Turning HATE Into LOVE: Homomorphic Ad Hoc Threshold Encryption for Scalable MPC

Leonid Reyzin*, Adam Smith**, and Sophia Yakoubov*

Boston University

Abstract. We explore large-scale fault-tolerant multiparty computation on a minimal communication graph. Our goal is to be able to privately aggregate data from thousands of users — for example, in order to obtain usage statistics from users' phones. To reflect typical phone deployments, we limit communication to the star graph (so that all users only talk to a single central server). To provide fault-tolerance, we require the computation to complete even if some users drop out mid-computation, which is inevitable if the computing devices are personally owned smartphones. Variants of this setting have been considered for the problem of secure aggregation by Chan *et al.* (Financial Cryptography 2012) and Bonawitz *et al.* (CCS 2017). We call this setting Large-scale One-server Vanishing-participants Efficient MPC (LOVE MPC).

We show that LOVE MPC requires at least three message flows, and that a three-message protocol requires some setup (such as a PKI). We then build LOVE MPC with optimal round- and communication- complexity (assuming semi-honest participants and a deployed PKI), using homomorphic ad hoc threshold encryption (HATE). We build the first HATE scheme with constant-size ciphertexts (although the public key length is linear in the number of users). Unfortunately, this construction is merely a feasibility result, because it relies on indistinguishability obfuscation.

We also construct more practical three- and five- message LOVE MPC in the PKI model for addition or multiplication. Unlike in the obfuscation-based construction, the per user message length in these protocols is linear in the number of users. However, the five-message protocol still has constant amortized message length, because only the first two messages are long, but they need to be exchanged only once (i.e., are input-independent and reusable) and thus can be viewed as setup.

* Leonid Reyzin and Sophia Yakoubov were supported in part by NSF grant 1422965.

** Adam Smith was supported in part by NSF awards IIS-1447700 and AF-1763786 and a Sloan Foundation Research Award.

Table of Contents

1 Introduction

Consider a service that has an app with a large smartphone user base. Suppose the service wants to collect aggregate usage statistics, but (for regulatory compliance, or for good publicity, or for fear of becoming a target for attackers and investigators) does not wish to learn the data of any individual user.

Let f be the function whose inputs are individual user data from up to n users and whose output is the aggregated information that the service wants to compute. Naturally, a secure multiparty computation protocol (MPC) for f can be used to provide the desired aggregate output to the service without revealing the inputs of any individual user.[1] However, in this setting, we cannot expect every phone to remain engaged for the duration of the protocol, as phones may go out of signal range or run out of charge. Thus, the protocol must be fault-tolerant: it must go on to completion even if some participants drop out. Moreover, given the large number of parties and the limitations on their computational power, the protocol needs to be efficient for every participating user. In particular, the users are assumed to be able to communicate directly only with the service provider.

On the other hand, this setting has its own advantages. The service collecting the data is already powerful enough to connect to and perform work for every user, and thus can be assumed to have a server (or server farm) that can perform a considerable amount of work in the protocol. Moreover, the users already trust the service to provide the code of the app and thus the implementation of the MPC code. Thus, an assumption that the server is semi-honest (aka honest-but-curious) is reasonable: the service itself does not want to have individual user data, for reasons outlined above, and the service itself is interested in arriving at the correct output. In other words, the service is honest, but does not want to know sensitive data, and thus we need to design protection against honest-but-curious servers. We will also assume that the users are honest-but-curious, as they run the app provided by the service. We call this setting Large-scale One-server Vanishing-participants Efficient MPC (LOVE MPC for short).

Of particular interest in this setting is the problem of computing the sum of the users' inputs for so-called *secure aggregation*. The problem of secure aggregation was first studied by Rastogi and Nath [RN10] and Shi *et al.* [SCR+11]. Chan *et al.* [CSS12] added fault-tolerance to the setting. Elahi *et al.* [EDG14] considered the problem of secure aggregation in the context of anonymous routing. Bonawitz *et al.* [BIK+17] considered the same model as we do here (without formalizing it) with the goal of achieving privacy-preserving federated learning. In this context, it is often the case that the users' inputs come from a constant-size space (e.g., binary), and thus the total sum is at most linear in the number of users.

[1] The question of what can be inferred about the individual users from the output of f is important, but orthogonal to our problem; this question is addressed at the point of choosing which f to compute—for example, by ensuring it is differentially private.

1.1 Our Contributions

In this paper, we formalize LOVE MPC and explore its limitations and possibilities. In Section 4.1, we show that three message flows are necessary, and that if LOVE MPC uses only three message flows, some setup (e.g. a PKI) is necessary.

We demonstrate two types of three-message (and therefore round-optimal) semi-honest LOVE MPC protocols for addition in the PKI model. The first type is simple and efficient, but requires messages whose size is linear in the number of users. The second type has constant-size messages but is not useable in practice because of heavy-weight tools (namely, indistinguishability obfuscation).

We also demonstrate a simple and efficient five-message semi-honest LOVE MPC protocol for addition over small message spaces (or multiplication) in the PKI model. This protocol consists of two phases: a two-message setup phase, and a three-message computation phase. The setup phase need only be performed once, after which the computation can be repeated many times. It requires linear-size messages only during the setup phase; after that, each computation uses only constant-size messages.

Our PKI model assumes each user has a public-private key pair, and users know the public keys of all the participants in the protocol. Note that we do not assume any correlated randomness: all the keys are generated separately and independently. How public keys are distributed is not important for our purposes; they can be assumed to be available from the semi-honest server, for example. This setup requires no additional trust assumptions and can be viewed simply as one initial communication round-trip that is reusable.

We now describe our technical approach in a bit more detail.

Three-Message LOVE MPC from HATE. To construct our three-message protocols for LOVE MPC, we rely on the following approach. Each user encrypts her input using the public keys of other users, in such a way that any subset of size $t+1$ users can decrypt it, but any smaller subset cannot. The users send their ciphertexts to the server, who homomorphically combines them in order to get a ciphertext corresponding to the output of f applied to the plaintexts. The server then sends the combined ciphertext to the users, who each decrypt to obtain shares of the output and send them back to the server; the server combines any $t+1$ of these shares to obtain the output.

Thus, the primitive we require is homomorphic ad hoc threshold encryption (HATE for short): homomorphic so the server can compute f without decrypting, ad hoc so the users can have uncorrelated keys, and threshold so $t+1$ users are necessary and sufficient to decrypt. In Section 3.2, we use indistinguishability obfuscation to construct the very first HATE with constant-size ciphertexts. However, since indistinguishability obfuscation is not useable in practice, in Section 3.1 we also build HATE schemes with ciphertexts linear in the number of users, from practical primitives like secret sharing and public key encryption.

Five-Message LOVE MPC from HTE. If we are willing to stray from round optimality, then we can use a homomorphic threshold encryption (HTE) scheme

that is not ad hoc by using an additional round-trip to set up correlated randomness. In particular, we use the ElGamal threshold encryption scheme described in Appendix A. Correlated randomness for this scheme can be set up using Shamir secret sharing in two message flows and $\Theta(n)$ communication per user. After these two rounds, the parties can use threshold ElGamal to compute as many multiplications (or additions over small message spaces) as they choose, at the cost of just three rounds and $\Theta(1)$ communication per user. So, the first computation requires five rounds and $\Theta(n)$ communication, but the amortized cost of a computation is just three rounds and $\Theta(1)$ communication per user.

1.2 Related Work

Work Related to LOVE MPC. Bonawitz *et al.* [BIK+17] present a LOVE MPC protocol for vector addition. Their honest-but-curious protocol can be viewed in the PKI model, similar to ours; it requires five messages and linear per-user communication complexity in the PKI model. In contrast, we present simpler protocols that require only three messages (Construction 3), and, at the cost of an additional two-message setup, achieve constant per-user communication (Construction 4). Our computational requirements on the users are also lighter in Construction 4, as we require amortized constant computation, while the protocol of Bonawitz *et al.* requires quadratic computation [BIK+17, Figure 3].

Many works have considered variants of our problem. For example, Shi *et al.* [SCR+11] present protocols in a similar client-server model that that are not fault-tolerant. Chan *et al.* [CSS12] present protocols in the same model that are fault-tolerant, but satisfy a different notion of privacy than MPC and require correlated setup. Tolerating vanishing participants in general MPC is considered by Badrinarayanan *et al.* [BJMS18] (who use the term "lazy" instead of "vanishing").

A number of papers propose the use of techniques similar to ours. In particular, the use of multi-key fully homomorphic encryption (FHE) for round-efficient multi-party computation has been explored by Mukherjee and Wichs [MW16]; the use of threshold FHE was considered by Boneh *et al.* [BGG+18]; and the combination of threshold and multi-key properties for FHE was considered by Badrinarayanan *et al.* [BJMS18]. None of these works consider the client-server communication model we consider. We use threshold homomorphic (but not fully homomorphic) encryption in all of our LOVE MPC constructions.

Work Related to HATE. Fully homomorphic ATE was demonstrated by Badrinarayanan *et al.* [BJMS18], but with polynomial-size ciphertexts. The share-and-encrypt approach that we use in Construction 1 has also appeared in the past (but without homomorphism)—e.g., in the work of Daza *et al.* [DHMR07].

A number of papers [BZ14,ABG+13,Zha14] use obfuscation to achieve constant-size ciphertexts in broadcast encryption; we use it in Constructions 2 and 5 to achieve constant-size ciphertexts in HATE.

2 Threshold Encryption (TE) Definitions

A threshold encryption scheme [DHMR07] is an encryption scheme where a message is encrypted to a group $\mathcal{R}$ of recipients, and decryption must be done collaboratively by at least $t+1$ members of that group. (This can be defined more broadly for general access structures, but we limit ourselves to the threshold access structure.) We give an example of a threshold encryption scheme (a threshold variant of ElGamal, due to Desmedt and Frankel [DF90]) in Appendix A.

2.1 Threshold Encryption Algorithms

A threshold encryption scheme consists of five algorithms, described below. This description is loosely based on the work of Daza *et al.* [DHMR07], but we modify the input and output parameters to focus on those we require in our primary constructions (Section 3.1), with some additional parameters discussed in the text.

$\mathsf{Setup}(1^\lambda, t) \to (\mathsf{params}, \mathsf{msk})$ is a randomized algorithm that takes in a security parameter λ as well as a threshold t and sets up the global public parameters params for the system, as well as the master secret key msk for key generation. If $\mathsf{msk} = \perp$, the scheme is *ad hoc*, meaning that there is no master secret key and that each party can set up their own public-private key pair.
For simplicity, we provide Setup with the threshold t, and assume that t is encoded in params from hereon out. However, in t-flexible schemes, t may be decided by each sender at encryption time, and should then be an input to Enc (and encoded in the resulting ciphertext). In keyed-sender schemes (where the sender must use their secret key to encrypt and recipients must use the sender's public key to decrypt), t may also be specified in the sender's public key.

$\mathsf{KeyGen}(\mathsf{params}, \mathsf{msk}) \to (pk, sk)$ is a randomized key generation algorithm that takes in the global public parameters params and the master secret key msk and returns a public-private key pair.
If the scheme is ad hoc, KeyGen does not require the master secret key msk. Omitting msk from ad hoc schemes enables individual parties to run KeyGen themselves.
Some schemes require the sender's public key for decryption; we call such schemes *keyed-sender*. If the scheme is keyed-sender, the public and secret keys may each have two parts. Informally, those are the parts of the public key necessary for encryption to that party (and the parts of the secret key necessary for decryption by that party), and the parts of the public key necessary for the decryption of a message from that party (and the parts of the secret key necessary for encryption by that party). We give a keyed-sender construction in Section 3.2.

$\mathsf{Enc}(\mathsf{params}, \{pk_i\}_{i \in \mathcal{R}, |\mathcal{R}| > t}, m) \to c$ is a randomized encryption algorithm that encrypts a message m to a set of public keys belonging to the parties in

the intended recipient set $\mathcal{R}$ in such a way that any size-$(t+1)$ subset of the recipient set should jointly be able to decrypt. We assume t is specified within params, but (if the scheme is keyed-sender) it may also be specified within the sender's public key, or (if the scheme is t-flexible) on the fly as an input to Enc itself.

In keyed-sender schemes, Enc may also require the sender's private key sk_{Sndr}.

$\mathsf{PartDec}(\mathsf{params}, \{pk_i\}_{i\in\mathcal{R}}, sk_j, c) \to d_j$ is an algorithm that uses a secret key sk_j belonging to one of the intended recipients to get a partial decryption d_j of the ciphertext c. This partial decryption can then be combined with t other partial decryptions to recover the message.

In keyed-sender schemes, $\mathsf{PartDec}$ may also require the sender's public key pk_{Sndr}.

$\mathsf{FinalDec}(\mathsf{params}, \{pk_i\}_{i\in\mathcal{R}}, c, \{d_i\}_{i\in\mathcal{R}'\subseteq\mathcal{R},|\mathcal{R}'|>t}) \to m$ is an algorithm that combines $t+1$ or more partial decryptions to recover the message m.

Not all threshold encryption schemes allow/require all of the algorithm inputs described above. Sometimes disallowing an input can make the scheme less flexible, but, on the other hand, sometimes schemes that do not rely on certain inputs have an advantage.

More Flexibility: Unneeded Inputs. Most threshold encryption schemes in the literature require a trusted central authority who holds the master secret key msk to be the one to run the key generation algorithm for every party. This is often not ideal; in many scenarios, such a trusted central authority does not exist. We call a threshold encryption scheme *ad hoc* if a public-private key pair can be generated without knowledge of a master secret key; that is, if each party is able to generate their keys independently. We use the acronym ATE to refer to an ad hoc threshold encryption scheme.

Secondly, requiring both decryption algorithms ($\mathsf{PartDec}$ and $\mathsf{FinalDec}$) to be aware of the set of public keys belonging to individuals in the set $\mathcal{R}$ of recipients can be limiting. We call a threshold encryption scheme *$\mathcal{R}$-oblivious* if neither partial decryption nor final decryption uses this information. On the surface, it looks like an $\mathcal{R}$-oblivious scheme should require less communication, since the sender would never need to communicate $\mathcal{R}$ to the recipients. However, in Appendix B we show a lower bound on the ciphertext size in an $\mathcal{R}$-oblivious scheme that is linear in the size of the recipient set.

More Flexibility: Additional Inputs. In describing the threshold encryption algorithms, for the most part we assumed that the threshold t was fixed within the global public parameters params (or, in a keyed-sender scheme, in the sender's public key). However, some schemes allow the sender to choose t at encryption time; we call such schemes t-flexible.

2.2 Homomorphic Threshold Encryption

In order to use ad hoc threshold encryption for multi-party computation, we need it to be *homomorphic*. There are three natural notions of homomorphism:

1. Homomorphism over ciphertexts, which is the notion typically considered;
2. Partial decryption homomorphism, which we introduce in this paper; and
3. Server-aided homomorphism, which we also introduce in this paper.

We use the acronym HATE to refer to an ad hoc threshold encryption scheme that has any of these notions of homomorphism. We informally describe all three notions below.

Definition 1 (Threshold Encryption: Homomorphism). *An* $\mathcal{F}$-homomorphic *threshold encryption scheme additionally has the following algorithm:*

$\mathsf{Eval}(\mathsf{params}, \{pk_i\}_{i\in\mathcal{R}}, [c_1, \ldots, c_l], f) \rightarrow c^*$ *is an algorithm that, given* l *ciphertexts and a function* $f \in \mathcal{F}$*, computes a new ciphertext* c^* *which decrypts to* $f(m_1, \ldots, m_l)$ *where each* c_q*,* $q \in [1, \ldots, l]$ *decrypts to* m_q*.*

Definition 2 (Threshold Encryption: Partial Decryption Homomorphism). *A* $\mathcal{F}$-partial decryption homomorphic *threshold encryption scheme additionally has the following algorithm:*

$\mathsf{PdecEval}(\mathsf{params}, \{pk_i\}_{i\in\mathcal{R}}, sk_i, [d_{i,1}, \ldots, d_{i,l}], f) \rightarrow d_i^*$ *is an algorithm that, given* l *partial decryptions and a function* $f \in \mathcal{F}$*, computes a new partial decryption* d_i^* *which can be used together with other partial decryptions to recover* $f(m_1, \ldots, m_l)$*, where each* $d_{i,q}$*,* $q \in [1, \ldots, l]$ *is party* P_i*'s partial decryption of an encryption of* m_q*.*

Both the ciphertext c^* produced by Eval and the partial decryption d_i^* produced by $\mathsf{PdecEval}$ should be small — that is, they should have size polynomial in $|\mathcal{R}|$ and λ but independent of f and l. Notice that this does not preclude ciphertext growth; for instance, in a homomorphic scheme, a fresh ciphertext might have size independent of $|\mathcal{R}|$, and the output of Eval might have size linear in $|\mathcal{R}|$. We draw a line by calling objects that have size polynomial in λ but independent of all other parameters *compact*, and objects that have size polynomial in both λ and $|\mathcal{R}|$ *semi-compact*. The outputs of Eval and $\mathsf{PdecEval}$ need only be semi-compact.

The third notion of homomorphism is *server-aided homomorphism*, which is homomorphism with an additional efficiency requirement. If a threshold encryption scheme is *$\mathcal{F}$-server-aided homomorphic*, then the output c^* of Eval (which itself may only be semi-compact) can be split into compact components $\{c_i^*\}_{i\in\mathcal{R}}$ such that every recipient P_i, $i \in \mathcal{R}$ should then be able to run $\mathsf{PartDec}$ given just one compact component c_i^*. Any homomorphic scheme that operates on compact ciphertexts and produces another compact ciphertext is also server-aided homomorphic; however, a homomorphic scheme that produces semi-compact ciphertexts may also be server-aided homomorphic, as long as the output ciphertexts can be split up into compact components.

