scooby, we need to remember that at the end of the $\text{Eval}_{\text{pk}}^{\text{Scooby}}$ procedure we will have a ciphertext at level U . This will satisfy our bound in Equation (1) if $2^{b_p \cdot U} \geq 2^{\text{sec}}$. Then, in executing $\text{Dec}_{\text{sk}_i}^{\text{Scooby}}$, at each level of the tree we notice that we are actually executing an operation equivalent to bootstrapping. This is because at each node $v = u \| b$, where u is the parent node and $b \in \{0, 1\}$, we are essentially either performing a homomorphic decryption with the key $(1, -\tilde{\text{sk}}_{u \| 1}^{\text{FHE}})$, or a homomorphic decryption with the key $(0, -\tilde{\text{sk}}_{u \| 0}^{\text{FHE}})$. Thus, at each stage of the execution of $\text{Dec}_{\text{sk}_i}^{\text{Scooby}}$ we have a ciphertext ct which is at level U .

Examining the bootstrappable BGV parameters proposed in [HS21] we see that in all cases we have $2^{b_p \cdot U} \geq 2^{128}$. Thus the Equation (1) does not actually result in any increase in parameters, at least in the case of the BGV scheme.

6 Multi-Party HEDS from Weaker Assumptions

We now present alternative constructions to the previous section, without relying on FHE with linear decryption and a super-polynomial modulus. In the first construction, in Section 6.1, we use any generic FHE scheme and a 2-party HSS scheme that supports homomorphic evaluation of the FHE decryption circuit. This means we no longer need the local decryption trick from Section 4.1, so can use FHE based on LWE with a polynomial modulus [BV14]. All LWE-based FHE constructions have decryption in NC1, so the 2-party HSS can be instantiated based on the Paillier assumption [OSY21] or on class groups [ADOS22], which support HSS for all of NC1.

In Section 6.2, we also give a variant of the construction that *only* requires 2-party HSS, and not FHE. This gives a way to bootstrap two-party HSS constructions to the multi-party setting. We show how it can be used to transform two-party HSS for branching programs, based on Paillier encryption, into 4-party HSS for homomorphic evaluation of constant-degree polynomials.

6.1 Scrapy: HEDS from Standard FHE + HSS for NC1

This construction, shown in Fig. 5, follows the tree-based structure of Scooby from the previous section. Previously, though, at each node of the tree, an FHE ciphertext was split into two ciphertexts encrypting shares of the message, by doing a special homomorphic decryption procedure tailored to the linear decryption property of the FHE scheme. In Scrapy, we instead do the

homomorphic decryption procedure inside a 2-party HSS scheme. Since most FHE schemes have decryption in NC1, it suffices to rely on HSS for NC1, which can be built from non-LWE-based assumptions. Of course, if done naively, this means we no longer get encrypted shares of the previous message, but would actually obtain the shares directly due to use of HSS. To avoid leaking all intermediate shares, we use an additional FHE scheme on each level of the tree, and use this to homomorphically evaluate the HSS evaluation procedure. The HSS evaluation keys are then only given out at the leaves of the tree, while at higher levels they are encrypted under FHE. Note that we only need the weaker, private-key form HSS, from Remark 2.1, where the sharing algorithm can be seen as done by a trusted dealer.

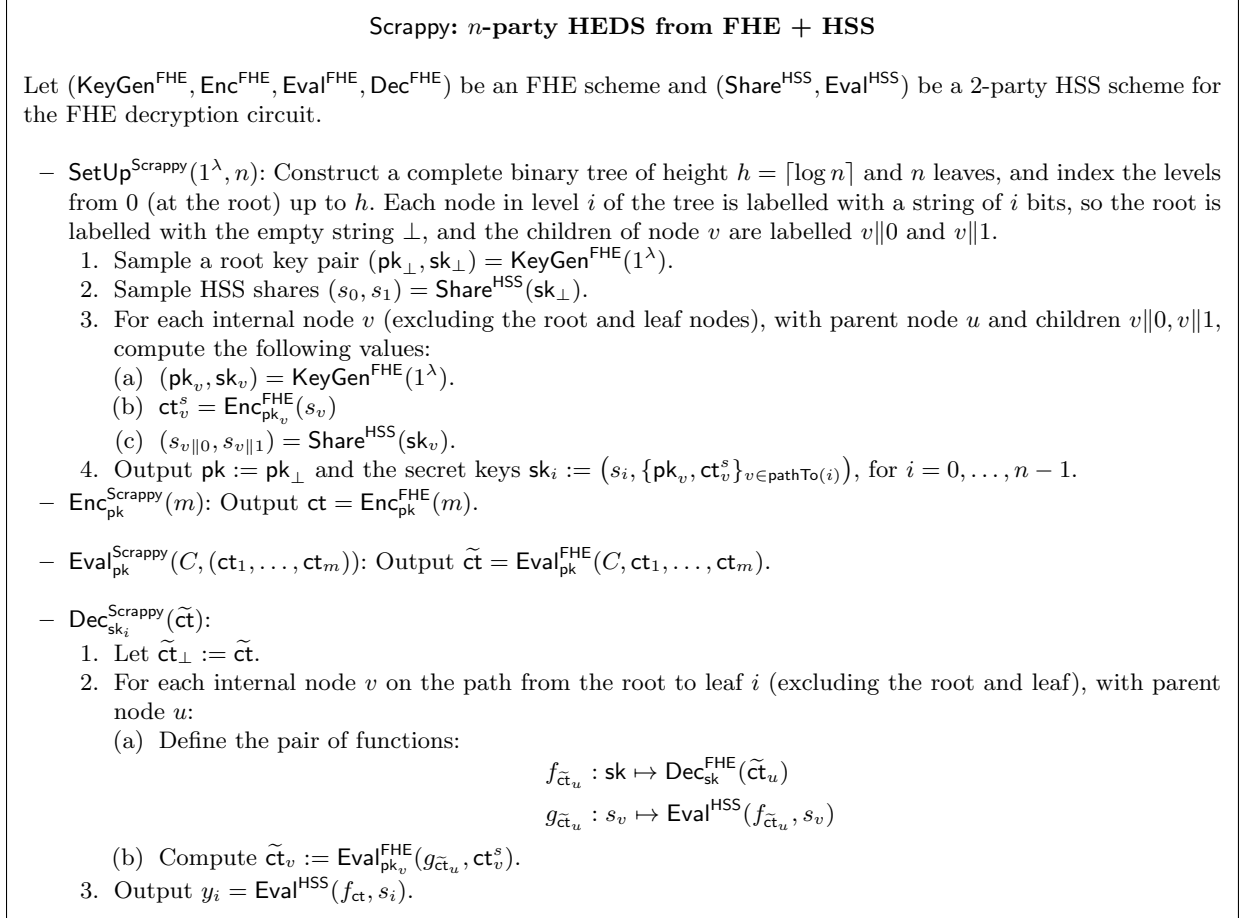


Fig. 5. Constructing n -party HEDS using standard FHE and 2-party HSS

Theorem 6.1. *Suppose there exists fully homomorphic encryption, and a 2-party HSS scheme that supports homomorphic evaluation of the FHE scheme’s decryption circuit. Then, there exists an n -party homomorphic encryption scheme with decryption to shares, for any $n = \text{poly}(\lambda)$.*

Proof. CORRECTNESS. We need to show that for any messages m_1, \dots, m_ρ and ciphertext $\tilde{\text{ct}}$ that is the output of Eval, on input a circuit C and encryptions of m_1, \dots, m_ρ , we have

$$y_1 + \dots + y_n = C(m_1, \dots, m_\rho) \pmod p$$

where $y_i = \text{Dec}_{\text{sk}_i}^{\text{Scrappy}}(\tilde{\text{ct}})$ and sk_i is the i -th secret key.

Consider the first step of the decryption algorithm, corresponding to the left and right children of the root node. For some $v \in \{0, 1\}$, this step homomorphically evaluates (in FHE) the function

$$g_{\tilde{\text{ct}}} : s_v \mapsto \text{Eval}^{\text{HSS}}(f_{\tilde{\text{ct}}}, s_v)$$

where here, $f_{\tilde{\text{ct}}}$ takes as input an FHE secret key, and uses it to decrypt $\tilde{\text{ct}}$. Note that this step is performed on the FHE ciphertext ct_v^s , where $v = 0$ for party i with $i < 2^h/2$, and $v = 1$ for the remaining parties. The parties obtain respective ciphertexts $\tilde{\text{ct}}_0, \tilde{\text{ct}}_1$, encrypted under pk_0 or pk_1 .