The motivation for this notion of homomorphism is that typically, it is desirable for any ciphertexts that are sent between parties to be as short as possible (preferably compact) in order to save on communication complexity. Semi-compact ciphertexts that need to be sent to multiple recipients can be expensive;

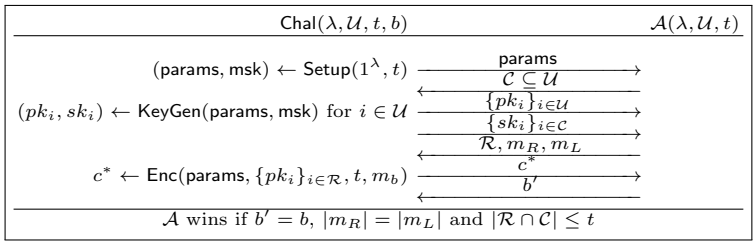

Fig. 1: Static Semantic Security Game for Threshold Encryption

however, even if the ciphertext is semi-compact, if each recipient only needs one compact component then in terms of communication complexity this can be as good as having compact ciphertexts. We describe a server-aided homomorphic ad hoc threshold encryption scheme in Section 3.2.2.

2.3 Threshold Encryption Security

We use the *semantic security* definition of Boneh *et al.* [BGG+18] for threshold encryption schemes.[2]

Definition 3 (Threshold Encryption: Static Semantic Security).
For $b \in \{R, L\}$, let $\mathsf{EXP}(\mathcal{A}, \lambda, n, t, b)$ denote the game described in Figure 1 played with the adversary $\mathcal{A}$, security parameter λ, number of existing parties $|\mathcal{U}| = n$, threshold t and fixed b. Let $\mathsf{WinProb}(\mathcal{A}, \lambda, n, t, b)$ denote the probability that the adversary $\mathcal{A}$ wins $\mathsf{EXP}(\mathcal{A}, \lambda, n, t, b)$.

A threshold encryption scheme $(\mathsf{Setup}, \mathsf{KeyGen}, \mathsf{Enc}, \mathsf{PartDec}, \mathsf{FinalDec})$ *is* (n, t)-statically semantically secure *if for all efficient adversaries $\mathcal{A}$, there exists a negligible function $negl$ such that*

$$|\mathsf{WinProb}(\mathcal{A}, \lambda, n, t, R) - \mathsf{WinProb}(\mathcal{A}, \lambda, n, t, L)| \leq \frac{1}{2} + negl(\lambda).$$

Note that this definition of security implies that even when the scheme is ad hoc (and therefore KeyGen can be run by participants independently instead of by a trusted central party), KeyGen is assumed to be run honestly; in particular, public keys cannot be generated based on the knowledge of other public keys. We leave the design of definitions and protocols such that public keys *can* be generated maliciously for future work.

In order to make Definition 3 more analogous to real world situations, it would make sense to additionally allow the adversary to query the challenger

[2] This is analogous to the static security definition of Gentry and Waters [GW09] for broadcast encryption. In fact, we borrow the adjective "static" from their definition. We can also define adaptive semantic security by allowing the adversary to provide the set of corrupt parties $\mathcal{C}$ after seeing the set of all public keys; however, in this paper we used static, not adaptive security.

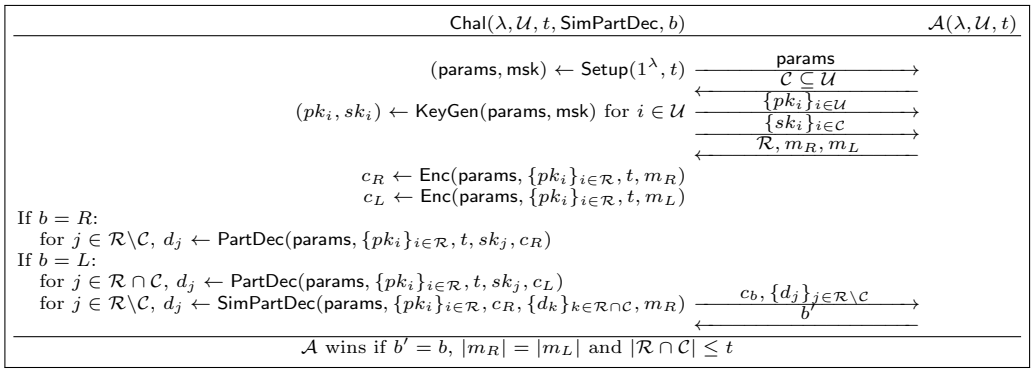

Fig. 2: Static Partial Decryption Simulatability Game for Threshold Encryption

on messages of its choice, and receive encryptions of those messages along with all corresponding partial decryptions. For the sake of simplicity, instead of modifying the static semantic security game in Figure 1, we add a second notion that we call *partial decryption simulatability* which implies that having the ability to make such queries will give the adversary no additional information. If a threshold encryption scheme is partial decryption simulatable, then it is possible to simulate remaining partial decryptions given t or fewer partial decryptions, a ciphertext, and a desired plaintext output. Our partial decryption simulatability is similar to, but stronger than, simulatability of partial decryption defined in [MW16], where only a single partial decryption can be simulated.

Definition 4 (Threshold Encryption: Static Simulatability).

For $b \in \{R, L\}$, let $\mathsf{EXP}(\mathcal{A}, \lambda, n, t, \mathsf{SimPartDec}, b)$ *denote the game described in Figure 2 played with the adversary $\mathcal{A}$, security parameter λ, number of existing parties $|\mathcal{U}| = n$, threshold t, simulation algorithm* SimPartDec *and fixed b. Let* $\mathsf{WinProb}(\mathcal{A}, \lambda, n, t, \mathsf{SimPartDec}, b)$ *denote the probability that the adversary $\mathcal{A}$ wins* $\mathsf{EXP}(\mathcal{A}, \lambda, n, t, \mathsf{SimPartDec}, b)$.

A threshold encryption scheme (Setup, KeyGen, Enc, PartDec, FinalDec) *is* (n, t)-statically partial decryption simulatable *if there exists an efficient algorithm* SimPartDec *such that for all efficient adversaries $\mathcal{A}$, there exists a negligible function negl such that*

$$|\mathsf{WinProb}(\mathcal{A}, \lambda, n, t, \mathsf{SimPartDec}, R) - \mathsf{WinProb}(\mathcal{A}, \lambda, n, t, \mathsf{SimPartDec}, L)| \leq \frac{1}{2} + negl(\lambda).$$

Putting it all together, we say that a threshold encryption scheme is has static security if it meets both of the above definitions.

Definition 5 (Threshold Encryption: Static Security). *A threshold encryption scheme* (Setup, KeyGen, Enc, PartDec, FinalDec) *is* (n, t)-statically secure *if it is both* (n, t)-statically semantically secure *(Definition 3) and* (n, t)-partial decryption simulatable *(Definition 4).*

3 Homomorphic Ad Hoc Threshold Encryption (HATE) Constructions

In this section, we describe some homomorphic ad hoc threshold encryption (HATE) constructions. The table in Figure 3 summarizes their properties. The first row of the table describes prior work, which focuses on *fully* homomorphic ad hoc threshold encryption (FHATE).[3]

In this paper, we consider two categories of additively-homomorphic ATE schemes: those with low *concrete* communication cost, and those with optimal *asymptotic* communication cost. Rows two and three of the table in Figure 3 describe two HATE instantiations — both based on share-and-encrypt (Construction 1) — which, despite their $\Theta(n)$-size ciphertexts, are efficient enough to be used in some scenarios.

The last row of the table ("obfuscation-based HATE", Construction 5) describes the first HATE scheme which has constant-size ciphertexts and partial decryptions. (It is also the first ATE scheme, homorphism or no, with these properties.) Unfortunately, it is a feasibility result more than anything else. It is not useable in practice, since it leverages indistinguishability obfuscation (iO) which currently has no practical instantiations.

3.1 HATE from Homomorphic Encryption and Secret Sharing

In this section, we describe our share-and-encrypt homomorphic ad hoc threshold encryption scheme which, despite its $\Theta(n)$-size ciphertexts, is efficient enough to be used in practice in some scenarios.

3.1.1 Background We leverage homomorphic public key encryption schemes and secret sharing. We assume familiarity with public key encryption schemes. We briefly review secret sharing in Appendix C. In particular, we use one non-standard property of secret sharing, which is *share simulatability*. Informally, a share simulatable t-out-of-n threshold secret sharing scheme ($\mathsf{SS.Share}$, $\mathsf{SS.Reconstruct}$) has a third algorithm $\mathsf{SS.SimShares}$ which takes in a message m_R and t or fewer honestly generated shares for a different message m_L, and generates the remaining shares such that all of the shares together are indistinguishable from an honestly generated sharing of m_R. $\mathsf{SS.SimShares}$ is only used in the proofs, not in the construction.

3.1.2 Building ATE from Homomorphic Encryption and Secret Sharing One natural way to build ATE is to use a threshold secret sharing scheme SS together with a public-key encryption scheme PKE, as in the work of Daza *et*

[3] There is another paper, due to Boneh *et al.* [BGG+18], that discusses fully homomorphic ad hoc threshold encryption, but because the homomorphism can only be applied to one ciphertext at a time in their scheme, it is not useable to instantiate LOVE MPC and we thus omit it from the table.

Name	pk size	sk size	ctext size	pdec size	homomorphism	message space size	assumption family	t-Flexible ?	$\mathcal{R}$-Oblivious ?
FHATE [BJMS18]	$\Theta(1)$	$\Theta(1)$	$poly(n)$	$poly(n, l)$	any	any	lattices	yes	yes
Shamir-and-ElGamal (Const. 1, Appendix E.1)	$\Theta(1)$	$\Theta(1)$	$\Theta(n)$	$\Theta(1)$	additive	small	DDH	yes	yes
CRT-and-Paillier (Const. 1, Appendix E.2)	$\Theta(1)$	$\Theta(1)$	$\Theta(n)$	$\Theta(1)$	additive	any	factoring	yes	no
obfuscation-based HATE (Const. 5, Section 3.2 and Appendix G)	$poly(n)$	$\Theta(1)$	$\Theta(1)$	$\Theta(1)$	additive*	small	iO	no	no

Fig. 3: A Summary of Homomorphic Ad Hoc Threshold Encryption Constructions. n refers to the number of parties, and l refers to the number of ciphertexts. (*) For the obfuscation-based scheme, homomorphism can be applied to an expanded ($\Theta(n)$) form of the ciphertext (the scheme is server-aided homomorphic).

al. [DHMR07]. The idea is to secret share the message, and to encrypt each share to a different recipient using their public key; therefore, we call this the *share-and-encrypt* construction. If the secret sharing and encryption schemes are homomorphic in compatible ways, the share-and-encrypt construction is a Homomorphic ATE (HATE).

Let $(\mathsf{PKE.KeyGen}, \mathsf{PKE.Enc}, \mathsf{PKE.Dec})$ be our public-key encryption scheme, and let $(\mathsf{SS.Share}, \mathsf{SS.Reconstruct})$ be our share simulatable threshold secret sharing scheme. We also allow algorithms $\mathsf{PKE.Setup}$ and $\mathsf{SS.Setup}$, which handle global setup for the encryption and secret sharing scheme, respectively. The share-and-encrypt ad hoc threshold encryption scheme is formally defined in Construction 1.

Theorem 1. *Share-and-encrypt (Construction 1) is a (n, t)-statically secure (Definition 5) ATE, as long as SS is a secure share simulatable t-out-of-n secret sharing scheme, and PKE is a CPA-secure public key encryption scheme.*

We prove Theorem 1 in Appendix D. In Appendix E, we describe two homomorphic instantiations of the share-and-encrypt ATE:

1. *Shamir-and-ElGamal* uses exponential Shamir secret sharing and the ElGamal public key encryption scheme, and
2. *CRT-and-Paillier* uses Chinese Remainder Theorem secret sharing and a variant of Paillier encryption.

$\mathsf{Setup}(1^\lambda)$:
- $\mathsf{params}_{\mathsf{PKE}} \leftarrow \mathsf{PKE.Setup}(1^\lambda)$
- $\mathsf{params}_{\mathsf{SS}} \leftarrow \mathsf{SS.Setup}(1^\lambda)$
- Return $\mathsf{params} = (\mathsf{params}_{\mathsf{PKE}}, \mathsf{params}_{\mathsf{SS}})$

$\mathsf{KeyGen}(\mathsf{params})$:
- Return $(pk, sk) \leftarrow \mathsf{PKE.KeyGen}(\mathsf{params}_{\mathsf{PKE}})$

$\mathsf{Enc}(\mathsf{params}, \{pk_i\}_{i\in\mathcal{R}}, t, m)$:
- $\{[m]_i\}_{i\in\mathcal{R}} \leftarrow \mathsf{SS.Share}(|\mathcal{R}|, t, m)$
- Return $c \leftarrow \{\mathsf{PKE.Enc}(pk_i, [m]_i)\}_{i\in\mathcal{R}}$

$\mathsf{PartDec}(\mathsf{params}, sk_i, c_i)$:
- Return $d_i \leftarrow \mathsf{PKE.Dec}(sk_i, c_i)$

$\mathsf{FinalDec}(\mathsf{params} = 1^\lambda, \{d_i\}_{i\in\mathcal{R}'\subseteq\mathcal{R},|\mathcal{R}'|>t})$:
- Return $m \leftarrow \mathsf{SS.Reconstruct}(\{d_i\}_{i\in\mathcal{R}'\subseteq\mathcal{R},|\mathcal{R}'|>t})$

Construction 1: Share-and-Encrypt Ad Hoc Threshold Encryption

Theorem 2. *Shamir-and-ElGamal (Appendix E.1) is an additively homomorphic ad hoc threshold encryption scheme for a polynomial-size message space.*

In Shamir-and-ElGamal we are limited to polynomial-size message spaces since final decryption uses brute-force search to find a discrete log. Jumping ahead to LOVE MPC, polynomial-size message spaces are still useful in many applications, as explained in the introduction. Moreover, the server already does work that is polynomial in the number of users, so asking it to perform another polynomial computation is not unreasonable.

Theorem 3. *CRT-and-Paillier (Appendix E.2) is an additively homomorphic ad hoc threshold encryption scheme.*

Both Shamir-and-ElGamal and CRT-and-Paillier are ad hoc threshold encryption schemes by Theorem 1; the homomorphisms in Theorems 2 and 3 follow from the homomorphisms of the underlying encryption and secret sharing schemes.

Communication Complexity. The share-and-encrypt ad hoc threshold encryption scheme has ciphertext size $\Theta(n)$ (asymptotics ignore the security parameter). On the other hand, all private and public keys remain constant-size.

Flexibility. If the secret sharing scheme SS does not require any setup (or requires setup independent of t), the share-and-encrypt scheme is t-flexible, since the sender can decide which threshold t to use at encryption time. Shamir-and-ElGamal is also $\mathcal{R}$-oblivious, since we are able to omit $\{pk_i\}_{i\in\mathcal{R}}$ as an input both to $\mathsf{PartDec}$ and to $\mathsf{FinalDec}$. However, CRT-and-Paillier is not $\mathcal{R}$-oblivious, since parties' moduli, which are part of their public keys, are necessary for $\mathsf{FinalDec}$.

Homomorphism. Depending on the homomorphisms of the underlying secret sharing and encryption schemes, the share-and-encrypt construction can have various homomorphisms; as described above, both Shamir-and-ElGamal and CRT-and-Paillier are additively homomorphic (with Shamir-and-ElGamal limitted to a polynomial-size message space).

Notice that we are able to omit all but the relevant part of the ciphertext as input to $\mathsf{PartDec}$ for each party (where the relevant part is the one encrypted under their key), making the scheme server-aided homomorphic. This further saves on communication in some contexts (Section 4.3).

Finally, since each partial decryption is simply a secret share, the scheme is partial decryption homomorphic in any way that the secret sharing scheme is homomorphic.

3.2 HATE from Indistinguishability Obfuscation

In this section, we introduce the first ad hoc threshold encryption construction with ciphertext size that is independent of the number of parties (at the expense of linear-size public keys). Because our construction is based on indistinguishability obfuscation (iO), its main purpose is to demonstrate that linear-size ciphertexts are not inherent, and a general ciphertext size lower bound is unlikely.

3.2.1 Background We leverage indistinguishability obfuscation [BGI+01], puncturable pseudorandom functions (PPRFs) [KPTZ13,BW13,BGI14b,SW14], and constrained signatures [BZ14], all of which we describe below. We also use Shamir secret sharing [Sha79], which we describe in Appendix C. Finally, we use pseudorandom generators [BM82], for which we do not provide a formal description since it is such a standard primitive.

Indistinguishability Obfuscation. Informally, indistinguishability obfuscation is a way to obfuscate a program in such a way that no efficient adversary can distinguish between the obfuscations of two programs of the same size as long as their input-output behaviors are the same. Indistinguishability obfuscation was first defined by Barak *et al.* [BGI+01], and a candidate construction was proposed by Garg *et al.* [GGH+13].

We restate the indistinguishability obfuscation definition of Garg *et al.* below (previously restated and rephrased by Boneh and Zhandry [BZ14]) with the simplification that we fix the parameter l to be circuit size.

Definition 6 (Definition 1 from Garg *et al.* [GGH+13] / Definition 2.1 from Boneh and Zhandry [BZ14]). *An* indistinguiability obfuscator iO *for circuits is an efficient algorithm satisfying the following conditions:*

- $\mathsf{iO}(C)$ *preserves the functionality of the circuit C. That is, for all circuits C, for all inputs x, we have that*

$$\Pr[\mathsf{ObfC}(x) = C(x) : \mathsf{ObfC} \leftarrow \mathsf{iO}(C)] = 1.$$

– *For any two circuits C_0, C_1 of the same size l with the same functionality, the circuits $\mathsf{ObfC}_0 = \mathsf{iO}(C_0)$ and $\mathsf{ObfC}_1 = \mathsf{iO}(C_1)$ are indistinguishable. That is, for any efficient distinguisher D, there exists a negligible function negl such that the following holds: For all circuit sizes $l \in \mathbb{N}$, for all pairs of circuits C_0, C_1 of size l, we have that if $C_0(x) = C_1(x)$ for all inputs x, then*

$$|\Pr[\mathsf{D}(\mathsf{iO}(C_0)) = 1] - \Pr[\mathsf{D}(\mathsf{iO}(C_1)) = 1]| \leq \frac{1}{2} + negl(l).$$

Puncturable Pseudorandom Functions. A puncturable pseudorandom function (PPRF) [KPTZ13,BW13,BGI14b,SW14] is a pseudorandom function (PRF) whose keys can be punctured. Let $k\{x\}$ denote the PPRF key k punctured at point x; then $\mathsf{PPRF}_k(x') = \mathsf{PPRF}_{k\{x\}}(x')$ for all $x' \neq x$, but given $k\{x\}$, $\mathsf{PPRF}_k(x)$ is indistinguishable from random.

We give a more formal definition of puncturable pseudorandom functions below.

Definition 7. *A* puncturable pseudorandom function PPRF *consists of three algorithms* KeyGen, Eval *and* Puncture.

$\mathsf{KeyGen}(1^\lambda) \to k$ *sets up the* PPRF *secret key* k.
$\mathsf{Eval}(k, x) \to y$ *evaluates the* PPRF *at point* $x \in$ domain *to obtain an output* $y \in$ range *for polynomial-size sets* domain *and* range. *In the rest of this paper, we use the alternative notation* $\mathsf{PPRF}_k(x)$ *to denote* $\mathsf{Eval}(k, x)$.
$\mathsf{Puncture}(k, x) \to k\{x\}$ *outputs a punctured* PPRF *key* $k\{x\}$.

Furthermore, the algorithms must satisfy the following conditions.

– (KeyGen, Eval) *must be a secure PRF.*
– *Functionality should be preserved over all unpunctured inputs. That is, for all inputs $x^* \in$ domain and keys k, if $k\{x^*\} \leftarrow \mathsf{Puncture}(k, x^*)$, then for all inputs $x \neq x^*$, $\mathsf{PPRF}_k(x) = \mathsf{PPRF}_{k\{x^*\}}(x)$.*
– *The true value of the* PPRF *at the punctured point x^* is indistinguishable from random given just $k\{x^*\}$.*

Constrained Signatures. Constrained signatures, introduced by Boneh and Zhandry [BZ14], are a special type of signature that supports constrained verification keys. Constrained verification keys reject all signatures on messages not matching some constraint C. Informally, constrained verification keys should be indistinguishable from real verification keys as long as no messages not matching the constraint have been signed.

We give a more formal definition of constrained signatures below.

Definition 8 (Definition 4.5 of Boneh and Zhandry [BZ14]). *A* constrained signature scheme SIG *consists of four algorithms* KeyGen, Sign, Verify *and* ConstrainedKeyGen.

$\mathsf{KeyGen}(1^\lambda, l) \to (pk, sk)$ *takes in the security parameter λ and an upper bound l on the constraint circuit. It outputs a valid verification/signing key pair (pk, sk). (We will often omit the parameter l.)*

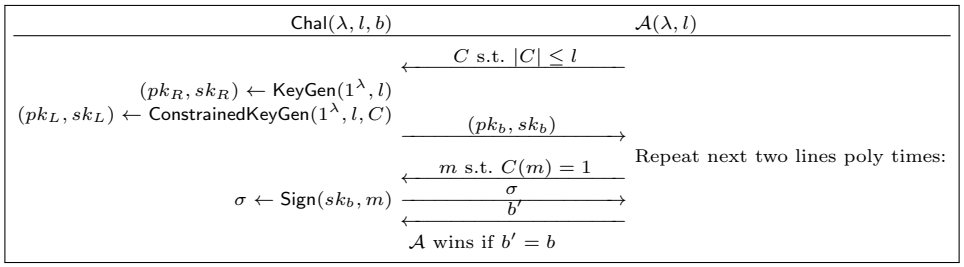

Fig. 4: Security Game for Constrained Signatures

$\mathsf{Sign}(sk, m) \to \sigma$ *signs the message* m *with the signing key* sk.

$\mathsf{Verify}(pk, m, \sigma) \to b$ *verifies the signature* σ *on the message* m *with the verification key* pk. *It returns* 1 *if the signature verifies, and* 0 *otherwise.*

$\mathsf{ConstrainedKeyGen}(1^\lambda, l, C) \to (pk, sk)$ *takes in the security parameter* λ, *an upper bound* l *on the constraint circuit, and a constraint circuit* C *such that* $|C| \leq l$. *It outputs a verification/signing key pair* (pk_C, sk_C) *such that* $\mathsf{Sign}(sk_C, m)$ *produces a valid signature relative to* pk_C *for all* m *such that* $C(m) = 1$, *but for any* m *where* $C(m) = 0$, *no valid signatures exist — that is,* $\mathsf{Verify}(pk_C, m, \sigma)$ *rejects all* σ.

Furthermore, the algorithms must satisfy the following two conditions.

- $(\mathsf{KeyGen}, \mathsf{Sign}, \mathsf{Verify})$ *must be a secure signature scheme (in the usual sense, e.g. existentially unforgeable [GMR88]).*
- *For* $b \in \{R, L\}$, *let* $\mathsf{EXP}(\mathcal{A}, \lambda, l, b)$ *denote the game described in Figure 4 played with the adversary* $\mathcal{A}$, *security parameter* λ, *constraint size upper bound* l *and fixed* b. *Let* $\mathsf{WinProb}(\mathcal{A}, \lambda, l, b)$ *denote the probability that the adversary* $\mathcal{A}$ *wins* $\mathsf{EXP}(\mathcal{A}, \lambda, l, b)$. *Then there exists a negligible function* $negl$ *such that for all* l *polynomial in* λ *and efficient adversaries* $\mathcal{A}$,

$$|\mathsf{WinProb}(\mathcal{A}, \lambda, l, R) - \mathsf{WinProb}(\mathcal{A}, \lambda, l, L)| \leq \frac{1}{2} + negl(\lambda).$$

3.2.2 Building ATE from Obfuscation The only asymptotic communication inefficiency in the share-and-encrypt HATE constructions of Section 3.1 comes from the $\Theta(n)$-size ciphertext. We can try to compress the ciphertext using obfuscation; instead of using the encrypted shares as the ciphertext, we can try to use an obfuscated program that *outputs* one encrypted share at a time given an appropriate input (such as receiver secret and public keys, and proof of the receiver's membership in the recipient set $\mathcal{R}$).

However, because it is difficult to obfuscate a secret sharing scheme without having the obfuscated program be of size linear in the threshold (because of the amount of randomness required for sharing), and we want our ciphertexts to be compact, we instead obfuscate a message- and recipient-set- agnostic program.

The program which each sender obfuscates, described in Algorithm 1, will be of size linear in the number of recipients[4], but it doesn't need to be part of the ciphertext. Instead, each sender can include such an obfuscated program just once in its public key. One can think of the obfuscated program in the sender's public key as a "horcrux".[5] The sender stores some of its secrets in this obfuscated program, and when it encrypts a message, the sender includes just enough information in the ciphertext that the obfuscated program can do the rest of the work.

Notice that having this obfuscated program as the sender's public key makes the obfuscation-based ATE scheme different from a typical public-key encryption scheme: the ATE scheme is *keyed-sender*, meaning that in order to encrypt a message the sender must use its secret key, and in order to decrypt a message, recipients need to use the sender's public key.

The obfuscation-based ATE is described in Construction 2. It uses an indistinguishability obfuscatior iO, puncturable pseudorandom function PPRF with range $\mathbb{Z}_p$, a constrained signature SIG, and a length-doubling pseudorandom generator PRG with domain $\{0,1\}^\lambda$ and range in $\{0,1\}^{2\lambda}$.

[4] At the cost of using differing-inputs obfuscation instead of indistinguishability obfuscation, the program can be made linear in the threshold t instead of in the number of recipients. This modification would additionally leverage cryptographic accumulators, e.g. a Merkle hash tree.

[5] A "horcrux" is a piece of one's soul stored in an external object, according to the fantasy series Harry Potter [Row05].

Let the public parameters $\mathsf{params} = (\lambda, p, t)$ consist of the security parameter λ, a large prime p such that the range of the puncturable pseudorandom function PPRF is in $\mathbb{Z}_p$, and the threshold t. For simplicity we omit params as input from the algorithms below.

$\mathsf{KeyGen}()$**:**

{The following generates the "receiver" portion of the keys. pv is used by others when sending messages to this party, and sv is used by this party to decrypt messages from others.}
$sv \leftarrow \{0,1\}^\lambda$
$pv \leftarrow \mathsf{PRG}(sv) \in \{0,1\}^{2\lambda}$
{The following generates the "sender" portion of the keys. $\mathsf{SIG}.sk$ and k_{Dec} are used by this party when sending messages to others, and $\mathsf{ObfFunc}$ is used by others to decrypt messages from this party.}
$(\mathsf{SIG}.pk, \mathsf{SIG}.sk) \leftarrow \mathsf{SIG.KeyGen}(1^\lambda)$
$k_{\mathsf{Dec}} \leftarrow \mathsf{PPRF.KeyGen}(1^\lambda)$
for $j \in [1, \ldots, t]$ **do**
  $k_{\mathsf{Share},j} \leftarrow \mathsf{PPRF.KeyGen}(1^\lambda)$
$k_{\mathsf{Share}} = (k_{\mathsf{Share},1}, \ldots, k_{\mathsf{Share},t})$
$\mathsf{ObfFunc} \leftarrow \mathsf{iO}(f_{k_{\mathsf{Dec}}, k_{\mathsf{Share}}, \mathsf{SIG}.pk})$
return $(pk = (pv, \mathsf{ObfFunc}), sk = (sv, \mathsf{SIG}.sk, k_{\mathsf{Dec}}))$

$\mathsf{Enc}(sk_{\mathsf{Sndr}} = (\mathsf{SIG}.sk, k_{\mathsf{Dec}}), \overrightarrow{pv} = \{pv_i\}_{i \in \mathcal{R}, |\mathcal{R}| \geq t}, m)$**:**

$\mathsf{nonce} \leftarrow \mathsf{PPRF.domain}$
$e = (\mathsf{PPRF}_{k_{\mathsf{Dec}}}(\mathsf{nonce}) + m) \bmod p$
$\sigma \leftarrow \mathsf{SIG.Sign}(\mathsf{SIG}.sk, (\overrightarrow{pv}, \mathsf{nonce}))$
return $c = (\mathsf{nonce}, e, \sigma)$

$\mathsf{PartDec}(pk_{\mathsf{Sndr}} = (pv, \mathsf{ObfFunc}), \overrightarrow{pv} = \{pv_i\}_{i \in \mathcal{R}}, sv_i, c)$**:**

Let idx be the index of the public value corresponding to the secret value sv_i in a lexicographic ordering of $\{pv_i\}_{i \in \mathcal{R}}$
$[m]_{\mathsf{idx}} \leftarrow \mathsf{ObfFunc}(\overrightarrow{pv}, \mathsf{idx}, sv_i, c)$
$d_i = (\mathsf{idx}, [m]_{\mathsf{idx}})$
return d_i

$\mathsf{FinalDec}(\{d_i\}_{i \in \mathcal{R}' \subset \mathcal{R}})$**:**

Perform Shamir reconstruction to recover the message m

Construction 2: Obfuscation-Based ATE

Algorithm 1 $f_{k_{\mathsf{Dec}},k_{\mathsf{Share}},\mathsf{SIG}.pk}(\overrightarrow{pv} = \{pv_i\}_{i\in\mathcal{R}}, \mathsf{idx}, sv, c)$

The following values are hardcoded in the program:

- $\mathsf{params} = (\lambda, p, t)$, where
  - λ is the security parameter,
  - p is large prime such that the range of the puncturable pseudorandom function PPRF is in $\mathbb{Z}_p$ and
  - t is the threshold
- A secret PPRF key k_{Dec} that is used to recover the message from the ciphertext c
- Secret PPRF keys $k_{\mathsf{Share}} = (k_{\mathsf{Share},1}, \ldots, k_{\mathsf{Share},t})$ that are used to produce randomness for sharing the message
- A signature verification key $\mathsf{SIG}.pk$

The following values are expected as input:

- public values $\overrightarrow{pv} = \{pv_i \in \{0,1\}^{2\lambda}\}_{i\in\mathcal{R}}$, ordered lexicographically
- index idx
- secret value $sv \in \{0,1\}^{\lambda}$
- ciphertext $c = (\mathsf{nonce}, e, \sigma)$

if $(\overrightarrow{pv}[\mathsf{idx}] = \mathsf{PRG}(sv))$ and $(\mathsf{SIG.Verify}(\mathsf{SIG}.pk, (\overrightarrow{pv}, \mathsf{nonce}), \sigma))$ **then**
  $w \leftarrow \mathsf{PPRF}_{k_{\mathsf{Dec}}}(\mathsf{nonce})$
  $m = (e - w) \bmod p$
  for $j \in [1, \ldots, t]$ **do**
    $\mathsf{coef}_j = \mathsf{PPRF}_{k_{\mathsf{Share},j}}(\mathsf{nonce})$
  $[m]_{\mathsf{idx}} = (m + \sum_{j\in[1,\ldots,t]} \mathsf{coef}_j \mathsf{idx}^j) \bmod p$ {This gives the idxth Shamir share of m}
  return $[m]_{\mathsf{idx}}$

It is important to check that the party invoking the program inputs a secret value that that is the preimage of one of the recipient public values, because otherwise any party could extract a secret share of the message. It is also important to check that the sender signed the set of recipient public values together with the nonce in order to bind the two objects together; if a nonce could be reused with multiple recipient sets, then an adversary could take a challenge ciphertext and submit it with a set of its own public values to decrypt it.

Informally, in order to prove security, we will have to show that given an obfuscation of this program, an adversary who only has t or fewer secret keys belonging to parties in the recipient set will not be able to tell the difference between an encryption of a message m_R and an encryption of a different message m_L. In order to do this, we will need to puncture k_{Dec} and k_{Share} on the challenge nonce to remove any critical information about the challenge plaintext from the program. It is crucial that, while we do this, the input-output behavior does not change at other possible nonces. In order for that to be possible, we must use a constrained signature scheme SIG. If we do, then in a hybrid, we can use a public

key such that no signatures exist on nonce together with any set of public keys other than the challenge recipient public key set, so that puncturing the keys at nonce can only affect the behavior when the challenge recipient public key set is used.

Security. We alter the static semantic security game (and partial decryption simulatability game) slightly to accommodate the obfuscation-based ATE construction. Informally, we require the adversary to commit to the recipient set $\mathcal{R}$ earlier (this is necessary for some of the hybrid games in our proof), and we provide the appropriate keys now that the scheme is keyed-sender. We describe the updated games in Appendix F. We call the new notion of security *super-static security.*

Theorem 4. *The obfuscation-based ATE (Construction 2) is* (n, t)*-super-statically secure (Definition 12) for any polynomial* n, t*, as long as* iO *is a secure indistinguishability obfuscator,* PPRF *is a secure puncturable* PRF *with range* $\mathbb{Z}_p$*,* SIG *is a constrained signature scheme, and* PRG *is a secure pseudorandom generator with domain* $\{0,1\}^\lambda$ *and range in* $\{0,1\}^{2\lambda}$*.*

We prove Theorem 4 in Appendix F.

Communication Complexity. The public keys in the obfuscation-based ATE (Construction 2) are large; because of the obfuscated program, which has input size $n = |\mathcal{R}|$, the public keys are of size polynomial in n. However, the ciphertexts are constant-size. This is the first ATE with constant-size ciphertexts.

Flexibility. The obfuscation-based ATE is not t-flexible, since the threshold t is fixed within the sender's public key. It is not $\mathcal{R}$-oblivious either, since each receiver has to input all receivers' public values to the obfuscated program.

Adding Server-Aided Homomorphism. Since partial decryptions are simply Shamir shares of the message, the obfuscation-based ATE is additively partial decryption homomorphic. However, in its current form, it is not homomorphic or server-aided homomorphic, since in order to extract anything homomorphic from the ciphertext, one must know a recipient secret value sv. Informally, in order to make the obfuscation-based ATE additively homomorphic, we can modify the obfuscated program to:

1. Use homomorphic public key encryption and decryption keys pk_i, sk_i instead of public and private values $pv_i = \mathsf{PRG}(sv_i), sv_i$,
2. Not require sk_i as input to the program, and
3. Return homomorphic encryptions of the secret shares instead of plaintext secret shares.