Let \tilde{m}_0, \tilde{m}_1 denote the messages in these two ciphertexts. Since s_0, s_1 are HSS shares of the FHE key sk_\perp , by the correctness of FHE and HSS, except with negligible probability it holds that

$$\begin{aligned} \tilde{m}_0 + \tilde{m}_1 &= g_{\tilde{\text{ct}}}(s_0) + g_{\tilde{\text{ct}}}(s_1) \\ &= \text{Eval}^{\text{HSS}}(f_{\tilde{\text{ct}}}, s_0) + \text{Eval}^{\text{HSS}}(f_{\tilde{\text{ct}}}, s_1) \\ &= f_{\tilde{\text{ct}}}(\text{sk}_\perp) \\ &= \text{Dec}_{\text{sk}_\perp}^{\text{FHE}}(\tilde{\text{ct}}) \\ &= \text{Dec}_{\text{sk}_\perp}^{\text{FHE}}(\text{Eval}_{\text{pk}_\perp}^{\text{FHE}}(\text{Enc}_{\text{pk}_\perp}^{\text{FHE}}(m_1), \dots, \text{Enc}_{\text{pk}_\perp}^{\text{FHE}}(m_\rho))) \\ &= \underbrace{C(m_1, \dots, m_\rho)}_{=\tilde{m}} \end{aligned}$$

This shows that, after the first layer of evaluation, the partition of parties ($P_0, \dots, P_{2^h/2-1}$) and ($P_{2^h/2}, \dots, P_{n-1}$) will each have obtained an encryption of \tilde{m}_0 or \tilde{m}_1 , under pk_0 or pk_1 , which are shares of the correct message \tilde{m} . By the same argument as above, at each subsequent node, the encryption of \tilde{m}_b will be split into two further encrypted shares, encrypted under the public keys for the next layer. This process continues until the penultimate layer of the tree, where for each non-leaf node v , any party $i \in \{v\|0, v\|1\}$ will compute a ciphertext $\tilde{\text{ct}}_v$, which encrypts under pk_v a share \tilde{m}_v .

At the final layer of the tree, if v has two children $v\|0, v\|1$, then each of the parties labelled with these leaves holds an HSS share s_i of sk_v , which can be used to homomorphically decrypt $\tilde{\text{ct}}_v$, where v is the parent of i , splitting \tilde{m}_v into two last shares. Otherwise, if v itself is a leaf node, then the party labelled with v uses its HSS share s_v to recover its share directly.

This shows that, with overwhelming probability, the parties obtain a correct sharing of the result.

SECURITY: We argue security for the strongest case of $n - 1$ corruptions, and assume for simplicity that n is a power of two. Let j be the index of the honest party. We consider an experiment where the adversary is given a public key pk together with the secret keys $\{\text{sk}_i\}_{i \neq j}$, and the encryption of a challenge message m . We gradually define a sequence of hybrid experiments, starting from the bottom layer h of the tree, where we eventually end up with one that is independent of the message m and indistinguishable from the first experiment. This implies the security property from Definition 3.3.

Hybrid $H_{h,1}$. This is the experiment defined above, where the adversary is given pk , $\{\text{sk}_i\}_{i \neq j}$ and an encryption of m .

Hybrids $H_{\ell,0}$ and $H_{\ell-1,1}$, for $\ell = h, \dots, 1$: Let v be the index of the node on layer ℓ of the tree that lies on $\text{pathTo}(j)$, let \bar{v} be its sibling and u their parent node. These hybrids are as follows:

- $H_{\ell,0}$: This hybrid is defined as in $H_{\ell,1}$, except we replace the HSS share $s_{\bar{v}}$, which was originally computed with $\text{Share}^{\text{HSS}}(\text{sk}_u)$, with a share from $\text{Share}^{\text{HSS}}(0)$. If v is not a leaf node, the ciphertext $\text{ct}_{\bar{v}}$ is now computed as an encryption of the new share $s_{\bar{v}}$.
- $H_{\ell-1,1}$: This hybrid is defined as in $H_{\ell,0}$, except here we modify the ciphertext ct_u , which was either an encryption of the HSS share s_u under pk_u , or (if $u = \perp$ is the root) an encryption of m under pk_{\perp} . In either case, we replace it with an encryption of zero under pk_u .

Hybrid $H_{\ell,0}$ is indistinguishable from $H_{\ell,1}$, because the second HSS share s_v is independent of the view of the adversary, so we can rely on the security property of HSS. Hybrid $H_{\ell-1,1}$ is indistinguishable from $H_{\ell,0}$, due to the IND-CPA security of the FHE scheme; this relies on the fact that in $H_{\ell,0}$, we have removed the FHE secret key sk_u from the view of the adversary.

In the final hybrid $H_{0,1}$, the root ciphertext ct_{\perp} has been replaced with an encryption of zero, so the view of the adversary is independent of the original message m . This concludes the proof. \square

The above theorem implies n -party HEDS assuming (1) LWE with a polynomial modulus [BV14], (2) circular security, and (3) HSS for NC1 circuits, which can be based on decisional composite residuosity [OSY21] or a DDH-like assumption in class groups [ADOS22]. If we only require n -party HEDS for bounded-depth circuits, we can remove the circular security assumption, since we only required levelled FHE.

6.2 Shaggy: Bootstrapping HEDS to More Parties

We now give a separate transformation that increases the number of parties in HEDS, *without* relying on fully homomorphic encryption. This construction, in Fig. 6, essentially applies one layer of the previous, tree-based construction, with a branching factor of n instead of 2. Additionally, instead of alternating between FHE and HSS evaluation, we always evaluate within an n -party HEDS scheme. This allows bootstrapping any sufficiently powerful n -party HEDS to support n^2 parties.

Theorem 6.2. *Let n -HEDS be an n -party HEDS for a class of circuits \mathcal{C} , whose decryption algorithm, when viewed as a function of sk_i , can be written as a circuit in \mathcal{C} . Then, n^2 -HEDS (in Figure 6) is an n^2 -party HEDS for \mathcal{C} . Its encryption and evaluation algorithms are the same as in n -Scooby, while the complexity of decryption increases by a polynomial factor.*

Proof. CORRECTNESS: Let $\text{ct}_1, \dots, \text{ct}_{\rho}$ be n^2 -HEDS encryptions of messages m_1, \dots, m_{ρ} . During decryption, party (i, j) first computes a ciphertext $\tilde{\text{ct}}_i$, a homomorphic decryption of ct , now encrypted under pk_i , and then uses $\text{sk}_{i,j}$ to recover $y_{i,j}$, a share of this.

By the correctness of n -HEDS, for each $i \in [n]$ we have

$$\begin{aligned} \sum_{j=1}^n y_{i,j} &= \sum_j \text{Dec}_{\text{sk}_{i,j}}^{\text{n-HEDS}}(\text{Eval}_{\text{pk}_i}^{\text{n-HEDS}}(f_{\text{ct}}, \text{ct}_i^s)) \\ &= f_{\text{ct}}(\text{sk}_i) \\ &= \text{Dec}_{\text{sk}_i}^{\text{n-HEDS}}(\text{ct}) \end{aligned}$$

It follows, again due to the correctness of n -HEDS, that $\sum_{i=1}^n \sum_{j=1}^n y_{i,j} = \sum_i \text{Dec}_{\text{sk}_i}^{n\text{-HEDS}}(\text{ct}) = C(m_1, \dots, m_\rho)$, as required.

SECURITY: The security can be argued in a similar way to the proof of Theorem 6.1, for just a single layer of the tree. Since the adversary corrupts at most $n^2 - 1$ parties, there is at least one pair (i, j) where the adversary does not know $\text{sk}_{i,j}$. This means that we can replace the ciphertext ct_i^s by an encryption of zero, which will be secure due to the security of n -HEDS. This in turn means that sk_i is unknown to the adversary, so we can consider another hybrid where the challenge ciphertext ct is replaced with an encryption of zero. This implies IND-CPA security to an adversary who sees up to $n^2 - 1$ secret keys. \square

Note that the decryption complexity of the bootstrapped construction n^2 -HEDS is increased by a polynomial factor. Depending on the original n -party scheme, then, it may not be possible to apply the transformation more than once, if the new decryption algorithm is no longer in the class \mathcal{C} .