This modification would make the construction additively *server-aided homomorphic*; a server can save the recipients work by evaluating the obfuscated

program to extract encryptions of all recipients' partial decryptions, do homomorphic computation on those partial decryptions (since our PKE scheme is homomorphic, and we already have partial decryption homomorphism), and send all parties their final encrypted partial decryption.

More concretely, we can use ElGamal encryption [ElG84]. Since ElGamal is multiplicatively homomorphic (not additively homomorphic), we use exponential Shamir sharing to make the homomorphisms play nicely together. Once the obfuscated program is evaluated, we are essentially using the Shamir-and-ElGamal HATE (described in detail in Appendix E.1). In particular, this implies that we are limited to polynomial-size message spaces, since final decryption uses brute-force search to find a discrete log (see discussion after Theorem 2).

In Appendix G we give more details about this modification. Construction 5 describes the new additively server-aided homomorphic HATE; Algorithm 5 describes the new program that needs to be obfuscated and included in each sender's public key.

Theorem 5. *The modified obfuscation-based ATE (Construction 5) is* (n, t)*-super-statically secure (Definition 12) for any polynomial* n, t*, as long as* iO *is a secure indistinguishability obfuscator,* PPRF *is a secure puncturable* PRF *with range* $\mathbb{Z}_p$*,* SIG *is a constrained signature scheme, and* EG *is a secure public-key encryption scheme. Moreover, it is additively server-aided homomorphic for a polynomial-size message space.*

4 Large-scale One-server Vanishing-participants Efficient MPC (LOVE MPC)

Large-scale One-server Vanishing-participants Efficient MPC (LOVE MPC) is different from more traditional MPC in two ways: (1) in addition to tolerating corruptions, it tolerates some number of parties who vanish (i.e., drop out mid-computation), and (2) only the server learns the output. Our model is influenced by the work of Badrinarayanan *et al.* [BJMS18], which introduces the notion of "lazy parties" who may drop out during the protocol execution.

4.1 Lower Bounds

We show lower bounds both for the number of message flows in a LOVE MPC construction, and for the setup requirements.

Theorem 6. *For many functions (including addition), a LOVE MPC cannot be instantiated in fewer than three message flows, and if only three flows are used, then setup (e.g. correlated randomness or PKI) is unavoidable.*

Proof. We prove this theorem in two parts.

Lower Bounds on Number of Message Flows. A one-message protocol (where each user sends the server a single message, as in non-interactive MPC (NIMPC) [BGI+14a]) is impossible in our setting for many functions f, for the following reason. In a one-message protocol, the users would all send a single message to the server, who would compute the desired output. However, if the protocol is fault-tolerant, the set of participating users cannot be known in advance. Thus, an honest-but-curious server would be able to compute f on many different subsets of participating users, simply by ignoring some of the received messages. For example, if f is simply the sum of the users' individual values, the server could compute f both with and without a particular user present, thus learning every user's input.

A two-message protocol does not make sense, since a second message flow would involve the server sending the users messages. A server-to-user message before the user-to-server message does not solve the above problem, and a server-to-user message after the user-to-server message cannot affect the output, since the server should be the one to arrive at the output. We conclude that a LOVE MPC construction requires at least three message flows.

Lower Bounds on Setup Assumptions. A three-message protocol without any joint setup (e.g. correlated randomness or PKI) allows the server to perform what is essentially a Sybil attack. By fault-tolerance, the output should still be computable if a few participants drop out after sending the first message. Moreover, the output should not change depending on which participants drop in between the first and third flows; otherwise, the honest-but-curious server can pretend some users dropped out and see how the output changes (just like in the argument against one-message protocols). Therefore, the output should be fixed as soon as the second flow messages are sent by the server. (Generalizing to more than three message flows, the output should be fixed as soon as the server sends its last message.) This feature enables an honest-but-curious server to compute f on any single real user's input combined with inputs of the server's choice, as follows. After receiving the first message from a real user, the server will simulate the first message of $n-1$ users with inputs of the server's choice, and then simulate the rest of the messages of the protocol as if the real user dropped out before sending its third message. As long as the protocol can tolerate a single user dropping out, the server will be able to compute the desired output. We conclude that a three-message LOVE MPC construction requires some setup.

4.2 Definitions

Our ideal functionality, described in Figure 5, is a variant of the trusted party functionality of Badrinarayanan *et al.* [BJMS18], modified to support only a single output party (the server Srvr) and to allow functions with more than a single bit of output.

Let $\mathcal{U}$ be the set of all parties, $\mathcal{P}$ be the set of parties who do not drop out by the end of the protocol execution, and $\mathcal{C}$ be the set of corrupt parties. For correctness, we require that $|\mathcal{P}| > t$ for a dropout threshold t. For security,

we require that $|\mathcal{C}| \leq t_c$ for a corruption threshold t_c. Note that in all of our constructions, we have $t = t_c$.

A static semi-honest adversary $\mathcal{A}$ specifies the following sets: $\mathcal{C} \subseteq \mathcal{U}$ of corrupt parties such that $|\mathcal{C}| \leq t_c$, $\mathcal{D}_{\text{INPUT}} \subseteq \mathcal{U}$ of parties who drop out before the end of the input phase, and $\mathcal{D}_{\text{OUTPUT}} \subseteq \mathcal{U}$ of parties who drop out after the input phase (where $\mathcal{P} = \mathcal{U} \backslash (\mathcal{D}_{\text{INPUT}} \cup \mathcal{D}_{\text{OUTPUT}})$). Only parties who do not drop out before the end of the input phase (that is, $i \in \mathcal{U} \backslash \mathcal{D}_{\text{INPUT}}$) have their inputs included in the computation. The adversary receives the view of all the parties in $\mathcal{C}$.

Informally, a LOVE MPC protocol has static semi-honest security if for all input vectors $\{x_i\}_{i \in \mathcal{U}}$ and all efficient adversaries $\mathcal{A}$, there exists a simulator $\mathcal{S}$ who interacts with the ideal functionality in Figure 5 and can simulate the view of $\mathcal{A}$. Badrinarayanan *et al.* [BJMS18] also discuss security against malicious parties, which we do not address here.

We present our protocols in the PKI model. Because we consider only honest-but-curious attackers, the PKI model does not require any additional trust: the clients could simply exchange public keys via the server in two additional message flows before the start of the protocol. The importance of the PKI model for our protocols is that this exchange is independent of the inputs and needs to happen only once; after that, the protocols can be run repeatedly with the same public keys.

There are multiple ways to model PKI formally: "global" setup (e.g., Canetti and Rabin [CR03], Canetti *et al.* [CDPW07] and Dodis *et al.* [DKSW09]), which uses key registration that is shared by multiple, possibly different, protocols; or "local" setup (e.g., Barak *et al.* [BCNP04]), in which key registration is per protocol instance. In any of these, since our adversary is semi-honest, the simulator is allowed to know the secret keys of the corrupted parties; in addition, local setup means that the security definition is weaker and the simulator is more powerful, because the simulator can simulate the setup and thus is able to know (or even decide) secret keys for the honest parties. The LOVE MPC protocol that we describe in Section 4.3 can be proven secure with global setup, unless it is instantiated with the keyed-sender obfuscation-based HATE (Section 3.2), in which case it requires local setup (but still can be run multiple times with the same PKI).

4.3 Three-Message LOVE MPC from HATE

Let $(\mathsf{HATE.Setup}, \mathsf{HATE.KeyGen}, \mathsf{HATE.Enc}, \mathsf{HATE.PartDec}, \mathsf{HATE.FinalDec}, \mathsf{HATE.Eval})$ be a homomorphic ad hoc threshold encryption scheme. Assume the HATE has been set up (and so params is publicly available, and contains t), and that each party has already run KeyGen and that everyone's public keys have been distributed through a public key infrastructure. We describe a three-message HATE-based LOVE MPC in Construction 3. When instantiated with Shamir-and-ElGamal or CRT-and-Paillier, we call it Shamir-and-ElGamal LOVE MPC or CRT-and-Paillier LOVE MPC, respectively. As written, Construction 3 does not use keyed-sender HATE, and so cannot be instantiated with obfuscation-

Functionality $\mathcal{F}_f$, interacting with server Srvr and parties P_i, $i \in \mathcal{U}$.

INIT: On input (INIT) from the simulator $\mathcal{S}$:
1. Initialize an empty map INPUTS from parties to their inputs.

INPUTABORT: On input (INPUTABORT, $\mathcal{D}_{\text{INPUT}}$) from the simulator $\mathcal{S}$, store $\mathcal{D}_{\text{INPUT}}$.
INPUT: On input (INPUT, x_i) from party P_i: Store INPUTS$[i] = x_i$.
OUTPUTABORT: On input (OUTPUTABORT, $\mathcal{D}_{\text{OUTPUT}}$) from the simulator $\mathcal{S}$, store $\mathcal{D}_{\text{OUTPUT}}$.
OUTPUT: On input (OUTPUT) from the simulator $\mathcal{S}$:
1. Remove i from INPUTS for $i \in \mathcal{D}_{\text{INPUT}}$.
2. If $|\mathcal{U} \backslash (\mathcal{D}_{\text{INPUT}} \cup \mathcal{D}_{\text{OUTPUT}})| > t$: compute $y = f(\text{INPUTS})$.
3. Else: set $y = \perp$.
4. Output y to the server Srvr.

Fig. 5: Ideal Functionality $\mathcal{F}_f$ for LOVE MPC Secure Against Semi-Honest Adversaries.

based HATE. However, in Section 4.4 we alter Construction 3 to use obfuscation-based HATE and call the result obfuscation-based LOVE MPC.

Theorem 7. *HATE-based LOVE MPC (Construction 3) in the global-setup PKI model returns the correct output of f if fewer than t parties drop out. It is secure against t static semi-honest corruptions as long as* HATE *is a (n, t)-statically secure (Definition 5) $\mathcal{F}$-homomorphic ATE construction such that $f \in \mathcal{F}$.*

Proof. Correctness is true by the correctness of the underlying HATE.

To prove security, we describe a simulator $\mathcal{S}$ in Figure 6. Since we require global setup, $\mathcal{S}$ does not have access to honest parties' secret keys. However, because of the honest-but-curious assumption, $\mathcal{S}$ does see corrupt parties' secret keys and randomness. $\mathcal{S}$ can simulate by encrypting 0 for each honest party in Flow 1 (without knowing their secret keys), and simulating the partial decryptions for each honest party in Flow 3 (again without knowing their secret keys), by first performing partial decryption on behalf of the corrupt parties using their secret keys and randomness.

Notice that the only points in which the simulation differs from a real execution view is Flow 1, when the simulator encrypts 0s instead of the actual inputs, and Flow 3, when the simulator simulates partial decryptions instead of using genuine ones. The simulated corrupt parties' view is indistinguishable from a real view by CPA security and partial decryption simulatability, respectively.

Efficiency. Shamir-and-ElGamal LOVE MPC and CRT-and-Paillier LOVE MPC require $\Theta(n)$ communication per party, where $n = |\mathcal{U}|$. Since ciphertexts are $\Theta(n)$ in size, each party sends a $\Theta(n)$-size message in Flow 1, and receives a $\Theta(n)$-size message in Flow 2. However, we can leverage the server-aided homomorphism of share-and-encrypt and save some concrete cost by having the server only send each party the relevant part of the ciphertext in Flow 2; that is, the encryption of their secret share.

Flow 1: Each party P_i sends a message to the server Srvr

Each party P_i, $i \in \mathcal{U}$ does the following:

1. Computes
$$c_i \leftarrow \mathsf{HATE.Enc}(\mathsf{params}, \{pk_i\}_{i\in\mathcal{U}}, x_i).$$
2. Sends c_i to Srvr.

Flow 2: Server Srvr sends a message to each party P_i

Let $\mathcal{D}_{\text{INPUT}} \subseteq \mathcal{U}$ be the set of parties from whom the server Srvr did not receive a ciphertext. Srvr computes the sum ciphertext

$$c \leftarrow \mathsf{HATE.Eval}(\mathsf{params}, \{pk_i\}_{i\in\mathcal{U}}, \{c_i\}_{i\in\mathcal{U}\setminus\mathcal{D}_{\text{INPUT}}}, f)$$

and sends c to all parties $i \in \mathcal{U}\setminus\mathcal{D}_{\text{INPUT}}$.

Flow 3: Each party P_i sends a message to the server Srvr

Each party P_i, $i \in \mathcal{U}\setminus\mathcal{D}_{\text{INPUT}}$ does the following:

1. Computes
$$d_i \leftarrow \mathsf{HATE.PartDec}(\mathsf{params}, \{pk_j\}_{j\in\mathcal{U}}, sk_i, c).$$
2. Sends d_i to Srvr.

The server Srvr computes the output

Let $\mathcal{D}_{\text{OUTPUT}} \subseteq \mathcal{U}\setminus\mathcal{D}_{\text{INPUT}}$ be the set of parties from whom the server Srvr got a ciphertext c_i, but not a partial decryption d_i. As long as $|\mathcal{P} = \mathcal{U}\setminus(\mathcal{D}_{\text{INPUT}} \cup \mathcal{D}_{\text{OUTPUT}})| > t$, Srvr computes

$$y \leftarrow \mathsf{HATE.FinalDec}(\mathsf{params}, \{pk_i\}_{i\in\mathcal{U}}, c, \{d_i\}_{i\in\mathcal{P}}).$$

Construction 3: LOVE MPC for Function f From HATE in Three Rounds

1. The simulator $\mathcal{S}$ sends (INIT) to the ideal functionality.
2. $\mathcal{S}$ runs the adversary $\mathcal{A}$ to determine $\mathcal{C}$, $\mathcal{D}_{\text{INPUT}}$, $\mathcal{D}_{\text{OUTPUT}}$, and sends those to the ideal functionality.
3. [Flow 1] For each honest party P_i, $i \in \mathcal{U} \backslash \mathcal{C}$, $\mathcal{S}$ encrypts 0 in place of the actual input:

$$c_i \leftarrow \mathsf{HATE.Enc}(\mathsf{params}, \{pk_j\}_{j \in \mathcal{U}}, 0).$$

4. [Flow 2] Whether the server Srvr is honest or semi-honest (that is, whether or not its role is played by the simulator), it will correctly compute the sum ciphertext c and send it to all parties.
5. [Flow 3] For each corrupt party P_i, $i \in \mathcal{C}$, since the simulator $\mathcal{S}$ knows sk_i and the randomness used by the adversary, $\mathcal{S}$ can compute the partial decryption $d_i \leftarrow \mathsf{HATE.PartDec}(\mathsf{params}, \{pk_j\}_{j \in \mathcal{U}}, sk_i, c)$, which is guaranteed to equal the one the corrupt party will send the server Srvr in the final flow. ($\mathcal{S}$ can do this even if the adversary is rushing.) However, since $\mathcal{S}$ does not know the honest parties' keys, $\mathcal{S}$ must simulate their partial decryptions. $\mathcal{S}$ sends (OUTPUT) to the ideal functionality (note that the output is fixed at this point, since it is fixed as soon as the server sent the ciphertext), learns the actual output y, and computes the simulated partial decryptions for the honest parties by using the partial decryption simulatability (Definition 4) of the HATE. That is, the simulator runs

$$\{d_i\}_{i \in \mathcal{U} \backslash \mathcal{C}} \leftarrow \mathsf{HATE.SimPartDec}(\mathsf{params}, \{pk_i\}_{i \in \mathcal{U}}, c, \{d_i\}_{i \in \mathcal{C}}, y).$$

   By partial decryption simulatability, these will be indistinguishable from genuine partial decryptions.
6. [Output Computation] Whether the server Srvr is honest or semi-honest, it will correctly compute the output from the partial decryptions it receives.

Fig. 6: Simulator $\mathcal{S}$ for LOVE MPC from HATE

4.4 Three-Message LOVE MPC from Keyed-Sender Server-Aided Homomorphic ATE

The LOVE MPC protocol above (Construction 3) uses HATE that is *not keyed-sender* (that is, the sender does not need to use their own secret key to encrypt); so, it fits perfectly with our share-and-encrypt HATE (Construction 1), but not with our obfuscation-based HATE (Construction 5). We can modify it to use a HATE that is *keyed-sender* by adding sk_i as an input to $\mathsf{HATE.Enc}$, and the sender public keys pk_{Sndr} as inputs to $\mathsf{HATE.Eval}$ and $\mathsf{HATE.PartDec}$. When instantiated with obfuscation-based HATE, we call it obfuscation-based LOVE MPC. Obfuscation-based LOVE MPC will still be secure, with the caveat that now, the simulator will need access to all secret keys, because otherwise it will not be able to simulate honest parties' encryptions. This means that obfuscation-based LOVE MPC requires local setup.

Note that obfuscation-based HATE has super-static security (Definition 12) instead of static security; this means that the adversary must commit to the recipient set before seeing public keys. In particular, when we build LOVE MPC out of this HATE construction, a given setup instance can only be used for LOVE MPC among a fixed set of recipients.

Efficiency. Obfuscation-based LOVE MPC requires only constant communication per party once public keys have been distributed.

4.5 Five-Message LOVE MPC from Homomorphic Threshold Encryption

Given a threshold encryption scheme that is homomorphic but lacks the ad hoc property, we can achieve a LOVE MPC five-message protocol for multiplication in two phases. In the first phase, the parties establish some correlated randomness (which can be reused). In the second phase, the parties leverage the correlated randomness and use (non ad hoc) multiplicatively-homomorphic threshold encryption to compute on their inputs in three message flows. Note that the first phase is reusable; the second phase can be re-executed multiple times.