Construction n^2 -HEDS

Let n -HEDS be an n -party HEDS. We build n^2 -party HEDS, and label the parties $P_{i,j}$, for $i, j \in [n]$

- **Setup** $^{n^2\text{-HEDS}}(1^\lambda, n^2)$:
 1. Let $(\text{pk}_\perp, \text{sk}_1, \dots, \text{sk}_n) = \text{Setup}^{n\text{-HEDS}}(1^\lambda, n)$.
 2. For $i \in [n]$:
 - (a) Sample $(\text{pk}_i, \text{sk}_{i,1}, \dots, \text{sk}_{i,n}) = \text{Setup}^{n\text{-HEDS}}(1^\lambda, n)$.
 - (b) Sample $\text{ct}_i^s = \text{Enc}_{\text{pk}_i}^{n\text{-HEDS}}(\text{sk}_i)$.
 3. Output $\text{pk} := (\text{pk}_\perp, \text{ct}_1^s, \dots, \text{ct}_n^s)$ and the n^2 secret keys $\text{sk}_{i,j} := (\text{ct}_i^s, \text{sk}_{i,j}, \text{pk}_i)$, for $i, j \in [n]$.
- **Enc** $_{\text{pk}}^{n^2\text{-HEDS}}(m)$: Output $\text{ct} = \text{Enc}_{\text{pk}_\perp}^{n\text{-HEDS}}(m)$.
- **Eval** $_{\text{pk}}^{n^2\text{-HEDS}}(C, (\text{ct}_1, \dots, \text{ct}_\rho))$:
 1. Compute $\text{ct} = \text{Eval}_{\text{pk}_\perp}^{n\text{-HEDS}}(C, \text{ct}_1, \dots, \text{ct}_\rho)$.
 2. Let f_{ct} be the function that takes as input sk_i and outputs $\text{Dec}_{\text{sk}_i}^{n\text{-HEDS}}(\text{ct})$.
 3. Compute $\tilde{\text{ct}}_i = \text{Eval}_{\text{pk}_i}^{n\text{-HEDS}}(f_{\text{ct}}, \text{ct}_i^s)$, for $i = 1, \dots, n$.
 4. Output $(\tilde{\text{ct}}_1, \dots, \tilde{\text{ct}}_n)$.
- **Dec** $_{\text{sk}_{i,j}}^{n^2\text{-HEDS}}(\tilde{\text{ct}}_i)$: Output $y_{i,j} = \text{Dec}_{\text{sk}_{i,j}}^{n\text{-HEDS}}(\tilde{\text{ct}}_i)$.

Fig. 6. Bootstrapping n -party HEDS to n^2 parties

Shaggy: Instantiation with HSS from Paillier. We now describe the Shaggy construction, which is obtained by applying the above transformation to two-party HSS based on the decisional composite residuosity assumption used in Paillier encryption. We can only apply the transformation once, so we obtain a 4-party HEDS/HSS scheme, which can support homomorphic evaluation of constant-degree polynomials. As the underlying two-party scheme n -HEDS, we use the HSS construction from [OSY21] or [RS21].

First, we need to frame the 2-party HSS constructions of [OSY21, RS21] in our HEDS framework. The constructions are given in a “public-key” flavour of HSS, with $\text{Setup}^{\text{HSS}}$ and Enc^{HSS} algorithms which are the same as in HEDS. The Eval^{HSS} algorithm, however, requires knowing a secret key,

unlike the syntax for $\text{Eval}^{\text{HEDS}}$. We can still make this fit in our HEDS framework, by adjusting the scheme so that homomorphic evaluation is performed in the $\text{Dec}_{\text{sk}_i}^{\text{HEDS}}$ step, which knows the secret key, instead of $\text{Eval}^{\text{HEDS}}$. Concretely, we define the $\text{Eval}^{\text{HEDS}}$ algorithm to simply output an “evaluated ciphertext” defined to be the set of input ciphertexts together with the circuit C . These are then passed to Dec^{HEDS} , together with the secret key, which then runs the HSS evaluation algorithm. This makes the resulting HEDS non-compact, since the complexity of decryption depends on the circuit, but it can still be used for the construction in Fig. 6.

Complexity of Evaluation in Paillier-based HEDS. We now analyze the circuit complexity of the resulting Dec^{HEDS} algorithm, which performs HSS evaluation of constant-degree polynomials. We can assume the polynomial is a simple monomial $f(x_1, \dots, x_c) = x_1 x_2 \cdots x_c$ for a constant number of inputs (since to handle sums of monomials, it’s enough to evaluate each monomial separately and add the shares).

With the methods of [OSY21, RS21], each input x_i is given as a Paillier encryption of x_i , together with encryptions of x_i multiplied with each bit of the secret key. In homomorphic evaluation, the parties perform $c - 1$ sequential multiplications, where in each of these, the core operation is a step that computes:

$$z = \text{DDLog}(C^d \pmod{N^2}) + F_k(\text{id}) \pmod{N}$$

Here, $N = pq$ is a public modulus, $C \in \mathbb{Z}_{N^2}^*$ is a ciphertext, d is a secret share that is known only to one party, and F is a pseudorandom function with key k known to both parties. The distributed discrete log function $\text{DDLog}(X)$ computes $\lfloor X/N \rfloor \cdot (X \pmod{N})^{-1} \pmod{N}$.

In general, modular exponentiation and inversion are not known to be in NC1. However, it turns out that $\text{DDLog}(C^d)$, when viewed as a function of d for fixed C , does lie in NC1. The idea is that since C is public, we can consider the powers $C^{2^j} \pmod{N^2}$ as hard-coded into the description of the function. Similarly, we hardcode $C^{-2^j} \pmod{N}$, for $j = 1, \dots, \ell$, where ℓ is the bit length of d . This allows computing

$$C^d = \prod_{j=1}^{\ell} (C^{2^j})^{d_j} \pmod{N^2}, \quad (C^d \pmod{N})^{-1} = \prod_{j=1}^{\ell} (C^{-2^j})^{d_j} \pmod{N}$$

Since iterated product, modular reduction, addition/subtraction and integer division are all in NC1 [BCH84], $\text{DDLog}(C^d)$ can be computed as an NC1 circuit. Furthermore, evaluation of a PRF based on factoring can be done in NC1 [NR04].

In the complete multiplication algorithm, the above step is repeated $O(\lambda)$ times in parallel, which does not affect the circuit depth. The multiplication algorithm is run c times sequentially, where the outputs of one multiplication are used as the private d shares input to the next. If c is a constant, it follows that the entire evaluation procedure is in NC1.

Plugging in two-party HSS for poly-sized branching programs (which includes NC1), we obtain the following.

Corollary 6.1. *Assume the decisional composite residuosity assumption holds. Then, there exists a 4-party (non-compact) homomorphic encryption scheme with decryption to shares for constant-degree polynomials.*

7 Casper: Friendly AFS-spooky Encryption from Multi-Key FHE

In this section we introduce *Casper*, a practical instantiation of an AFS-spooky encryption from a multi-key FHE (MK-FHE) scheme and a tree-based construction similar to the one described in Section 5. We recall that to obtain a spooky encryption, we need to avoid all interactions between the parties in the key generation phase, contrarily to *Scooby*, where the setup phase is either performed by a trusted third party or emulated via an MPC protocol. As mentioned before, our technique is essentially the same as for *Scooby*, except we now require an MK-FHE scheme to obtain a non-interactive setup. In FHE, such MK variants exist for many schemes, and we outline in this section two different spooky constructions, one for TFHE style MK-FHE schemes [CCS19], and one for B/FV-style MK-FHE schemes [CDKS19].

Motivation. Since these MK-FHE based constructions already exist for spooky encryption, we might wonder why we should try to build a HEDS like functionality in the first place. The reason is very simple: since our tree-based construction is more efficient compared to the spooky scheme of Dodis et al., we would like to apply our improvement to an adaptation of *Scooby* with MK-TFHE and obtain a more efficient spooky encryption. However, note that the resulting scheme *Casper* will be less efficient than *Scooby*, since MK-FHE is much less efficient than single key FHE.