We can use the multiplicatively-homomorphic ElGamal-based threshold encryption construction TEG due to Desmedt and Frankel [DF90], described in Appendix A. The scheme operates over a group $\mathcal{G}$ of prime order p with generator g in which the decisional Diffie-Hellman problem is assumed to be hard. We assume that $\mathcal{G}$, p, and g are known to everyone; we also assume that each party P_i already has a key pair (pk_i, sk_i) for some semantically secure encryption scheme, and that pk_i is known to everyone.

We show this protocol, which we call Reusable Threshold ElGamal LOVE MPC, in Construction 4. In the first phase the parties run TEG.Setup as well as TEG.KeyGen in two message flows. In the second phase, all parties use the resulting instance of threshold ElGamal to encrypt their values to the joint public key pk, the server homomorphically multiplies them, broadcasts the resulting short ciphertext, and decrypts using the short partial decryptions it gets back.

In Construction 4 we show how the parties can compute multiplication. If the parties want to compute addition over small message spaces instead of multiplication, each party should encrypt g^{x_i} instead of x_i, and after TEG decryption the server can recover the output through brute-force search (same as in Theorem 2).

Theorem 8. *Reusable Threshold ElGamal LOVE MPC (Construction 4) in the global-setup PKI model returns the product of the parties' inputs in $\mathcal{G}$ if fewer than t parties drop out. It is secure against t static semi-honest corruptions under the decisional Diffie-Hellman assumption, as long as* PKE *is CPA-secure.*

Proof. Correctness is true by the correctness of the underlying primitives.

We informally prove security by describing a simulator $\mathcal{S}$. $\mathcal{S}$ has access to all parties' threshold ElGamal keys $[sk]_i$: because we have a semi-honest adversary, the simulator $\mathcal{S}$ can see all of the corrupt parties' randomness and computation, and because the simulator plays the honest parties' roles in the first phase, $\mathcal{S}$ learns the honest parties' keys as well. $\mathcal{S}$ behaves honestly in Phase 1. $\mathcal{S}$ can encrypt 0 for each honest party in Flow 2.1, and simulate the partial decryptions for each honest party in Flow 2.3 using the partial decryption simulatability of threshold ElGamal.

Efficiency. Notice that a single execution of phase 1 suffices for multiple executions of phase 2, and that while in phase 1 the communication is $\Theta(n)$ per party, in phase 2 it is constant. So, the amortized communication complexity per party over multiple executions of phase 2 is constant.

Flow 1.1 Each party P_i, $i \in \mathcal{U}$ does the following:
1. Picks a random $r_i \leftarrow^{\$} [p]$.
2. Shamir-shares r_i by picking a random polynomial f_i of degree t over the field $\mathbb{Z}_p$ with r_i as its y-intercept, and computing $[r_i]_j = f_i(j)$ for $j \in \mathcal{U}$.
3. Computes $c_{i,j} = \mathsf{PKE.Enc}(\mathsf{PKE}.pk_j, [r_i]_j)$ for $j \in \mathcal{U}$.
4. Sends $(g^{r_i}, \{c_{i,j}\}_{j \in \mathcal{U}})$ to the server Srvr.

Flow 1.2 The server Srvr does the following:
1. Computes the shared public key $pk = \prod_{i \in \mathcal{U}} g^{r_i}$
2. For each $i \in \mathcal{U}$, forwards the ciphertexts $\{c_{j,i}\}_{i \in \mathcal{U}}$ (as well as the shared public key pk) to P_i

Phase 1 Post-Processing Each party P_i, $i \in \mathcal{U}$ stores the shared public key pk, decrypts $c_{j,i}$ to obtain $[r_j]_i$ for all $j \in \mathcal{U}$, and computes its secret key share as $[sk]_i = \sum_{j \in \mathcal{U}} [r_j]_i$.

Flow 2.1 Each party P_i, $i \in \mathcal{U}$ sends $c_i = \mathsf{TEG.Enc}(pk, x_i)$ to Srvr.

Flow 2.2 The server Srvr computes the product ciphertext $c = \mathsf{TEG.Eval}(pk, \{c_i\}_{i \in \mathcal{U}}, \times)$, and sends c to all parties.

Flow 2.3 The parties compute their partial decryptions as $d_i = \mathsf{TEG.PartDec}([sk]_i, c)$ and send d_i to the server Srvr.

Phase 2 Post-Processing Srvr combines the partial decryptions as $y = \mathsf{TEG.FinalDec}(\{pk'_i\}_{i \in \mathcal{R}}, \{d_i\}_{i \in \mathcal{R}' \subseteq \mathcal{R}}, c)$ to obtain the output y.

Construction 4: Reusable Threshold ElGamal LOVE MPC

Acknowledgements

We would like to thank Ran Canetti for helpful discussions.

References

AB06. C. Asmuth and J. Bloom. A modular approach to key safeguarding. *IEEE Trans. Inf. Theor.*, 29(2):208–210, September 2006.

ABG+13. Prabhanjan Ananth, Dan Boneh, Sanjam Garg, Amit Sahai, and Mark Zhandry. Differing-inputs obfuscation and applications. Cryptology ePrint Archive, Report 2013/689, 2013. `http://eprint.iacr.org/2013/689`.

BCNP04. Boaz Barak, Ran Canetti, Jesper Buus Nielsen, and Rafael Pass. Universally composable protocols with relaxed set-up assumptions. In *45th FOCS*, pages 186–195. IEEE Computer Society Press, October 2004.

BGG+18. Dan Boneh, Rosario Gennaro, Steven Goldfeder, Aayush Jain, Sam Kim, Peter M. R. Rasmussen, and Amit Sahai. Threshold cryptosystems from threshold fully homomorphic encryption. In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO 2018, Part I*, volume 10991 of *LNCS*, pages 565–596. Springer, Heidelberg, August 2018.

BGI+01. Boaz Barak, Oded Goldreich, Russell Impagliazzo, Steven Rudich, Amit Sahai, Salil P. Vadhan, and Ke Yang. On the (im)possibility of obfuscating programs. In Joe Kilian, editor, *CRYPTO 2001*, volume 2139 of *LNCS*, pages 1–18. Springer, Heidelberg, August 2001.

BGI+14a. Amos Beimel, Ariel Gabizon, Yuval Ishai, Eyal Kushilevitz, Sigurd Meldgaard, and Anat Paskin-Cherniavsky. Non-interactive secure multiparty computation. In Juan A. Garay and Rosario Gennaro, editors, *CRYPTO 2014, Part II*, volume 8617 of *LNCS*, pages 387–404. Springer, Heidelberg, August 2014.

BGI14b. Elette Boyle, Shafi Goldwasser, and Ioana Ivan. Functional signatures and pseudorandom functions. In Hugo Krawczyk, editor, *PKC 2014*, volume 8383 of *LNCS*, pages 501–519. Springer, Heidelberg, March 2014.

BIK+17. Keith Bonawitz, Vladimir Ivanov, Ben Kreuter, Antonio Marcedone, H. Brendan McMahan, Sarvar Patel, Daniel Ramage, Aaron Segal, and Karn Seth. Practical secure aggregation for privacy-preserving machine learning. In Bhavani M. Thuraisingham, David Evans, Tal Malkin, and Dongyan Xu, editors, *ACM CCS 17*, pages 1175–1191. ACM Press, October / November 2017.

BJMS18. Saikrishna Badrinarayanan, Aayush Jain, Nathan Manohar, and Amit Sahai. Secure MPC: Laziness leads to GOD. Cryptology ePrint Archive, Report 2018/580, 2018. `https://eprint.iacr.org/2018/580`.

BM82. Manuel Blum and Silvio Micali. How to generate cryptographically strong sequences of pseudo random bits. In *23rd FOCS*, pages 112–117. IEEE Computer Society Press, November 1982.

BW13. Dan Boneh and Brent Waters. Constrained pseudorandom functions and their applications. In Kazue Sako and Palash Sarkar, editors, *ASIACRYPT 2013, Part II*, volume 8270 of *LNCS*, pages 280–300. Springer, Heidelberg, December 2013.

BZ14. Dan Boneh and Mark Zhandry. Multiparty key exchange, efficient traitor tracing, and more from indistinguishability obfuscation. In Juan A. Garay and Rosario Gennaro, editors, *CRYPTO 2014, Part I*, volume 8616 of *LNCS*, pages 480–499. Springer, Heidelberg, August 2014.

CDPW07. Ran Canetti, Yevgeniy Dodis, Rafael Pass, and Shabsi Walfish. Universally composable security with global setup. In Salil P. Vadhan, editor, *TCC 2007*, volume 4392 of *LNCS*, pages 61–85. Springer, Heidelberg, February 2007.

CFY17. Robert K. Cunningham, Benjamin Fuller, and Sophia Yakoubov. Catching MPC cheaters: Identification and openability. In Junji Shikata, editor, *ICITS 17*, volume 10681 of *LNCS*, pages 110–134. Springer, Heidelberg, November / December 2017.

CR03. Ran Canetti and Tal Rabin. Universal composition with joint state. In Dan Boneh, editor, *CRYPTO 2003*, volume 2729 of *LNCS*, pages 265–281. Springer, Heidelberg, August 2003.

CS03. Jan Camenisch and Victor Shoup. Practical verifiable encryption and decryption of discrete logarithms. In Dan Boneh, editor, *CRYPTO 2003*, volume 2729 of *LNCS*, pages 126–144. Springer, Heidelberg, August 2003.

CSS12. T.-H. Hubert Chan, Elaine Shi, and Dawn Song. Privacy-preserving stream aggregation with fault tolerance. In Angelos D. Keromytis, editor, *FC 2012*, volume 7397 of *LNCS*, pages 200–214. Springer, Heidelberg, February / March 2012.

DF90. Yvo Desmedt and Yair Frankel. Threshold cryptosystems. In Gilles Brassard, editor, *CRYPTO'89*, volume 435 of *LNCS*, pages 307–315. Springer, Heidelberg, August 1990.

DHMR07. Vanesa Daza, Javier Herranz, Paz Morillo, and Carla Ràfols. CCA2-secure threshold broadcast encryption with shorter ciphertexts. In Willy Susilo, Joseph K. Liu, and Yi Mu, editors, *ProvSec 2007*, volume 4784 of *LNCS*, pages 35–50. Springer, Heidelberg, November 2007.

DKSW09. Yevgeniy Dodis, Jonathan Katz, Adam Smith, and Shabsi Walfish. Composability and on-line deniability of authentication. In Omer Reingold, editor, *TCC 2009*, volume 5444 of *LNCS*, pages 146–162. Springer, Heidelberg, March 2009.

EDG14. Tariq Elahi, George Danezis, and Ian Goldberg. PrivEx: Private collection of traffic statistics for anonymous communication networks. In Gail-Joon Ahn, Moti Yung, and Ninghui Li, editors, *ACM CCS 14*, pages 1068–1079. ACM Press, November 2014.

ElG84. Taher ElGamal. A public key cryptosystem and a signature scheme based on discrete logarithms. In G. R. Blakley and David Chaum, editors, *CRYPTO'84*, volume 196 of *LNCS*, pages 10–18. Springer, Heidelberg, August 1984.

GGH+13. Sanjam Garg, Craig Gentry, Shai Halevi, Mariana Raykova, Amit Sahai, and Brent Waters. Candidate indistinguishability obfuscation and functional encryption for all circuits. In *54th FOCS*, pages 40–49. IEEE Computer Society Press, October 2013.

GMR88. Shafi Goldwasser, Silvio Micali, and Ronald L. Rivest. A digital signature scheme secure against adaptive chosen-message attacks. *SIAM Journal on Computing*, 17(2):281–308, April 1988.

GW09. Craig Gentry and Brent Waters. Adaptive security in broadcast encryption systems (with short ciphertexts). In Antoine Joux, editor, *EUROCRYPT 2009*, volume 5479 of *LNCS*, pages 171–188. Springer, Heidelberg, April 2009.

KPTZ13. Aggelos Kiayias, Stavros Papadopoulos, Nikos Triandopoulos, and Thomas Zacharias. Delegatable pseudorandom functions and applications. In Ahmad-Reza Sadeghi, Virgil D. Gligor, and Moti Yung, editors, *ACM CCS 13*, pages 669–684. ACM Press, November 2013.

MW16. Pratyay Mukherjee and Daniel Wichs. Two round multiparty computation via multi-key FHE. In Marc Fischlin and Jean-Sébastien Coron, editors, *EUROCRYPT 2016, Part II*, volume 9666 of *LNCS*, pages 735–763. Springer, Heidelberg, May 2016.

RN10. Vibhor Rastogi and Suman Nath. Differentially private aggregation of distributed time-series with transformation and encryption. In Ahmed K. Elmagarmid and Divyakant Agrawal, editors, *Proceedings of the ACM SIGMOD International Conference on Management of Data, SIGMOD 2010, Indianapolis, Indiana, USA, June 6-10, 2010*, pages 735–746. ACM, 2010.

Row05. J.K. Rowling. *Harry Potter and the Half-Blood Prince*. Bloomsbury, 2005.

SCR+11. Elaine Shi, T.-H. Hubert Chan, Eleanor G. Rieffel, Richard Chow, and Dawn Song. Privacy-preserving aggregation of time-series data. In *NDSS 2011*. The Internet Society, February 2011.

Sha79. Adi Shamir. How to share a secret. *Communications of the Association for Computing Machinery*, 22(11):612–613, November 1979.

SW14. Amit Sahai and Brent Waters. How to use indistinguishability obfuscation: deniable encryption, and more. In David B. Shmoys, editor, *46th ACM STOC*, pages 475–484. ACM Press, May / June 2014.

Zha14. Mark Zhandry. Adaptively secure broadcast encryption with small system parameters. Cryptology ePrint Archive, Report 2014/757, 2014. `http://eprint.iacr.org/2014/757`.

A Threshold Encryption Scheme: Threshold ElGamal

One simple example of a (non ad hoc) threshold encryption scheme is the threshold ElGamal scheme TEG due to Desmedt and Frankel [DF90], described in Figure 7. TEG is defined over a group $\mathcal{G}$ of prime order p with generator g in which the decisional Diffie-Hellman problem is assumed to be hard.

$\mathsf{Setup}(1^\lambda, t)$:
- Pick a secret key $sk \leftarrow^{\$} [p]$, and a random polynomial f of degree t with sk as its y-intercept.
- Return $pk = g^{sk}, \mathsf{msk} = f$.

$\mathsf{KeyGen}(\mathsf{msk})$:
- Pick a random $i \leftarrow^{\$} [1, \ldots, p-1]$.
- Return $sk_i = f(i)$.

$\mathsf{Enc}(pk, m \in \mathcal{G})$:
- Pick a random $y \in [p]$.
- $u = g^y$.
- $v = (pk)^y m$.
- Return $c = (u, v)$.

$\mathsf{PartDec}(sk_j, c = (u, v))$:
- Return $d_j = u^{sk_j}$.

$\mathsf{FinalDec}(\{d_i\}_{i \in \mathcal{R}' \subseteq \mathcal{R}}, c = (u, v))$:
- Interpolate the partial decryptions in the exponent to get u^{sk}: $y = \prod_{i \in \mathcal{R}' \subseteq \mathcal{R}} d_j^{\lambda_i}$, where λ_i is the appropriate Lagrange coefficient. (The Lagrange coefficients used depend on the identities i of the parties who participate. However, given that a certain threshold of parties do, the Lagrange coefficients do not affect the output.)
- Return $m = \frac{v}{y}$.

$\mathsf{Eval}(pk, c_1 = (u_1, v_1), c_2 = (u_2, c_2), \times)$:
- $u = u_1 u_2$
- $v = v_1 v_2$
- Return $c = (u, v)$.

Fig. 7: Threshold ElGamal Multiplicatively Homomorphic Encryption Scheme (TEG)

Lemma 1. *Threshold ElGamal is (n, t)-statically secure (Definition 5, modified to use pk instead of $\{pk_i\}_{i \in \mathcal{U}}$) for any polynomial n, t as long as the Decisional Diffie-Hellman (DDH) assumption holds in $\mathcal{G}$.*

Informally, this lemma follows by a standard reduction from the DDH assumption.

Lemma 2. *Threshold ElGamal is (n, t)-partial decryption simulatable (Definition 4, modified to use pk instead of $\{pk_i\}_{i\in\mathcal{U}}$) as long as the Decisional Diffie-Hellman assumption holds in $\mathcal{G}$.*

Informally, this lemma follows since partial decryptions can easily be simulated by interpolation in the exponent.

B Lower Bounds on Ciphertext Size for $\mathcal{R}$-Oblivious Ad Hoc Threshold Encryption Schemes

Theorem 9. *In any $\mathcal{R}$-oblivious ad hoc threshold encryption scheme* (Setup, KeyGen, Enc, PartDec, FinalDec)*, the average size of a ciphertext c produced as*

$$(\mathsf{params}) \leftarrow \mathsf{Setup}(1^\lambda, t = 0)$$

$$\{(pk_i, sk_i) \leftarrow \mathsf{KeyGen}(\mathsf{params})\}_{i\in[u]}$$

$$\mathcal{R} \leftarrow \text{ a random size-}n\text{ subset of } [u]$$

$$c \leftarrow \mathsf{Enc}(\mathsf{params}, \{pk_i\}_{i\in\mathcal{R}}, m)$$

for any m in the message space is $O(\log_2 \binom{u}{n})$.