To give a rough idea on the slow-down in MK-FHE operations, we can consider the TFHE versus MK-TFHE example: fixing the security level to 128 bits, the evaluation of an homomorphic NAND gate (which performs a bootstrapping) in single-key TFHE costs about 20ms, while in MK-TFHE the cost for the same operation may vary from 270 ms (for a 2-party scenario) to several seconds (for a 8-party scenario).

As already remarked, an AFS-spooky solution based on MK-FHE has already been proposed by Dodis et al. [DHRW16]. This solution is based on the LWE assumption with super-polynomial modulus and its complexity is $O(n^2 \cdot |F|)$, where n is the number of parties and $|F|$ is the size of the circuit evaluated. The main problem of this approach is that even simple operations in MK-FHE seem have complexity $O(n^2)$. The scheme proposed in [DHRW16] was essentially based on MK-FHE schemes proposed by Clear et al. [CM15] and Mukherjee et al. [MW16], which were just theoretical constructions at the time and far from being practical. The two solutions we propose in this section are based on two new MK-FHE constructions, which are the multi-key variants of the schemes TFHE and B/FV.

Multi-key FHE. The MK-TFHE scheme in [CCS19] has been proposed as an improvement of many of previous MK-FHE solutions, including [CM15] and [MW16] on which the first spooky scheme in [DHRW16] was based. The construction proposed in [CM15], and simplified in [MW16], was single-hop for keys (static), meaning that the number of parties has to be known before the computation begins, and after finishing a computation with a fixed set of keys, the output cannot be easily used in more computations involving additional keys. The line of work by [CM15] and [MW16] has been improved by [PS16] and [BP16], which proposed a multi-hop for keys (dynamic) constructions. The multi-key variant of TFHE proposed by Chen et al. [CCS19] simplifies the constructions proposed in previous works and is more efficient both in terms of memory and execution time. In terms of ciphertext size, the work by [PS16] has ciphertexts that have a square factor in the number of parties, while [BP16] and [CCS19] are linear. In terms of public and evaluation keys, [CCS19] proposes smaller public keys and does not require to store the bootstrapping key, since bootstrapping can be performed directly by using public keys. In terms of asymptotic complexity, the work

Construction	Leveled operations	Bootstrapping	(Expected) Complexity
[DHRW16]	yes	yes	$O(n^2 \cdot F)$
[CCS19]	no but possible	yes	$O(n^2 \cdot F + n^2 \cdot \log n)$
[KKL ⁺ 22]	yes	yes	$O(n \cdot F + n^2 \cdot \log n)$

Table 2. Comparing AFS-Spooky solutions based on MK-FHE constructions.

by [PS16] is $O(n^{2.37})$ and the one by [BP16] is polynomial in the number of parties, while [CCS19] the operations have a complexity of order $O(n^2)$. Unfortunately, no practical comparison can be done between [CCS19] and previous works because [CCS19] is the only work that comes with a proof-of-concept implementation. Previous constructions were too impractical to be implemented. This makes us believe that, even if the asymptotic costs are similar, concretely the scheme instantiated with the MK-TFHE construction by [CCS19] is more efficient.

By using the MK-TFHE construction by [CCS19], the execution requires to perform a bootstrapping for each homomorphic NAND gate in the circuit that is evaluated, followed by $\log n$ bootstrappings for the distributed decryption part. Furthermore, the solution presented in [CCS19] can be largely improved by evaluating levelled integer operations (not only binary gates and no bootstrapping after each operation). With this approach, the solution could have a huge improvement in terms of complexity.

The MK-B/FV solution [CDKS19], [KKL⁺22] is better than that used in [DHRW16] in terms of complexity. The levelled operations (in particular multiplication) can be in fact evaluated in $O(n)$ (from [KKL⁺22]), instead of $O(n^2)$, and the bootstrapping can be evaluated in $O(n^2)$ (from [CDKS19]). Again, this would lead to a better complexity for the resulting AFS-spooky construction than previous works. More concretely, if the circuit to be evaluated does not require bootstrapping (i.e., the parameters for the noise allow to perform all the circuit in a levelled mode), then bootstrapping is required only in the distributed decryption part; otherwise, the complexity should take into account the evaluation of periodic bootstrappings to reduce the noise.

We summarize the complexity of MK-FHE-based spooky encryption in Table 2 and give details on our new MK-FHE based AFS-spooky solutions in the following sections.

7.1 Spooky from MK-TFHE

The MK-TFHE scheme proposed in 2019 by Chen et al. [CCS19], is an extension to a multi-key setting of the gate bootstrapping technique proposed by TFHE. A gate bootstrapping allows to evaluate a binary gate and to perform a fast bootstrapping in order to bring the noise back to a fixed level.

The resulting MK-FHE scheme is non-interactive: after a setup phase where the secure parameters are fixed and a public (uniformly random) common reference string (CRS) is provided to parties, all the parties perform the key generation without interacting with each other, i.e., they generate both their secret and public keys locally. Public keys are then published and shared with all the parties. We describe these two phases, setup and key generation, in Figure 7.

Each party can now encrypt their inputs by using their own secret key, and a third party could do the same by using the public keys. Once the encrypted inputs are provided, every party involved in the process, or a third party, can independently perform homomorphic computations on these encrypted inputs, and all will produce the same final result. We describe the encryption

MK-TFHE construction (setup and key generation)

Setup(λ): Given in input the security parameter λ

1. Generate the LWE public parameters: let $n \in \mathbb{Z}$ be the LWE key size, q be the ciphertext modulus, ϕ_{LWE} be the LWE secret key distribution over \mathbb{Z} , χ_{LWE} be the error distribution over \mathbb{Z}_q , $B_{\text{LWE}} \geq 2$ and $d_{\text{LWE}} \in \mathbb{N}$ be the base and degree of decomposition, respectively.
2. Generate the RLWE public parameters: let N be the ring dimension, we assume N being a power of 2, ϕ_{RLWE} be the secret key distribution over R , χ_{RLWE} be the error distribution over R_q , $B_{\text{RLWE}} \geq 2$ and $d_{\text{RLWE}} \in \mathbb{N}$ be respectively the base and degree of decomposition for the RGSW gadget vector $\vec{g} = (q/B^1, \dots, q/B^d)$. Moreover, generate a uniformly random Common Reference String (CRS) $\vec{a} \leftarrow \mathcal{U}(R_q^d)$.
3. Set as $pp_{\text{LWE}} = (n, \phi_{\text{LWE}}, \chi_{\text{LWE}}, B_{\text{LWE}}, d_{\text{LWE}})$ as the LWE public parameters, and as $pp_{\text{RLWE}} = (N, \phi_{\text{RLWE}}, \chi_{\text{RLWE}}, B_{\text{RLWE}}, d_{\text{RLWE}})$, \vec{a}) as the RLWE public parameters.
4. Return the public parameters $pp = (pp_{\text{LWE}}, pp_{\text{RLWE}})$.

Key Generation(pp): Given in input the public parameters pp , each party P_i does the following:

1. Sample the LWE and RLWE secrets $\mathbf{s}_i \leftarrow \phi_{\text{LWE}}$ and the RLWE secret $z_i \leftarrow \phi_{\text{RLWE}}$;
2. Sample an error vector $\mathbf{e}^{(i)} \leftarrow \chi_{\text{RLWE}}^{d_{\text{RLWE}}}$;
3. Set the RLWE public key as $PK_i = \mathbf{b}^{(i)} = \mathbf{s}_i \cdot \mathbf{a} + \mathbf{e}^{(i)} \in R_q^{d_{\text{RLWE}}}$.
4. Given in input the public parameters, its own LWE and RLWE secret keys, each P_i generates the **bootstrapping key** as a uni encryption. For each of the n elements $s_{i,j}$ (for $j \in [n]$) of the LWE secret key:
 - (a) Sample $r_j \leftarrow \phi_{\text{RLWE}}$;
 - (b) Sample $\mathbf{e}_{j,1} \leftarrow \chi_{\text{RLWE}}^{d_{\text{RLWE}}}$ and $\mathbf{e}_{j,2} \leftarrow \chi_{\text{RLWE}}^{d_{\text{RLWE}}}$;
 - (c) Sample $\mathbf{f}_{j,1} \leftarrow \mathcal{U}(R_q^{d_{\text{RLWE}}})$;
 - (d) Compute $\mathbf{d}_j = r_j \cdot \mathbf{a} + s_{i,j} \cdot \mathbf{g} + \mathbf{e}_{j,1} \in R_q^{d_{\text{RLWE}}}$;
 - (e) Compute $\mathbf{f}_{j,0} = z_i \cdot \mathbf{f}_{j,1} + r_j \cdot \mathbf{g} + \mathbf{e}_{j,2} \in R_q^{d_{\text{RLWE}}}$.
5. Output $BK_{i,j} = (\mathbf{d}_j, F_j) \in R_q^{d_{\text{RLWE}}} \times R_q^{2d_{\text{RLWE}}}$, with $F_j = [\mathbf{f}_{j,0}, \mathbf{f}_{j,1}]$. Note the bootstrapping key of party i as BK_i .
6. Given in input the public parameters, its own LWE and RLWE secret keys, each party i generates the **key switching key** as a list of LWE encryptions under the secret key \vec{s}_i of all the coefficients of the RLWE secret key z_i . Note the key switching key of party i as KS_i .
7. Each P_i publishes the tuple (PK_i, BK_i, KS_i) .

Fig. 7. MK-TFHE Construction (setup and key generation) [CCS19]

MK-TFHE construction (encryption and evaluation)

Encrypt($m \in \{0, 1\}, \mathbf{s}_i$): Party i encrypts its input message $m \in \{0, 1\}$ under the secret key \mathbf{s}_i as an LWE ciphertext $\text{ct} = (\beta, \alpha) \in \mathbb{Z}_q^{n+1}$, such that $\beta - \alpha \cdot \mathbf{s}_i \approx q/4 \cdot m$.

A ciphertext encrypted under a single LWE secret key, can be extended to a ciphertext encrypted in a multi-key fashion (i.e., under all the k parties keys), by setting up to 0 all the other uniformly random parts (e.g., $\text{ct} = (\beta, \alpha) \in \mathbb{Z}_q^{n+1}$ encrypted under party P_1 secret key, becomes $\vec{\text{ct}} = (\beta, \alpha, \vec{0}, \dots, \vec{0}) \in \mathbb{Z}_q^{nk+1}$).

Homomorphic NAND gate($\text{ct}_1, \text{ct}_2, \{(PK_i, BK_i, KS_i)\}_{i=1, \dots, k}$): The evaluator (that can be a third party or each of the parties involved in the protocol) performs a multi-key NAND gate evaluation by:

1. Extending the input ciphertexts ct_1, ct_2 to multi-key fashion ciphertexts $\vec{\text{ct}}_1, \vec{\text{ct}}_2$;
2. Performing a linear combination $\vec{\text{ct}} = (5q/8, 0, \dots, 0) - \vec{\text{ct}}_1 - \vec{\text{ct}}_2$;
3. Performing a multi-key bootstrapping to $\vec{\text{ct}}$ followed by a multi-key key switching. For more details on these steps, please check [CCS19, Section 4].

Return an MK-LWE ciphertext $\text{ct} = (\beta, \alpha_1, \dots, \alpha_k) \in \mathbb{Z}_q^{nk+1}$ such that $\beta - \sum_{i=1}^k \alpha_i \cdot \vec{s}_i \approx q/4 \cdot (m_1 \wedge m_2)$ with overwhelming probability.

Fig. 8. MK-TFHE Construction (Encryption and Evaluation) [CCS19]

step and the evaluation of the homomorphic NAND gate – main homomorphic operation used in

MK-TFHE – in Figure 8. For more details on all these steps, we refer to [CCS19]. Observe that the scheme allows for more parties to join the computations afterwards, without the need for the old participants to regenerate keys. However, the parameter set chosen during setup phase will impose a maximal number of parties.

While all the previous algorithms are exactly the same proposed in [CCS19], the distributed decryption that we propose is different from the original construction. We show in Figure 10 how to perform such a decryption, without the parties interacting, in a similar way as we did for Scooby in Figure 11. The high level idea remains the same: we start from a ciphertext encrypting the final result under all the secret keys of the parties, and we construct a binary tree having this ciphertext as the root. Every node in the tree represents an encryption of a share of the final result s under a certain number of secret keys. When we split this node in two branches, the new nodes will contain a new encryption of one of the two shares s_1 and s_2 of s , each one encrypted under one of the two shares of the key used in the previous node. The computation of these new encryptions of these shares is performed by evaluating a multi-key bootstrapping, using only the public keys of the parties involved in that node. This means that any party, or any third party, can do the entire computation independently. When we reach the leaves of the tree, these leaves contain encryptions of the shares of each party, where each share is encrypted under a single secret key. At this point, each party can decrypts and obtain their own share.

Notice all the previous algorithms are exactly the same proposed in [CCS19], whereas the distributed decryption that we propose is different from the original construction. We show, in Figure 10, how to perform such a decryption, without the parties interacting, in a similar way as we did for Scooby in Figure 11. The high level idea remains the same: we start from a ciphertext encrypting the final result under all the secret keys of the parties, and we construct a binary tree having this ciphertext as the root. Every node in the tree represents an encryption of a share of the final result s under a certain number of secret keys. When we split this node in two branches, the new nodes will contain a new encryption of one of the two shares s_1 and s_2 of s , each one encrypted under one of the two shares of the key used in the previous node. The computation of these new encryptions of these shares is performed by evaluating a multi-key bootstrapping, using only the public keys of the parties involved in that node. This means that any party, or any third party, can do the entire computation independently. When we reach the leaves of the tree, these leaves contain encryptions of the shares of each party, where each share is encrypted under a single secret key. At this point, each party can decrypts and obtain their own share.

7.2 Spooky from MK-BFV

An example of MK variant for the schemes BFV and CKKS has been proposed in 2019 by Chen et al. [CDKS19], and improved in 2022 by Kim et al. [KKL⁺22]. In this section, we only focus on the BFV-style solution and do not consider CKKS, since in CKKS the noise is part of the message and cannot be reduced.

We are not aware of any practical implementation of a multi-key variant for BGV, however, since BGV and BFV are very similar constructions, we think it should be possible to have a MK solution for BGV too.

In [CDKS19] and [KKL⁺22], the authors extend addition, multiplication and bootstrapping operations of BFV to a multi-key setting. As for MK-TFHE, the schemes are non-interactive and the decryption is distributed.

Distributed decryption *à la Scooby* for MK B/FV

Setup(λ):

1. Given in input the security parameter λ , generate the RLWE public parameters: set N , a power of 2, to be the ring dimension, t to be the plaintext modulus, Q to be the ciphertext modulus, ϕ to be the secret key distribution over R , χ to be the error distribution over R_Q ;
2. Generate a uniformly random common reference string (CRS) $\mathbf{a} \leftarrow \mathcal{U}(R_Q^d)$;
3. Choose a gadget decomposition $h : R_Q \rightarrow R_d$ with a gadget vector $\mathbf{g} \in R_Q^d$;
4. Compute $\Delta = \lfloor Q/t \rfloor$. Return the public parameters $pp = (N, t, Q, \phi, \chi, \mathbf{a}, h, \mathbf{g})$.

Key Generation(pp): Given in input the public parameters pp , each party i generates its secret and public keys as follows:

1. Sample the RLWE secret $s_i \leftarrow \phi$;
2. Sample an error vector $\mathbf{e}_{0,i} \leftarrow \chi^d$ and compute $\mathbf{b}_i = -s_i \cdot \mathbf{a} + \mathbf{e}_{0,i} \in R_Q^d$;
3. Sample $r_i \leftarrow \phi$ and $\mathbf{e}_{1,i} \leftarrow \chi^d$ and compute $\mathbf{d}_i = -r_i \cdot \mathbf{a} + s_i \cdot \mathbf{g} + \mathbf{e}_{1,i} \in R_Q^d$;
4. Sample $\mathbf{u}_i \leftarrow \mathcal{U}(R_Q^d)$ and $\mathbf{e}_{2,i} \leftarrow \chi^d$ and compute $\mathbf{v}_i = -s_i \cdot \mathbf{u}_i - r_i \cdot \mathbf{g} + \mathbf{e}_{2,i} \in R_Q^d$;
5. Return the public key $PK_i = (\mathbf{b}_i, \mathbf{d}_i, \mathbf{u}_i, \mathbf{v}_i)$ and the evaluation key $EK_i = (\mathbf{b}_i[0], \mathbf{a}[0])$.