Proof. To see this, imagine that a challenger runs all four lines described above (that is, generates u key pairs, a random $\mathcal{R}$ and a ciphertext). The challenger then sends the key pairs to the adversary, whose task is to identify the keys belonging to $\mathcal{R}$. The challenger sends c to the adversary; the adversary attempts decryption with each key pair (for $t = 0$ one key pair is sufficient to decrypt), and identifies those for which decryption yields m as belonging to $\mathcal{R}$. Note that the probability of correct decryption with a key that does not belong to $\mathcal{R}$ should be negligible, or the threshold encryption scheme is not secure. This allows the adversary to learn $\mathcal{R}$. Since there are $\binom{u}{n}$ possibilities for $\mathcal{R}$, it should require at least $\log_2 \binom{u}{n}$ bits to communicate. Since the key pairs are generated independently of $\mathcal{R}$, they don't count towards those bits; thus, the ciphertext should be at least $\log_2 \binom{u}{n}$ bits long. In the case when $u = 2n$, this is lower-bounded by 2^n.

Theorem 10. *For any $t < n$, any $\mathcal{R}$-oblivious ad hoc threshold encryption scheme* (Setup, KeyGen, Enc, PartDec, FinalDec)*, the average size of a ciphertext c produced as*

$$(\mathsf{params}) \leftarrow \mathsf{Setup}(1^\lambda, t)$$

$$\{(pk_i, sk_i) \leftarrow \mathsf{KeyGen}(\mathsf{params})\}_{i\in[u]}$$

$$\mathcal{R} \leftarrow \text{ a random size-}n\text{ subset of } [u]$$

$$c \leftarrow \mathsf{Enc}(\mathsf{params}, \{pk_i\}_{i\in\mathcal{R}}, m)$$

for any m in the message space is $O(\log_2 \binom{u-t}{n-t})$.

Proof. The proof goes exactly as it does for Theorem 9, but the challenger identifies to the adversary t key pairs in $\mathcal{R}$, so that the adversary only needs to identify the $n-t$ remaining ones. The adversary does this by computing partial decryptions using the t identified key pairs, and testing each of the remaining $u-t$ key pairs. It tests a key pair by using it to generate a partial decryption and seeing whether that partial decryption combines with the ones it already has to produce m.

C Background: Secret Sharing

Secret sharing was introduced by Shamir [Sha79]. Informally, a t-out-of-n threshold secret sharing of a secret m is an encoding of the secret into n pieces, or *shares*, such that any $t+1$ shares together can be used to reconstruct the secret m, but t or fewer shares give no information at all about m. A secret sharing scheme SS consists of two algorithms: $\mathsf{SS.Share}$ and $\mathsf{SS.Reconstruct}$.

- $\mathsf{SS.Share}(n, t, m) \to ([m]_1, \ldots, [m]_n)$ takes in a secret m and produces the n secret shares.
- $\mathsf{SS.Reconstruct}([m]_{i_1}, \ldots, [m]_{i_{t+1}}) \to \tilde{m}$ takes in $t+1$ secret shares and returns the reconstructed secret $\tilde{m}$.

Informally, correctness requires that $\tilde{m} = m$, and privacy requires that given t or fewer shares of either m_R or m_L, no efficient adversary can guess which message was shared.

C.1 Share Simulatability

We additionally use a property which we call *share simulatability*, which requires that given t or fewer honestly generated shares of m_R and given m_L, there exists an efficient algorithm $\mathsf{SS.SimShares}$ which generates the rest of the shares in such a way that the resulding sharing is indistinguishable from a fresh sharing of m_L.

Definition 9 (Secret Sharing: Share Simulatability).
For $b \in \{0,1\}$, let $\mathsf{EXP}(\mathcal{A}, n, t, \mathsf{SimShares}, b)$ denote the game described in Figure 8 played with the adversary $\mathcal{A}$, number of shares n, threshold t, simulation algorithm $\mathsf{SimShares}$ and fixed b. Let $\mathsf{WinProb}(\mathcal{A}, n, t, \mathsf{SimShares}, b)$ denote the probability that the adversary $\mathcal{A}$ wins $\mathsf{EXP}(\mathcal{A}, n, t, \mathsf{SimShares}, b)$.

*A secret sharing scheme scheme $(\mathsf{SS.Share}, \mathsf{SS.Reconstruct})$ is (n,t)-*share simulatable *if there exists an efficient simulation algorithm $\mathsf{SS.SimShares}$ such that for all efficient adversaries $\mathcal{A}$, there exists a negligible function $negl$ such that*

$$|\mathsf{WinProb}(\mathcal{A}, n, t, \mathsf{SS.SimShares}, 0) - \mathsf{WinProb}(\mathcal{A}, n, t, \mathsf{SS.SimShares}, 1)| \leq \frac{1}{2} + negl(\lambda).$$

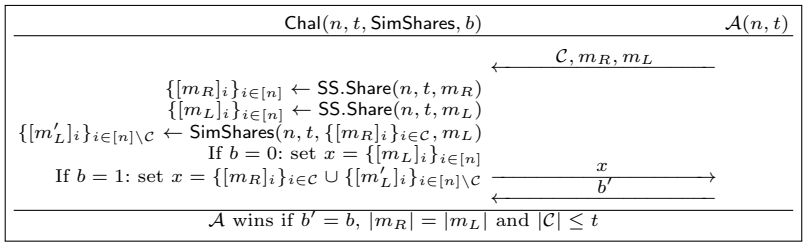

Fig. 8: Share Simulatability Game for Secret Sharing

C.2 Shamir Secret Sharing [Sha79]

Shamir t-out-of-n secret sharing (Shamir) uses degree-(t) polynomials over some field. $\mathsf{Shamir.Share}(n, t, m)$ generates a random degree-(t) polynomial f with m as its y-intercept; each share $[m]_i$ is a point $(x_i, f(x_i))$ on the polynomial (with $x_i \neq 0$). Any $t + 1$ shares can be used to interpolate the polynomial, reconstructing m. Any t or fewer shares give no information about m.

Shamir secret sharing is share simulatable; any t or fewer points can be interpolated with $(0, m_L)$ (and optionally with some additional random points) to obtain a degree-t polynomial.

Additionally, Shamir secret sharing is linearly homomorphic: a shared value m can be multiplied by a constant, or added to another shared value m', by separately operating on the individual shares.

D Proofs of Properties of the Share-and-Encrypt Ad Hoc Threshold Encryption Construction

In this appendix, we prove Theorem 1.

Theorem 11 (Restated from Theorem 1). *The share-and-encrypt ATE (Construction 1) is (n, t)-statically secure (Definition 5), as long as* SS *is a secure share simulatable t-out-of-n secret sharing scheme, and* PKE *is a CPA-secure public key encryption scheme.*

In order to prove Theorem 1, in Appendix D.1 we show that the share-and-encrypt ATE is (n, t)-statically semantically secure, and in Appendix D.2 we show that it is (n, t)-partial decryption simulatable.

D.1 Proof that Share-and-Encrypt is Statically Semantically Secure

Proof. We show a sequence of indistinguishable games between a static security challenger Chal and an adversary $\mathcal{A}$. The sequence starts with $\mathsf{EXP}(\mathcal{A}, \lambda, n, t, R)$ from Definition 3 (that is, a challenger who always uses $b = R$), and ends with $\mathsf{EXP}(\mathcal{A}, \lambda, n, t, L)$ (that is, a challenger who always uses $b = L$).

Game 1 This is the game as described in Figure 1 with $b = R$ (that is, this is $\mathsf{EXP}(\mathcal{A}, \lambda, n, t, R)$).

Game 2 This game is the same as the previous game, but when computing the challenge ciphertext c^*, Chal does the following:

- $\{[m_L]_i\}_{i\in\mathcal{R}} \leftarrow \mathsf{SS.Share}(n, t, m_L)$
- $\{[m_R]_i\}_{i\in\mathcal{R}\cap\mathcal{C}} \leftarrow \mathsf{SS.SimShares}(n, t, \{[m_L]_i\}_{i\in\mathcal{R}\setminus\mathcal{C}}, m_L)$
- $c^* \leftarrow \{\mathsf{PKE.Enc}(pk_i, [m_L]_i)\}_{i\in\mathcal{R}\setminus\mathcal{C}} \cup \{\mathsf{PKE.Enc}(pk_i, [m_R]_i)\}_{i\in\mathcal{R}\cap\mathcal{C}}$

If $\mathcal{A}$ can tell the difference between this game and the previous game, then we can design another adversary $\mathcal{B}$ that uses $\mathcal{A}$ to break the share simulatability property of the secret sharing scheme SS (Definition 9). $\mathcal{B}$ forwards $(m_R, m_L, \mathcal{R}\cap\mathcal{C})$ to the share simulatability challenger (Figure 8). Upon receiving shares from the share simulatability challenger, $\mathcal{B}$ encrypts them and sends them to $\mathcal{A}$. If the share simulatability challenger flips $b = 0$, $\mathcal{A}$'s view will be as in the previous game; if the share simulatability challenger flips $b = 1$, $\mathcal{A}$'s view will be as in this game. $\mathcal{B}$ sends the share simulatability challenger $b'' = 0$ if $\mathcal{A}$ submits $b' = R$, and $b'' = 1$ if $\mathcal{B}$ submits $b' = L$.

Game 3.idx **for** $\mathsf{idx} \in [1, \ldots, n-t]$

This game is the same as the previous game, but for the idxth honest party (without loss of generality, let that be P_i with public key pk_i), the challenger encrypts $[m_L]_i$.

Game 3.1 is indistinguishable from Game 2, and Game 3.idx is indistinguishable from Game 3.$(\mathsf{idx}-1)$ for $\mathsf{idx} \in [2, \ldots, n-t]$, by the CPA security of the public key encryption scheme PKE. If $\mathcal{A}$ can tell the difference between these games, then we can design another adversary $\mathcal{B}$ that uses $\mathcal{A}$ to break the CPA security of PKE.

$\mathcal{B}$ honestly generates all of the PKE keys except for the idxth honest key pair. $\mathcal{B}$ talks to a CPA PKE challenger to get pk_i. It then does everything as before, except it sends $m_0 = [m_R]_i$ and $m_1 = [m_L]_i$ to the CPA challenger, and uses the challenge ciphertext it gets as part of c^*. Note that when the CPA challenger uses $b = 0$ we are in the previous game, and when the CPA challenger uses $b = 1$ we are in this game. $\mathcal{B}$ passes on the guess made by $\mathcal{A}$ to the CPA challenger.

Note that Game 3.$(n-t)$ is $\mathsf{EXP}(\mathcal{A}, \lambda, n, t, L)$.

D.2 Proof that Share-and-Encrypt is Partial Decryption Simulatable

Proof. $\mathsf{SimPartDec}$ can simulate partial decryptions in the share-and-encrypt ad hoc threshold encryption scheme in Construction 1 simply by running $\{d_i\}_{i\in\mathcal{R}\setminus\mathcal{C}} \leftarrow \mathsf{SS.SimShares}(\{d_i\}_{i\in\mathcal{R}\cap\mathcal{C}}, m_R)$, and returning $\{d_i\}_{i\in\mathcal{R}\setminus\mathcal{C}}$.

Any adversary $\mathcal{A}$ who can win the static partial decryption simulatability game described in Figure 2 when played with $\mathsf{SimPartDec}$ with non-negligible probability can be used to break the share simulatability of SS and win the share simulatability game described in Figure 8.

E Share-and-Encrypt HATE Instantiations

In this appendix, we instantiate the share-and-encrypt HATE (Construction 1) in two ways.

E.1 Shamir-and-ElGamal

We build share-and-encrypt HATE out of ElGamal encryption [ElG84] and a variant of Shamir secret sharing. We need to use a *variant* of Shamir secret sharing (which we call exponential Shamir secret sharing), and not Shamir secret sharing itself, because Shamir secret sharing is additively homomorphic (and the homomorphism is applied via addition of individual shares), but ElGamal is multiplicatively homomorphic (and the homomorphism is applied via multiplication of ciphertexts), so if we attempt to apply a homomorphism on encrypted shares, it will not work. What we need in order to get an additively homomorphic ATE scheme is to use ElGamal encryption with a secret sharing scheme which is additively homomorphic, but whose homomorphism is applied via multiplication. Therefore, we need to alter our Shamir secret sharing scheme by moving the shares to the exponent; then, taking a product of two shares will result in a share of the sum of the two shared values. Below we describe the ElGamal encryption scheme and the exponential Shamir secret sharing scheme which we use.

ElGamal Multiplicatively Homomorphic Encryption. Figure 9 describes the ElGamal multiplicatively homomorphic encryption scheme (EG). Note that we split the key generation algorithm into two algorithms: Setup and KeyGen. This is because when we use ElGamal as part of our HATE scheme, it is important that all parties share the same modulus and generator, so we factor out part of KeyGen into Setup, which will only be run once globally.

Exponential Shamir Secret Sharing. Figure 10 describes the exponential Shamir secret sharing scheme (EShamir).

Notice that the reconstruction uses brute force search; this means that this secret sharing scheme can only be used for very small (polynomial-size in λ) message spaces. However, HATE is interesting even in this setting. For instance, if all we want to do is take a poll by summing encryptions of 0s and 1s, this HATE scheme enables us to do it. It is reasonable to assume that the server can manage to do brute force search over a polynomial space, since it is already doing quadratic work in this computation.

Lemma 3. *The exponential Shamir secret sharing scheme (*EShamir*) described in Figure 10 is share simulatable.*

Proof. Informally, given a message m and t or fewer shares, we obtain correctly distributed remaining shares by interpolating the given with $(0, g^m)$ (and possibly with random values, if fewer than t shares are provided) in the exponent. This is done in a manner similar to the first step of reconstruction.

$\mathsf{Setup}(1^\lambda)$**:**
- Pick a prime-order group $\mathcal{G}$ with generator g in which the decisional Diffie-Hellman problem is assumed to be hard. Let p be the order of that group.
- Publish $\mathsf{params} = (\mathcal{G}, p, g)$.

$\mathsf{KeyGen}(\mathsf{params})$**:**
- Pick a random $sk \in [p]$.
- Publish $pk = g^{sk}$.

$\mathsf{Enc}(\mathsf{params}, pk, m \in \mathcal{G})$**:**
- Pick a random $y \in [p]$.
- $u = g^y$.
- $v = (pk)^y m$.
- Return $c = (u, v)$.

$\mathsf{Dec}(\mathsf{params}, sk, c = (u, v))$**:**
- Return $m = \frac{v}{u^{sk}}$.

$\mathsf{Eval}(\mathsf{params}, c_1 = (u_1, v_1), c_2 = (u_2, c_2), \times)$**:**
- $u = u_1 u_2$.
- $v = v_1 v_2$.
- Return $c = (u, v)$.

Fig. 9: ElGamal Multiplicatively Homomorphic Public Key Encryption Scheme (EG) [ElG84]

$\mathsf{Setup}(1^\lambda)$**:** same as $\mathsf{EG.Setup}$.

$\mathsf{Share}(n, t \leq n, m \in \mathbb{Z}_p)$**:**
- Pick a random degree-t polynomial f in $\mathbb{Z}_p$ which has m as its y-intercept. This can be done by picking t random coefficients $\mathsf{coef}_1, \ldots, \mathsf{coef}_{t-1} \in \mathcal{G}$, and setting $f(x) = m + \sum_{j=1}^{t-1} \mathsf{coef}_j x^j$.
- Return $\{\mathsf{Share}_i\}_{i \in [1, \ldots n]}$ where $\mathsf{Share}_i = (i, g^{f(i)})$.

$\mathsf{Reconstruct}(\{\mathsf{Share}_i = (i, y_i)\}_{i \in \mathcal{R}' \subseteq [n]})$**:**
- Perform polynomial interpolation over the shares in the exponent to recover $g^{\tilde{m}}$. As long as $|\mathcal{R}'| > t$, this can be done by throwing out values in $\mathcal{R}'$ until $|\mathcal{R}'| = t + 1$, and doing the following:

$$g^{\tilde{m}} = \prod_{i \in \mathcal{R}'} y_i^{\prod_{j \in \mathcal{R}', j \neq i} \frac{j}{j-i}}$$

- Recover $\tilde{m}$ by brute force search.

$\mathsf{Eval}(\mathsf{Share}_i = (i, y_i), \mathsf{Share}'_i = (i, y'_i), +)$**:** $\mathsf{Share}_i^+ = (i, y_i y'_i)$

Fig. 10: Exponential Shamir Secret Sharing Scheme ($\mathsf{EShamir}$)

E.2 CRT-and-Paillier

We also build share-and-encrypt HATE out of Camenisch-Shoup encryption and Chinese Remainder Theorem based secret sharing. Unlike Shamir-and-ElGamal (Section E.1), this HATE allows us to use large message spaces.

CS Additively Homomorphic Encryption. We use a slightly modified version of the Paillier-style verifiable encryption scheme described by Camenisch and Shoup [CS03].[6] Figure 11 describes the this scheme. Our modifications consist solely of removing elements from the ciphertext, so the modified scheme naturally inherits the CPA security of the original (but not its CCA security).[7].

$\mathsf{KeyGen}(1^\lambda)$**:**
- Let $N = pq$ where $p = 2p' + 1$ and $q = 2q' + 1$, and p' and q' are λ-bit primes.
- Let $h = 1 + N$.
- Choose a random $g' \in Z^*_{N^2}$, and set $g = (g')^{2N} \bmod N^2$. (g is a generator of a size-$p'q'$ subgroup with high probability.)
- Choose a random secret key $sk \in \{1, \ldots, \lfloor (N^2)/4 \rfloor\}$.
- Return $pk = g^{sk} \bmod N^2$.