Encryption($m \in R_t, EK$):

1. Sample $w \leftarrow \phi$ and $e_0, e_1 \leftarrow \chi$.
2. Return $\text{ct} = w \cdot EK + (\Delta m + e_0, e_1) \in R_Q^2$. A ciphertext encrypted under a single secret key, can be extended to a ciphertext encrypted in a multi-key fashion (i.e., under all the k parties keys), by setting up to 0 all the other uniformly random parts (e.g., $\text{ct} = (\beta, \alpha) \in R_Q^2$ encrypted under party P_1 secret key, becomes $\overline{\text{ct}} = (\beta, \alpha, 0, \dots, 0) \in R_Q^{k+1}$).

Homomorphic ADD($\overline{\text{ct}}_1 \in R_Q^{k+1}, \overline{\text{ct}}_2 \in R_Q^{k+1}$): The evaluator (that can be a third party or each of the parties involved in the protocol) returns $\overline{\text{ct}}_+ = \overline{\text{ct}}_1 + \overline{\text{ct}}_2 \in R_Q$.

Homomorphic MULT($\overline{\text{ct}}_1 \in R_Q^{k+1}, \overline{\text{ct}}_2 \in R_Q^{k+1}, \{(PK_i)\}_{i=1, \dots, k}$): Let $\overline{\text{ct}}_i = (\alpha_{0,i}, \alpha_{1,i}, \dots, \alpha_{k,i}) \in R_Q^{k+1}$. The evaluator (that can be a third party or each of the parties involved in the protocol) performs:

1. Compute $c_{i,j} = \lfloor (t/Q) \cdot \alpha_{i,1} \cdot \alpha_{j,2} \rfloor \in R_Q$;
2. Set $\overline{\text{ct}} = (c_{i,j})_{0 \leq i, j \leq k}$;
3. Perform a re-linearization on $\overline{\text{ct}}$ by using the public keys $\{(PK_i)\}_{i=1, \dots, k}$. For more details on this operation, we refer to [KKL⁺22].

Distributed Decrypt($\overline{\text{ct}}, s_i$): This step is the same as the one described in Figure 11 and uses the bootstrapping technique for MK-BFV proposed in [CDKS19].

Fig. 9. MK variant of the B/FV scheme [CDKS19, KKL⁺22] with distributed decryption *à la Scooby*

For completeness, we describe the different algorithms of the MK version of B/FV in Figure 9 and refer to [CDKS19] and [KKL⁺22] for more details. As for MK-TFHE, the setup phase generates all the parameters and a uniformly random CRS. The key generation algorithm can be performed independently by each party, i.e., each party generates its own secret and public keys and then broadcasts the public keys. The levelled homomorphic operations that can be performed publicly are additions and multiplications. All these algorithms are exactly the same as the ones proposed in the literature. As for MK-TFHE, the distributed decryption algorithm that we use to be similar to the one we propose in Figure 11.

Acknowledgements

The work of authors from KU Leuven was supported by CyberSecurity Research Flanders with reference number VR20192203 and by the FWO under an Odysseus project GOH9718N. Peter Scholl was supported by the Independent Research Fund Denmark (DFF) under project number 0165-00107B (C3PO) and the Aarhus University Research Foundation (AUFF).

Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of Cyber Security Research Flanders, the FWO, DFF or AUFF.

Distributed decryption *à la Scooby* for MK-TFHE

Scooby Decrypt($\text{ct}, s_i, \{(PK_j, BK_j, KS_j)\}_{j=1, \dots, k}\)$): As in Figure 11, the parties perform a spooky decryption to retrieve the shares. Party i does:

1. Assign $\text{ct}^{\text{Root}} \leftarrow \text{ct}$.
2. Party i now iterates over the path from the root node to the i -th leaf, until it reaches the internal node just above the leaf.
 - (a) Let $v = (u||b)$ be the current internal node (initially $v = \text{Root}$), with left node $v_L = u||1$ and right node $v_R = u||0$. By assumption, except with negligible probability, ct^v is a valid encryption of a message m_v such that $\sum m_v = m$ for all internal nodes of the same height in the tree.
 - (b) If $s_i \in v_L$ then party P_i sets $\text{ct}_L = (\beta, \alpha_1, \dots, \alpha_{k/2})$. They can therefore homomorphically evaluate the MK bootstrapping as if the secret key encrypting it was $(1, -s_1, \dots, -s_{k/2})$, by using $\{(PK_j, BK_j, KS_j)\}_{j=1, \dots, k/2}$.
 - (c) If $s_i \in v_R$ then party i sets $v_R = (0, \alpha_{k/2+1}, \dots, \alpha_k)$. They can therefore homomorphically evaluate the MK bootstrapping as if the secret key encrypting it was $(0, -s_{k/2+1}, \dots, -s_k)$, by using $\{(PK_j, BK_j, KS_j)\}_{j=k/2+1, \dots, k}$.
 - (d) Set $k = k/2$ and iterate until $k = 2$.
3. Each party P_i now has a ciphertext ct^v for an internal node v for which there are only two secret keys (s_i and another s_h) which are children of v .
4. Party P_i can now locally decrypt a message m_i such that $m_i + m_h = m_v$ via the method from Appendix A.

Fig. 10. Distributed decryption *à la Scooby*, adapted to the AFS-Spooky from MK-TFHE [CCS19]

References

- ADOS22. Damiano Abram, Ivan Damgård, Claudio Orlandi, and Peter Scholl. An algebraic framework for silent preprocessing with trustless setup and active security. *Cryptology ePrint Archive*, Report 2022/363, 2022. <https://ia.cr/2022/363>.
- BCG⁺19. Elette Boyle, Geoffroy Couteau, Niv Gilboa, Yuval Ishai, Lisa Kohl, and Peter Scholl. Efficient pseudorandom correlation generators: Silent OT extension and more. In Alexandra Boldyreva and Daniele Micciancio, editors, *Advances in Cryptology – CRYPTO 2019, Part III*, volume 11694 of *Lecture Notes in Computer Science*, pages 489–518, Santa Barbara, CA, USA, August 18–22, 2019. Springer, Heidelberg, Germany.
- BCG⁺20. Elette Boyle, Geoffroy Couteau, Niv Gilboa, Yuval Ishai, Lisa Kohl, and Peter Scholl. Correlated pseudorandom functions from variable-density LPN. In *61st Annual Symposium on Foundations of Computer Science*, pages 1069–1080, Durham, NC, USA, November 16–19, 2020. IEEE Computer Society Press.
- BCH84. P.W. Beame, S.A. Cook, and H.J. Hoover. Log depth circuits for division and related problems. In *25th Annual Symposium on Foundations of Computer Science, 1984.*, pages 1–6, 1984.
- BDGM19. Zvika Brakerski, Nico Döttling, Sanjam Garg, and Giulio Malavolta. Leveraging linear decryption: Rate-1 fully-homomorphic encryption and time-lock puzzles. In Dennis Hofheinz and Alon Rosen, editors, *TCC 2019: 17th Theory of Cryptography Conference, Part II*, volume 11892 of *Lecture Notes in Computer Science*, pages 407–437, Nuremberg, Germany, December 1–5, 2019. Springer, Heidelberg, Germany.
- BGI15. Elette Boyle, Niv Gilboa, and Yuval Ishai. Function secret sharing. In Elisabeth Oswald and Marc Fischlin, editors, *Advances in Cryptology – EUROCRYPT 2015, Part II*, volume 9057 of *Lecture Notes in Computer Science*, pages 337–367, Sofia, Bulgaria, April 26–30, 2015. Springer, Heidelberg, Germany.
- BGI16a. Elette Boyle, Niv Gilboa, and Yuval Ishai. Breaking the circuit size barrier for secure computation under DDH. In Matthew Robshaw and Jonathan Katz, editors, *Advances in Cryptology – CRYPTO 2016, Part I*, volume 9814 of *Lecture Notes in Computer Science*, pages 509–539, Santa Barbara, CA, USA, August 14–18, 2016. Springer, Heidelberg, Germany.
- BGI16b. Elette Boyle, Niv Gilboa, and Yuval Ishai. Function secret sharing: Improvements and extensions. In Edgar R. Weippl, Stefan Katzenbeisser, Christopher Kruegel, Andrew C. Myers, and Shai Halevi, editors, *ACM CCS 2016: 23rd Conference on Computer and Communications Security*, pages 1292–1303, Vienna, Austria, October 24–28, 2016. ACM Press.

- BGI⁺18. Elette Boyle, Niv Gilboa, Yuval Ishai, Huijia Lin, and Stefano Tessaro. Foundations of homomorphic secret sharing. In Anna R. Karlin, editor, *ITCS 2018: 9th Innovations in Theoretical Computer Science Conference*, volume 94, pages 21:1–21:21, Cambridge, MA, USA, January 11–14, 2018. LIPIcs.
- BGV12. Zvika Brakerski, Craig Gentry, and Vinod Vaikuntanathan. (Leveled) fully homomorphic encryption without bootstrapping. In Shafi Goldwasser, editor, *ITCS 2012: 3rd Innovations in Theoretical Computer Science*, pages 309–325, Cambridge, MA, USA, January 8–10, 2012. Association for Computing Machinery.
- BKS19. Elette Boyle, Lisa Kohl, and Peter Scholl. Homomorphic secret sharing from lattices without FHE. In Yuval Ishai and Vincent Rijmen, editors, *Advances in Cryptology – EUROCRYPT 2019, Part II*, volume 11477 of *Lecture Notes in Computer Science*, pages 3–33, Darmstadt, Germany, May 19–23, 2019. Springer, Heidelberg, Germany.
- BP16. Zvika Brakerski and Renen Perlman. Lattice-based fully dynamic multi-key FHE with short ciphertexts. In Matthew Robshaw and Jonathan Katz, editors, *Advances in Cryptology – CRYPTO 2016, Part I*, volume 9814 of *Lecture Notes in Computer Science*, pages 190–213, Santa Barbara, CA, USA, August 14–18, 2016. Springer, Heidelberg, Germany.
- BV14. Zvika Brakerski and Vinod Vaikuntanathan. Lattice-based FHE as secure as PKE. In Moni Naor, editor, *ITCS 2014: 5th Conference on Innovations in Theoretical Computer Science*, pages 1–12, Princeton, NJ, USA, January 12–14, 2014. Association for Computing Machinery.
- CCS19. Hao Chen, Ilaria Chillotti, and Yongsoo Song. Multi-key homomorphic encryption from TFHE. In Steven D. Galbraith and Shihō Moriai, editors, *Advances in Cryptology – ASIACRYPT 2019, Part II*, volume 11922 of *Lecture Notes in Computer Science*, pages 446–472, Kobe, Japan, December 8–12, 2019. Springer, Heidelberg, Germany.
- CDKS19. Hao Chen, Wei Dai, Miran Kim, and Yongsoo Song. Efficient multi-key homomorphic encryption with packed ciphertexts with application to oblivious neural network inference. In Lorenzo Cavallaro, Johannes Kinder, XiaoFeng Wang, and Jonathan Katz, editors, *ACM CCS 2019: 26th Conference on Computer and Communications Security*, pages 395–412. ACM Press, November 11–15, 2019.
- CGGI16. Ilaria Chillotti, Nicolas Gama, Mariya Georgieva, and Malika Izabachène. Faster fully homomorphic encryption: Bootstrapping in less than 0.1 seconds. In Jung Hee Cheon and Tsuyoshi Takagi, editors, *Advances in Cryptology – ASIACRYPT 2016, Part I*, volume 10031 of *Lecture Notes in Computer Science*, pages 3–33, Hanoi, Vietnam, December 4–8, 2016. Springer, Heidelberg, Germany.
- CGGI20. Ilaria Chillotti, Nicolas Gama, Mariya Georgieva, and Malika Izabachène. TFHE: Fast fully homomorphic encryption over the torus. *Journal of Cryptology*, 33(1):34–91, January 2020.
- CM15. Michael Clear and Ciaran McGoldrick. Multi-identity and multi-key leveled FHE from learning with errors. In Rosario Gennaro and Matthew J. B. Robshaw, editors, *Advances in Cryptology – CRYPTO 2015, Part II*, volume 9216 of *Lecture Notes in Computer Science*, pages 630–656, Santa Barbara, CA, USA, August 16–20, 2015. Springer, Heidelberg, Germany.
- DHRW16. Yevgeniy Dodis, Shai Halevi, Ron D. Rothblum, and Daniel Wichs. Spooky encryption and its applications. In Matthew Robshaw and Jonathan Katz, editors, *Advances in Cryptology – CRYPTO 2016, Part III*, volume 9816 of *Lecture Notes in Computer Science*, pages 93–122, Santa Barbara, CA, USA, August 14–18, 2016. Springer, Heidelberg, Germany.
- Ds17. Jack Doerner and abhi shelat. Scaling ORAM for secure computation. In Bhavani M. Thuraisingham, David Evans, Tal Malkin, and Dongyan Xu, editors, *ACM CCS 2017: 24th Conference on Computer and Communications Security*, pages 523–535, Dallas, TX, USA, October 31 – November 2, 2017. ACM Press.
- FV12. Junfeng Fan and Frederik Vercauteren. Somewhat practical fully homomorphic encryption. Cryptology ePrint Archive, Report 2012/144, 2012. <https://eprint.iacr.org/2012/144>.
- GH19. Craig Gentry and Shai Halevi. Compressible FHE with applications to PIR. In Dennis Hofheinz and Alon Rosen, editors, *TCC 2019: 17th Theory of Cryptography Conference, Part II*, volume 11892 of *Lecture Notes in Computer Science*, pages 438–464, Nuremberg, Germany, December 1–5, 2019. Springer, Heidelberg, Germany.
- GI14. Niv Gilboa and Yuval Ishai. Distributed point functions and their applications. In Phong Q. Nguyen and Elisabeth Oswald, editors, *Advances in Cryptology – EUROCRYPT 2014*, volume 8441 of *Lecture Notes in Computer Science*, pages 640–658, Copenhagen, Denmark, May 11–15, 2014. Springer, Heidelberg, Germany.
- GSW13. Craig Gentry, Amit Sahai, and Brent Waters. Homomorphic encryption from learning with errors: Conceptually-simpler, asymptotically-faster, attribute-based. In Ran Canetti and Juan A. Garay, editors, *Advances in Cryptology – CRYPTO 2013, Part I*, volume 8042 of *Lecture Notes in Computer Science*, pages 75–92, Santa Barbara, CA, USA, August 18–22, 2013. Springer, Heidelberg, Germany.

- HS21. Shai Halevi and Victor Shoup. Bootstrapping for HELib. *Journal of Cryptology*, 34(1):7, January 2021.
- KKL⁺22. Taechan Kim, Hyesun Kwak, Dongwon Lee, Jinyeong Seo, and Yongsoo Song. Asymptotically faster multi-key homomorphic encryption from homomorphic gadget decomposition. Cryptology ePrint Archive, Report 2022/347, 2022. <https://eprint.iacr.org/2022/347>.
- MW16. Pratyay Mukherjee and Daniel Wichs. Two round multiparty computation via multi-key FHE. In Marc Fischlin and Jean-Sébastien Coron, editors, *Advances in Cryptology – EUROCRYPT 2016, Part II*, volume 9666 of *Lecture Notes in Computer Science*, pages 735–763, Vienna, Austria, May 8–12, 2016. Springer, Heidelberg, Germany.
- NR04. Moni Naor and Omer Reingold. Number-theoretic constructions of efficient pseudo-random functions. *Journal of the ACM*, 51(2):231–262, 2004.
- OSY21. Claudio Orlandi, Peter Scholl, and Sophia Yakoubov. The rise of paillier: Homomorphic secret sharing and public-key silent OT. In Anne Canteaut and François-Xavier Standaert, editors, *Advances in Cryptology – EUROCRYPT 2021, Part I*, volume 12696 of *Lecture Notes in Computer Science*, pages 678–708, Zagreb, Croatia, October 17–21, 2021. Springer, Heidelberg, Germany.
- PS16. Chris Peikert and Sina Shiehian. Multi-key FHE from LWE, revisited. In Martin Hirt and Adam D. Smith, editors, *TCC 2016-B: 14th Theory of Cryptography Conference, Part II*, volume 9986 of *Lecture Notes in Computer Science*, pages 217–238, Beijing, China, October 31 – November 3, 2016. Springer, Heidelberg, Germany.
- RS21. Lawrence Roy and Jaspal Singh. Large message homomorphic secret sharing from DCR and applications. In Tal Malkin and Chris Peikert, editors, *Advances in Cryptology – CRYPTO 2021, Part III*, volume 12827 of *Lecture Notes in Computer Science*, pages 687–717, Virtual Event, August 16–20, 2021. Springer, Heidelberg, Germany.