$\mathsf{Enc}(pk, m)$**:**
- Choose a random $r \in [N/4]$.
- Return $c = (g^r \bmod N^2, pk^r h^m \bmod N^2)$.

$\mathsf{Dec}(sk, c = (u, v))$**:**
- $z = \frac{v}{u^{sk}} \bmod N^2$. (Note that $z = h^m$ if c is an encryption of m.)
- Return $m' = \frac{z-1}{N}$ with division over integers. (Note that m' is the discrete log of z w.r.t. h.)

Fig. 11: Camenisch-Shoup Additively Homomorphic Encryption Scheme (CS). We omit Setup, since this scheme does not require setup.

CRT Secret Sharing. We use a classic secret sharing scheme based on the Chinese Remainder Theorem, which allows each party to operate homorphically on shares in a different group. This version is due to Asmuth and Bloom [AB06]. We describe it in Figure 12.

The scheme is perfectly correct. Furthermore, it supports a limited number (currently set to n) of homomorphic additions. The setting of parameters in the setup phase in Figure 12 ensures that $n \cdot A \leq N_+$, where each individual sharing corresponds to a vector of modular reductions of an integer less than A. This means that n sharings, added coordinate-wise, will lead to the reconstruction

[6] Their scheme is designed it to be secure against chosen ciphertext attacks, which is unnecessary for our purposes.

[7] A similarly modified version of this scheme was used by Cunningham *et al.* [CFY17]

of an integer less than $n \cdot A$. Every set of more than t shares contains enough information for that reconstruction.

The scheme's statistical security relies on the requirement that each unauthorized set of shares can reconstruct the secret integer $\hat{a}$ only modulo some integer that is (a) relatively prime to N_0 and (b) at most N_-, which itself is at most $\frac{A}{N_0 2^k}$. By the following lemma, those conditions ensure that the view of any unauthorized set is within statistical difference 2^{-k} of uniform.

Lemma 4. *Let a, n, t be positive integers, and let A be uniformly random in $\{a' \in [a] : a' \bmod n = t\}$. Then for all positive integers $m < a$ that are relatively prime to n, the distribution of the random variable $B = A \bmod m$ is within statistical difference $\frac{nm}{a}$ of uniform.*

Proof. Consider the number of $a' \in [a]$ that solve both the equations $a' \bmod n = t$ and $a' \bmod m = u$ (for some u). Since m and n are relatively prime, this system is equivalent to $a' \bmod mn = v$ for some particular v. The number of solutions to this is $\lfloor \frac{a}{mn} \rfloor$ or $\lceil \frac{a}{mn} \rceil$. Thus, the probability that $B = u$ is always with $1 \pm \frac{mn}{a}$ of a uniform element of $\mathbb{Z}_m$. The total variation distance from uniform is thus at most $\frac{mn}{a}$.

Lemma 5. *The CRT secret sharing scheme (*CRTss*) described in Figure 12 is share simulatable.*

Proof. Recall the share simulatability game from Figure 8. On input a set of unauthorized shares $\{\mathsf{Share}_i\}_{i \in \mathcal{C}}$ which were created as a sharing of m_L, and a target message m_R, first find a nonegative integer $a < N_0 \prod_{i \in \mathcal{C}} N_i$ such that $a \bmod N_0 = m_R$ and $a \bmod N_i = \mathsf{Share}_i$ for $i \in \mathcal{C}$. Such an integer exists since the moduli are all relatively prime. Next, select a random $\hat{a}_R \in [A]$ such that $\hat{a}_R \bmod (N_0 \prod_{i \in \mathcal{C}} N_i) = a$. The correctness condition of the secret sharing scheme implies that $A > N_0 \prod_{i \in \mathcal{C}} N_i$, so this step is always possible. Finally, we produce the new shares as $\mathsf{Share}'_i = \hat{a} \bmod N_i$ for $i \in \mathcal{R} \setminus \mathcal{C}$.

By Lemma 4 above, the distribution of t or fewer shares of m_L are statistically indistinguishable from the corresponding distibution for m_R. The share simulation algorithm above selects a uniformly random sharing of m_R that is consistent with the unauthorized shares of m_L. The joint distribution is therefore statistically close to that of a fresh sharing of m_L.

F Security of the Obfuscation-Based Ad Hoc Threshold Encryption Construction

In order to make the threshold encryption definitions play nice with the obfuscation-based ATE (Construction 2), we need to alter the static semantic security game (and partial decryption simulatability game) in two ways.

First, we need to make the game even more static (what we call *super-static*) by forcing the adversary to commit not only to the set $\mathcal{C}$ of corrupt parties, but

$\mathsf{Share}(n, N_1, ..., N_n, N_0, t \leq n, m \in Z_{N_0}, 1^\lambda)$:

- Let $\begin{cases} N_+ \stackrel{\text{def}}{=} \min_{J \subseteq [n]:|J|=t+1} \prod_{i \in J} N_i \\ A \stackrel{\text{def}}{=} \lfloor N_+/n \rfloor \\ N_- \stackrel{\text{def}}{=} \max_{J \subseteq [n]:|J|=t} \prod_{i \in J} N_i \end{cases}$.
- If $N_0 \cdot 2^\lambda \cdot N_- > A$ then stop and return "Error: message space too large for λ bits of security."
- Select $a \in_R \{a' \in [A] : a' \bmod N_0 = m\}$
- Return $(\mathsf{Share}_1, \ldots, \mathsf{Share}_n)$ where $\mathsf{Share}_i = (i, a \bmod N_i)$.

$\mathsf{Reconstruct}(\mathsf{Share}_{i_1} = (i_1, y_{i_1}), \ldots, \mathsf{Share}_{i_t} = (i_t, y_{i_t}))$:

- Find the unique $\hat{a} \in Z_{\prod_j N_j}$ such that $\hat{a} \equiv y_i \mod N_i$ for all $i \in \{i_1, ..., i_t\}$.
- Return $\hat{m} = \hat{a} \bmod N_0$.

$\mathsf{Eval}(+, \mathsf{Share}_i = (i, y), \mathsf{Share}'_i = (i, y'))$: $\mathsf{Share}^+_i = (i, y + y' \bmod N_i)$

Fig. 12: Chinese Remainder Secret Sharing Scheme (CRTss). We omit Setup, since this scheme does not require setup.

also the set $\mathcal{R}$ of challenge ciphertext recipients. This is necessary for the proof of Theorem 4. Note that when $\mathcal{U} = \mathcal{R}$ (that is, when all parties are recipients), super-static security is equivalent to static security.

Second, since Construction 2 is a keyed-sender scheme, we need to (a) provide the adversary with the sender public key and use the corresponding sender secret key to encrypt, and (b) in order to get CPA security, we need to allow the adversary to make multiple encryption queries, since encryption is no longer a public operation. In typical public key encryption a game which allows multiple encryption queries is equivalent to one that does not, but this is not true in a keyed-sender setting.

The new super-static semantic security game for keyed-sender ATE is described in Figure 13. (and the correspondingly modified super-static partial decryption simulatability game is described in Figure 14).

Definition 10 (super-static semantic security for keyed-sender ad hoc threshold encryption).

For $b \in \{R, L\}$, let $\mathsf{EXP}(\mathcal{A}, \lambda, n, t, b)$ denote the game described in Figure 13 played with the adversary $\mathcal{A}$, security parameter λ, number of existing parties $|\mathcal{U}| = n$, threshold t and fixed b. Let $\mathsf{WinProb}(\mathcal{A}, \lambda, n, t, b)$ denote the probability that the adversary $\mathcal{A}$ wins $\mathsf{EXP}(\mathcal{A}, \lambda, n, t, b)$.

A keyed-sender ad hoc threshold encryption scheme (Setup, KeyGen, Enc, PartDec, FinalDec) *is* (n, t)-super-statically semantically secure *if for all efficient adversaries $\mathcal{A}$ there exists a negligible function negl such that*

$$|\mathsf{WinProb}(\mathcal{A}, \lambda, n, t, R) - \mathsf{WinProb}(\mathcal{A}, \lambda, n, t, L)| \leq \frac{1}{2} + negl(\lambda).$$

Definition 11 (super-static partial decryption simulatability for keyed-sender ad hoc threshold encryption).

For $b \in \{R, L\}$, let $\mathsf{EXP}(\mathcal{A}, \lambda, n, t, b)$ denote the game described in Figure 14 played with the adversary $\mathcal{A}$, security parameter λ, number of existing parties

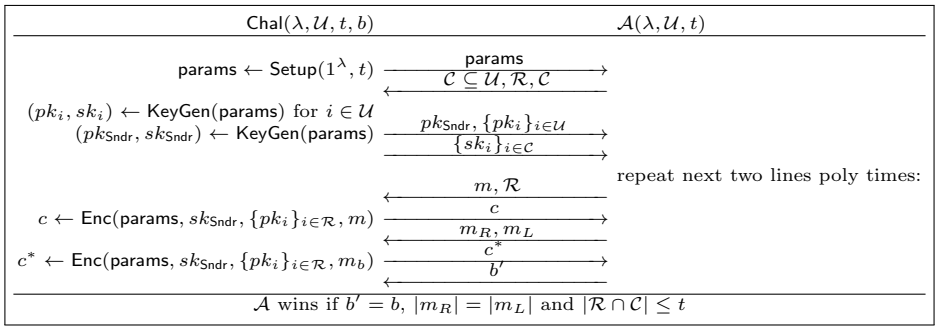

Fig. 13: Super-Static Semantic Security Game for Keyed-Sender Ad Hoc Threshold Encryption.

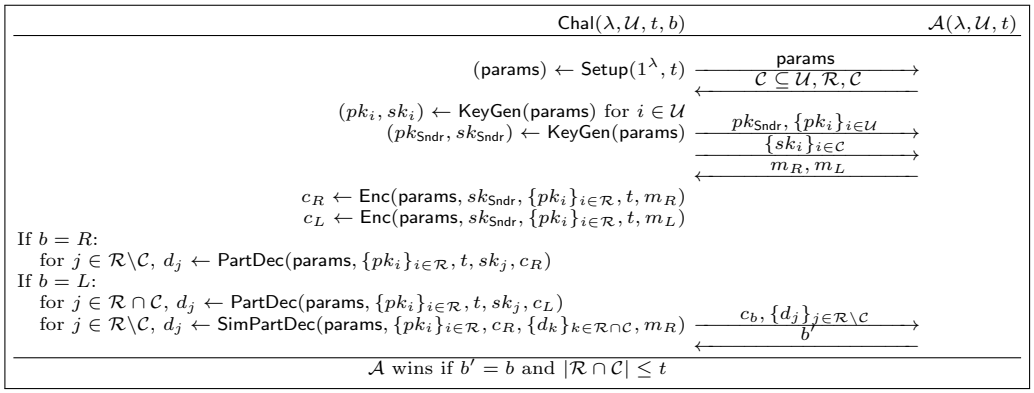

Fig. 14: Super-Static Partial Decryption Simulatability Game for Keyed-Sender Ad Hoc Threshold Encryption

$|\mathcal{U}| = n$, *threshold* t *and fixed* b. *Let* $\mathsf{WinProb}(\mathcal{A}, \lambda, n, t, b)$ *denote the probability that the adversary* $\mathcal{A}$ *wins* $\mathsf{EXP}(\mathcal{A}, \lambda, n, t, b)$.

A keyed-sender ad hoc threshold encryption scheme (Setup, KeyGen, Enc, PartDec, FinalDec) *is* (n, t)-super-statically semantically secure *if there exists an efficient algorithm* SimPartDec *such that for all efficient adversaries* $\mathcal{A}$ *there exists a negligible function negl such that*

$$|\mathsf{WinProb}(\mathcal{A}, \lambda, n, t, R) - \mathsf{WinProb}(\mathcal{A}, \lambda, n, t, L)| \leq \frac{1}{2} + negl(\lambda).$$

Definition 12 (super-static security for keyed-sender ad hoc threshold encryption). *A keyed-sender ad hoc threshold encryption scheme* (Setup, KeyGen, Enc, PartDec, FinalDec) *is* (n, t)-super-statically secure *if it is both* (n, t)-super-statically semantically secure *(Definition 10) and* (n, t)-super-statically partial decryption simulatable *(Definition 11).*

Theorem 12 (Restated from Theorem 4). *The obfuscation-based ATE (Construction 2) is (n,t)-super-statically secure (Definition 12) for any polynomial n,t, as long as* iO *is a secure indistinguishability obfuscator,* PPRF *is a secure puncturable* PRF *with range $\mathbb{Z}_p$,* SIG *is a constrained signature scheme, and* PRG *is a secure pseudorandom generator with domain $\{0,1\}^\lambda$ and range in $\{0,1\}^{2\lambda}$.*

In order to prove Theorem 4, we must show the obfuscation-based Homomorphic Ad Hoc Threshold Encryption construction is super-statically semantically secure and super-statically partial decryption simulatable. We prove super-static semantic security in Appendix F.1; we prove partial decryption simulatability in Appendix F.2.

F.1 Proof that Obfuscation-Based Homomorphic Ad Hoc Threshold Encryption Share-and-Encrypt is Super-Statically Semantically Secure

Notation. In our sequence of games, we use $c^* = (\mathsf{nonce}^*, e^*, \sigma^*)$ to denote the challenge ciphertext. In particular, nonce^* denotes the value to which the PPRF is applied in order to generate the mask w^*, and e^* denotes the "one time pad" encryption of the challenge message with the generated mask ($e^* = (m^* + w^*) \bmod p$).

We use $\overrightarrow{pv}$ to denote a vector of public values. $\overrightarrow{pv}$ is always assumed to be ordered lexicographically (that is, all programs implicitly reject vectors of public values which are out of order).

Proof. We show a sequence of indistinguishable games between a super-static semantic security challenger Chal and an adversary $\mathcal{A}$. The sequence starts with $\mathsf{EXP}(\mathcal{A}, \lambda, n, t, R)$ from Definition 10 (that is, a challenger who always uses $b = R$), and ends with $\mathsf{EXP}(\mathcal{A}, \lambda, n, t, L)$ (that is, a challenger who always uses $b = L$). We summarize this sequence of games in Figure 16. Instead of showing all of the games, we show that the first game is indistinguishable from a game where the challenge ciphertext encrypts a random message; to get to the last game, the shown sequence of games is reversed with m_L instead of m_R.

For the sake of simplicity, we assume that the number of corrupt challenge ciphertext recipients $|\mathcal{R}^* \cap \mathcal{C}|$ is always t; security in this case trivially implies security for smaller $\mathcal{R}^* \cap \mathcal{C}$.

Game 4 This is $\mathsf{EXP}(\mathcal{A}, \lambda, n, t, R)$ as described in Definition 13 and Figure 13.

Game 5 In this game, when creating the sender's public key (specifically, the sender's signature verification key $\mathsf{SIG}.pk$), the challenger Chal generates a constrained verification key instead of a regular one. The constraint is designed to ensure that the challenge nonce nonce^* on which the PPRFs are called can only ever be associated with the challenge recipient set. That is, the challenger picks nonce^* at random when generating the sender public key, and sets the signature constraint to be

$$C(\overrightarrow{pv}, \mathsf{nonce}) = \begin{cases} 0, & \text{if } \mathsf{nonce} = \mathsf{nonce}^* \text{ and } \overrightarrow{pv} \neq \{pv_i\}_{i \in \mathcal{R}^*} \\ 1, & \text{otherwise} \end{cases}$$

When computing the challenge ciphertext c^*, Chal uses the nonce nonce^*. Game 5 is indistinguishable from Game 4 by the security of constrained signatures. Since the challenger chooses each new nonce at random, and with overwhelming probability will never sign anything not satisfying the constraint (since no nonce other than the challenge nonce will be equal to nonce^*), then if the adversary can distinguish Game 5 from Game 4, then the challenger can use that adversary to distinguish between a constrained and unconstrained public verification key.

Game 6.i **for** $i \in [1, \ldots, n-t]$
In this game, when choosing the ith honest party's public value pv, the challenger chooses it at random from $\{0,1\}^{2\lambda}$, so that with overwhelming probability it has no preimages relative to the pseudorandom generator PRG. Game 6.1 is indistinguishable from Game 5, and Game 6.i is indistinguishable from Game 6.$(i-1)$ for $i \in [2, \ldots, n-t]$, by the security of pseudorandom generators.
We let Game 6 denote Game 6.$(n-t)$.

Game 7 In this game, the challenger Chal changes the way encryption takes place in the obfuscated program, as shown in Algorithm 2. In particular, instead of computing $[m]_{\mathsf{idx}} = (m + \sum_{j \in [1,\ldots,t]} \mathsf{coef}_j \mathsf{idx}^j) \bmod p$, the program now computes $[-w]_{\mathsf{idx}} = (-w + \sum_{j \in [1,\ldots,t]} \mathsf{coef}_j \mathsf{idx}^j) \bmod p$, followed by $[m]_{\mathsf{idx}} = ([-w]_{\mathsf{idx}} + e) \bmod p$. Note that this gives the exact same share as computing $[m]_{\mathsf{idx}}$ directly; in fact, the only reason we did not use these instructions explicitly to begin with is clarity.