Auxiliary Supplementary Material

A Two Party Distributed Decryption: Type **msb**

For the **msb**-type encryption schemes we can effectively use the same method described for the **lsb** version in Section 4.1. To prove our result, we will use the following lemma.

Lemma A.1 ([BKS19], Lemma 1). *Let $p, q \in \mathbb{N}$, $R = \mathbb{Z}[X]/(X^N + 1)$ for N a power of 2. Let $m, e \in R_p$ with $\|e\|_\infty \leq B$, and $t_0, t_1 \in R_q$ be uniformly random subject to $t_0 + t_1 = (q/p) \cdot m + e \pmod q$. Then, it holds that*

$$\left\lfloor (p/q) \cdot t_0 \right\rfloor + \left\lfloor (p/q) \cdot t_1 \right\rfloor = m \pmod p,$$

with probability at least $1 - N \cdot (B + 1) \cdot p/q$ over the choice of the shares t_0, t_1 .

Proof. Let $I = [-q/(2p), q/(2p))$. Let $t_0 = (q/p) \cdot z_0 + r_0$, such that $z_0 \in R_p$ and $r_0 \in R|_I$. Let $l \in R$ such that $t_1 = (q/p) \cdot (m - z_0) + e - r_0 + q \cdot l \in R$. Then we have

$$\begin{aligned} \left\lfloor (p/q) \cdot t_0 \right\rfloor &= \left\lfloor (p/q) \cdot ((q/p) \cdot z_0 + r_0) \right\rfloor \pmod p \\ &= \left\lfloor z_0 + (p/q) \cdot r_0 \right\rfloor \pmod p \end{aligned}$$

By the definition of r_0 we have that $(p/q) \cdot r_0 \in R|_{[-1/2, 1/2)}$, which means that rounding merely reduces the above to z_0 modulo p . Moreover,

$$\begin{aligned} \left\lfloor (p/q) \cdot t_1 \right\rfloor &= \left\lfloor (p/q) \cdot ((q/p) \cdot (m - z_0) + e - r_0 + q \cdot l) \right\rfloor \pmod p \\ &= \left\lfloor m - z_0 + p \cdot l + (p/q) \cdot (e - r_0) \right\rfloor \pmod p. \end{aligned}$$

Now assume that $e - r_0 \in R|_I$. Then it holds that $(p/q) \cdot (e - r_0) \in R|_{[-1/2, 1/2)}$, hence

$$\left\lfloor (p/q) \cdot t_1 \right\rfloor = m - z_0 \pmod p$$

If instead $e - r_0 \notin R|_I$, then the coefficients of r_0 are too close to the boundaries of I . As t_0 is chosen uniformly at random, we have that the distribution of z_0 is the uniform distribution over R_p and the distribution of r_0 is the uniform distribution of $R|_I$. For every $j \in \{1, \dots, N\}$ the j -th component of r_0 the probability that it is outside the interval

$$I_j := (-q/(2p) + e_j, q/(2p) + e_j]$$

is at most

$$(\|e\|_\infty + 1) \cdot p/q.$$

As, by assumption, $\|e\|_\infty < B$, a union over all the components of r_0 yields the bound of correct rounding with a probability

$$1 - N \cdot (B + 1) \cdot p/q.$$

□

As done before, we suppose that \mathbf{sk} is split into two keys \mathbf{sk}_1 and \mathbf{sk}_2 such that P_1 holds \mathbf{sk}_1 and P_2 holds \mathbf{sk}_2 , and that $\mathbf{sk} = \mathbf{sk}_1 + \mathbf{sk}_2 \pmod{q}$. We can then, without interaction, compute an additive sharing of $m = m_1 + m_2$ given a valid ciphertext ct encrypting m . This can be achieved as follows.

2-party DistDec^{msb}: Let $R \leftarrow \mathbb{Z}_q$ be a public random nonce.

1. P_1 computes $d_1 \leftarrow \text{ct} \cdot (1, -\mathbf{sk}_1) \pmod{q}$ and then $m_1 \leftarrow \lfloor d_1 \cdot p/q \rfloor$.
2. P_2 computes $d_2 \leftarrow \text{ct} \cdot (0, -\mathbf{sk}_2) \pmod{q}$ and then $m_2 \leftarrow \lfloor d_2 \cdot p/q \rfloor$.

We claim that this correctly computes $m = m_1 + m_2$ with overwhelming probability.

Proposition A.1. *Given an LD-based FHE scheme of type **msb** (Definition 4.1), where $(q, N, B, d, \Delta, S, \mathbf{pk}, \mathbf{sk}) \leftarrow \text{KeyGen}^{\text{FHE}}(1^\lambda, p)$, with $q > 2 \cdot p \cdot (B + 1) \cdot 2^{\text{sec}}$ and $\mathbf{sk}_1 + \mathbf{sk}_2 = \mathbf{sk}$. Let (ct, m) be a pair of ciphertext/plaintext messages and m_1 and m_2 values obtained with the procedure described above. Then, it holds that*

$$m = m_1 + m_2 \pmod{p},$$

with probability at least $1 - N \cdot 2^{-\text{sec}}$.

Proof. By the assumption on B , namely Equation 1, and by Lemma A.1, we have

$$1 - \frac{N \cdot (B + 1) \cdot p}{q} \leq 1 - \frac{N \cdot (B + 1) \cdot p}{2 \cdot p \cdot (B + 1) \cdot 2^{\text{sec}}} \leq 1 - N \cdot 2^{-\text{sec}}.$$

This concludes the proof. □

A.1 n -party Scooby Decryption: the **msb** case

In Figure 11, we give the procedure for n -party decryption to shares in the **msb** case of the Scooby construction.

Algorithm $\text{Dec}_{\text{sk}_i}^{\text{Scooby}}(\text{ct})$ (for msb-based construction)

Let $F : \{0, 1\}^\lambda \times [n] \rightarrow R_q$ be a pseudorandom function.

$\text{Dec}_{\text{sk}_i}^{\text{Scooby}}(\text{ct})$:

1. Parse sk_i as $\tilde{\text{sk}}_i$, K^{prf} and $(\text{pk}_v^{\text{FHE}}, \text{ct}_v)$, for every node v from the root to leaf i .
2. Let $\tilde{\text{ct}}_\perp := \text{ct}$.
3. For each internal node v on the path from the root to leaf i (excluding the root and leaf):
 - (a) Write $v = u||b$, where u is the parent of v (so $b = 0$ if v is a left child and $b = 1$ otherwise).
 - (b) Define the function:

$$f_{\tilde{\text{ct}}_u}^b : \text{sk} \mapsto \left[\tilde{\text{ct}}_u \cdot (b, -\text{sk}) + (-1)^b \cdot F(K^{\text{prf}}, u) \pmod{q} \right]$$

(c) Compute $\tilde{\text{ct}}_v := \text{Eval}_{\text{pk}_v}^{\text{FHE}}(f_{\tilde{\text{ct}}_u}^b, \text{ct}_v)$.

4. Write $i = u||b$, then take the leaf ciphertext $\tilde{\text{ct}}_i$ and output the share

$$m_i = \left[\tilde{\text{ct}}_i \cdot (b, -\tilde{\text{sk}}_i) + (-1)^b \cdot F(K^{\text{prf}}, u) \pmod{q} \right]$$

Fig. 11. Decryption to shares for msb-based Scooby