Algorithm 2 $f^{\mathsf{Game}\ 7}_{k_{\mathsf{Dec}}, k_{\mathsf{Share}}, \mathsf{SIG}.pk}(\overrightarrow{pv}, \mathsf{idx}, sv, c)$

if $(\overrightarrow{pv}[\mathsf{idx}] = \mathsf{PRG}(sv))$ and $(\mathsf{SIG.Verify}(\mathsf{SIG}.pk, (\overrightarrow{pv}, \mathsf{nonce}), \sigma))$ **then**
  $w \leftarrow \mathsf{PPRF}_{k_{\mathsf{Dec}}}(\mathsf{nonce})$
  $m = (e - w) \bmod p$
  for $j \in [1, \ldots, t]$ **do**
    $\mathsf{coef}_j = \mathsf{PPRF}_{k_{\mathsf{Share},j}}(\mathsf{nonce})$
  $[-w]_{\mathsf{idx}} = (-w + \sum_{j \in [1,\ldots,t]} \mathsf{coef}_j \mathsf{idx}^j) \bmod p$ {This gives the idxth Shamir share of $-w$}
  $[m]_{\mathsf{idx}} = ([-w]_{\mathsf{idx}} + e) \bmod p$ {This gives the idxth Shamir share of m}
  return $[m]_{\mathsf{idx}}$

Game 7 is indistinguishable from Game 6 by the security of indistinguishability obfuscation; the programs have identical input-output behavior.

Game 8 In this game, when creating the sender's public key (specifically, the obfuscation of Algorithm 1 it contains), the challenger punctures the PPRF

keys k_{Dec} and $k_{\mathsf{Share},j}$ for all $j \in [1,\ldots,t]$ at nonce^*. To preserve the input-output behavior of the program, Chal computes $w^* = \mathsf{PPRF}_{k_{\mathsf{Dec}}}(\mathsf{nonce}^*)$ and $\mathsf{coef}_j^* = \mathsf{PPRF}_{k_{\mathsf{Share},j}}(\mathsf{nonce}^*)$ for $j \in [1,\ldots,t]$, and modifies the program to set $w = w^*$ and $\mathsf{coef}_j = \mathsf{coef}_j^*$ when $\mathsf{nonce} = \mathsf{nonce}^*$, as shown in Algorithm 3.

Algorithm 3 $f^{\mathsf{Game}\ 8}_{k_{\mathsf{Dec}}\{\mathsf{nonce}^*\},k_{\mathsf{Share}}\{\mathsf{nonce}^*\},\mathsf{SIG}.pk,\mathsf{nonce}^*,w^*,\mathsf{coef}_1^*,\ldots,\mathsf{coef}_t^*}(\overrightarrow{pv},\mathsf{idx},sv,c)$

if $(\overrightarrow{pv}[\mathsf{idx}] = \mathsf{PRG}(sv))$ and $(\mathsf{SIG.Verify}(\mathsf{SIG}.pk,(\overrightarrow{pv},\mathsf{nonce}),\sigma))$ **then**
  if $\mathsf{nonce} = \mathsf{nonce}^*$ **then**
    $w = w^*$
    for $j \in [1,\ldots,t]$ **do**
      $\mathsf{coef}_j = \mathsf{coef}_j^*$
  else
    $w \leftarrow \mathsf{PPRF}_{k_{\mathsf{Dec}}}(\mathsf{nonce})$
    for $j \in [1,\ldots,t]$ **do**
      $\mathsf{coef}_j = \mathsf{PPRF}_{k_{\mathsf{Share},j}}(\mathsf{nonce})$
  $[-w]_{\mathsf{idx}} = (-w + \sum_{j\in[1,\ldots,t]} \mathsf{coef}_j \mathsf{idx}^j) \bmod p$ {This gives the idxth Shamir share of $-w$}
  $[m]_{\mathsf{idx}} = ([-w]_{\mathsf{idx}} + e) \bmod p$ {This gives the idxth Shamir share of m}
  return $[m]_{\mathsf{idx}}$

Game 8 is indistinguishable from Game 7 by the security of indistinguishability obfuscation; the programs have identical input-output behavior.

Game 9.j for $j \in [1,\ldots,t]$
In this game, the challenger Chal chooses coef_j^* truly at random.

Game 9.1 is indistinguishable from Game 8, and Game 9.j is indistinguishable from Game 9.$(j-1)$ for $j \in [2,\ldots,t]$, by the security of puncturable PRFs.

We let Game 9 denote Game 9.t.

Game 10 In this game, the challenger Chal chooses w^* truly at random.

Chal also computes the one time pad component of the challenge ciphertext e^* as $e^* = (m_R + w^*) \bmod p$.

Game 10 is indistinguishable from Game 9 by the security of puncturable PRFs.

Game 11 In this game, the challenger Chal modifies the obfuscated program to hardcode the secret shares $\{[-w^*]_{\mathsf{idx}}\}_{\mathsf{idx}\in[n]}$ for the challenge ciphertext, instead of w^* and $\mathsf{coef}_1^*,\ldots,\mathsf{coef}_t^*$, as described in Algorithm 4. To preserve the input-output behavior of the program, Chal computes the shares $\{[-w^*]_{\mathsf{idx}}\}_{\mathsf{idx}\in[n]}$ exactly as they would have been computed in Algorithm 3.

Algorithm 4 $f^{\text{Game 11}}_{k_{\mathsf{Dec}}\{\mathsf{nonce}^*\},k_{\mathsf{Share}}\{\mathsf{nonce}^*\},\mathsf{SIG}.pk,[-w^*]_1,\ldots,[-w^*]_n}(\overrightarrow{pv},\mathsf{idx},sv,c)$

if $(\overrightarrow{pv}[\mathsf{idx}] = \mathsf{PRG}(sv))$ and $(\mathsf{SIG.Verify}(\mathsf{SIG}.pk,(\overrightarrow{pv},\mathsf{nonce}),\sigma))$ **then**
  if $\mathsf{nonce} = \mathsf{nonce}^*$ **then**
    $[-w]_{\mathsf{idx}} = [-w^*]_{\mathsf{idx}}$
  else
    $w \leftarrow \mathsf{PPRF}_{k_{\mathsf{Dec}}}(\mathsf{nonce})$
    for $j \in [1,\ldots,t]$ **do**
      $\mathsf{coef}_j = \mathsf{PPRF}_{k_{\mathsf{Share},j}}(\mathsf{nonce})$
    $[-w]_{\mathsf{idx}} = (-w + \sum_{j\in[1,\ldots,t]} \mathsf{coef}_j \mathsf{idx}^j) \bmod p$ {This gives the idxth Shamir share of $-w$}
  $[m]_{\mathsf{idx}} = ([-w]_{\mathsf{idx}} + e) \bmod p$ {This gives the idxth Shamir share of m}
  return $[m]_{\mathsf{idx}}$

Game 11 is indistinguishable from Game 10 by the security of indistinguishability obfuscation; the programs have identical input-output behavior.

Game 12 In this game, the challenger Chal picks $\{[-w^*]_{\mathsf{idx}}\}_{\mathsf{idx}\in[n]}$ at random.

Game 12 is indistinguishable from Game 11 by the security of indistinguishability obfuscation.

- On nonce $\mathsf{nonce} \neq \mathsf{nonce}^*$, the program behavior is unchanged.
- We guarantee that the program returns nothing in both games for nonce^* and $\overrightarrow{pv} \neq \{pv_i\}_{i\in\mathcal{R}^*}$ since no signatures exist on $(\overrightarrow{pv},\mathsf{nonce}^*)$ for $\overrightarrow{pv} \neq \{pv_i\}_{i\in\mathcal{R}^*}$.
- We guarantee that the program returns nothing in both games on nonce^* and indices idx corresponding to honest pv_i, $i \in \mathcal{R}^*\backslash\mathcal{C}$ since no values sv exist such that $pv_i = \mathsf{PRG}(sv)$ for honest parties i.
- Finally, consider nonce^* and corrupt pv_i ($i \in \mathcal{R}^* \cap \mathcal{C}$).
  Let $[\mathsf{idx}_1,\ldots,\mathsf{idx}_t]$ be the indices in the lexicographic ordering of $\{pv_i\}_{i\in\mathcal{R}^*}$ corresponding to corrupt pv_i, and let

$$A = \begin{pmatrix} \mathsf{idx}_1^t & \mathsf{idx}_1^{t-1} & \ldots & \mathsf{idx}_1^2 & \mathsf{idx}_1 \\ \mathsf{idx}_2^t & \mathsf{idx}_2^{t-1} & \ldots & \mathsf{idx}_2^2 & \mathsf{idx}_2 \\ \ldots & & & & \\ \mathsf{idx}_{t-1}^t & \mathsf{idx}_{t-1}^{t-1} & \ldots & \mathsf{idx}_{t-1}^2 & \mathsf{idx}_{t-1} \\ \mathsf{idx}_t^t & \mathsf{idx}_t^{t-1} & \ldots & \mathsf{idx}_t^2 & \mathsf{idx}_t \end{pmatrix}.$$

  Previously, the shares of $-w^*$ were computed as follows:

$$\begin{pmatrix} [-w^*]_{\mathsf{idx}_1} \\ [-w^*]_{\mathsf{idx}_2} \\ \ldots \\ [-w^*]_{\mathsf{idx}_{t-1}} \\ [-w^*]_{\mathsf{idx}_t} \end{pmatrix} = A \begin{pmatrix} \mathsf{coef}_t^* \\ \mathsf{coef}_{t-1}^* \\ \ldots \\ \mathsf{coef}_2^* \\ \mathsf{coef}_1^* \end{pmatrix} + \begin{pmatrix} -w^* \\ -w^* \\ \ldots \\ -w^* \\ -w^* \end{pmatrix}$$

  Choosing the coefficients as well as w^* at random and computing the shares of $-w^*$ as above is equivalent to choosing the shares of $-w^*$ as well as w^* at random and computing the coefficients as A^{-1} times the random shares of $-w^*$ plus w^*.

Game 13 In this game, the challenger switches to using a random message m^*. Game 13 is indistinguishable from Game 12 because the distributions do not change at all; e^* is still uniformly random, and the obfuscated program, which no longer contains any information about w^*, is unaffected.

The rest of the games are what we did before, but in reverse, with m_L instead of m_R.

Game	Justification	$\mathsf{SIG}.pk$	Honest pv_i	Obfuscated Program	c^*	m^*
4		real	real	real	real	m_R
5	Constrained Signatures	constrained to only verify on $(\overrightarrow{pk}, \mathsf{nonce}^*)$ when $\overrightarrow{pk} = \{pv_j\}_{j \in \mathcal{R}^*}$				
6	PRG		no matching secrets			
7	iO			semantic changes		
8	iO			puncture k_{Dec} and k_{Share} at nonce^*; hardcode correct values w^* and $\{\mathsf{coef}_j^*\}_{j \in [1,\ldots,t]}$		
9	PPRF			hardcode random $\{\mathsf{coef}_j^*\}_{j \in [1,\ldots,t]}$		
10	PPRF			hardcode random mask w^*	compute e^* using the random mask	
11	iO			hardcode shares of $-w^*$ instead of w^* and $\{\mathsf{coef}_j^*\}_{j \in [1,\ldots,t]}$		
12	iO			hardcode random values as shares of $-w^*$		
13	identical distributions					random

Fig. 15: Summary of Hybrids in Proof of Theorem 4

F.2 Proof that Obfuscation-Based Homomorphic Ad Hoc Threshold Encryption Share-and-Encrypt is Super-Partial Decryption Simulatable

SimPartDec is simply the Shamir secret sharing SimShares algorithm. An adversary who distinguishes such simulated shares from real shares can be used to break the super-static semantic security of the scheme.

G Additively Server-Aided Homomorphic Obfuscation-Based HATE

In this appendix, we describe the additively server-aided homomorphic obfuscation-based HATE scheme. The program each sender must obfuscate and include in their public key is described in Algorithm 5. The obfuscation-based HATE is described in Construction 5.

Algorithm 5 $f_{k_{\mathsf{Dec}},k_{\mathsf{Share}},k_{\mathsf{Enc}},\mathsf{SIG}.pk}(\overrightarrow{pk} = \{\mathsf{EG}.pk_j\}_{j\in\mathcal{R}}, \mathsf{idx}, c)$

The following values are hardcoded in the program:

- $\mathsf{params} = (\lambda, \mathsf{params}_{\mathsf{EG}}, t)$, where
  - λ is the security parameter,
  - $\mathsf{params}_{\mathsf{EG}} = (\mathcal{G}, p, g)$ consists of a p-order group $\mathcal{G}$ with generator g (where p is large prime and the range of the puncturable pseudorandom function PPRF is in $\mathbb{Z}_p$), and
  - t is the threshold
- A secret PPRF key k_{Dec} that is used to recover the message from the ciphertext c
- Secret PPRF keys $k_{\mathsf{Share}} = (k_{\mathsf{Share},1}, \ldots, k_{\mathsf{Share},t})$ that are used to produce randomness for sharing the message
- Secret PPRF keys $k_{\mathsf{Enc}} = (k_{\mathsf{Enc},1}, \ldots, k_{\mathsf{Enc},n})$ that are used to produce randomness for encrypting the shares
- A signature verification key $\mathsf{SIG}.pk$

The following values are expected as input:

- public encryption keys $\overrightarrow{pk} = \{\mathsf{EG}.pk_j\}_{j\in\mathcal{R}}$
- an index idx
- ciphertext $c = (\mathsf{nonce}, e, \sigma)$

if $\mathsf{SIG.Verify}(\mathsf{SIG}.pk_{\mathsf{Sndr}}, (\overrightarrow{pk}, \mathsf{nonce}), \sigma)$ **then**
  $w \leftarrow \mathsf{PPRF}_{k_{\mathsf{Dec}}}(\mathsf{nonce})$
  $m = (e - w) \bmod p$
  for $j \in [1, \ldots, t]$ **do**
    $\mathsf{coef}_j = \mathsf{PPRF}_{k_{\mathsf{Share},j}}(\mathsf{nonce})$
  $[m]_{\mathsf{idx}} = g^{(\sum_{j\in[1,\ldots,t]} \mathsf{coef}_j \mathsf{idx}^j)+m}$ {This gives the idxth exponential Shamir share of m}
  $r_{\mathsf{idx}} = \mathsf{PPRF}_{k_{\mathsf{Enc},\mathsf{idx}}}(\mathsf{nonce})$
  $c = \mathsf{EG.Enc}(\mathsf{params}_{\mathsf{EG}}, \overrightarrow{pk}_{\mathsf{idx}}, [m]_{\mathsf{idx}}; r_{\mathsf{idx}})$
  {This returns an ElGamal encryption of the idxth exponential Shamir share of m. Encryption uses randomness r_{idx}.}
  return c

Let the public parameters $\mathsf{params} = (\lambda, \mathsf{params}_{\mathsf{EG}} = (\mathcal{G}, p, g), t)$ consist of the security parameter λ, a large p-order group $\mathcal{G}$ with generator g (such that p is prime and the range of the puncturable pseudorandom function PPRF is in $\mathbb{Z}_p$), and the threshold t. For simplicity we omit params as input from the algorithms below.

$\mathsf{KeyGen}(t)$**:**

This is exactly as in Construction 2, except that the sender generates n additional PPRF keys $k_{\mathsf{Enc},1}, \ldots, k_{\mathsf{Enc},n}$, and instead of obfuscating f from Algorithm 1 to get $\mathsf{ObfFunc}$, the sender obfuscates f from Algorithm 5

$\mathsf{Enc}(sk_{\mathsf{Sndr}} = (\mathsf{SIG}.sk, k_{\mathsf{Dec}}), \overrightarrow{pk} = \{\mathsf{EG}.pk_j\}_{j \in \mathcal{R}, |\mathcal{R}| \geq t}, m)$**:**

This is exactly as in Construction 2

$\mathsf{PartDec}(\mathsf{ObfFunc}_{\mathsf{Sndr}}, \{\mathsf{EG}.pk_j\}_{j \in \mathcal{R}}, \mathsf{EG}.sk_i, c)$**:**

if the ciphertext c is an output of a homomorphic evaluation **then**
$c' = c$
else
Let idx be the index of the public key corresponding to the secret key $\mathsf{EG}.sk$ in a lexicographic ordering of $\overrightarrow{pk} = \{\mathsf{EG}.pk_j\}_{j \in \mathcal{R}}$
$c' = \mathsf{ObfFunc}_{\mathsf{Sndr}}(\{\mathsf{EG}.pk_j\}_{j \in \mathcal{R}}, \mathsf{idx}, c)$
$[m]_{\mathsf{idx}} \leftarrow \mathsf{EG.Dec}(\mathsf{params}_{\mathsf{EG}}, \mathsf{EG}.sk_i, c')$
$d_i = (\mathsf{idx}, [m]_{\mathsf{idx}})$
return d_i

$\mathsf{FinalDec}(\{d_i\}_{i \in \mathcal{R}' \subset \mathcal{R}})$**:**

Perform exponential Shamir reconstruction $\mathsf{EShamir.Reconstruct}(\{d_i\}_{i \in \mathcal{R}'})$ as described in Figure 10 to recover m

$\mathsf{Eval}(\{pk_{\mathsf{Sndr}}\}_{\mathsf{Sndr} \in \mathcal{S}}, \overrightarrow{pk} = \{pk_i\}_{i \in \mathcal{R}}, [c_1, \ldots, c_l], +)$**:**

{Note that this algorithm receives the public keys for all senders Sndr (and thus their obfuscated programs). Without loss of generality, let ciphertext c_q be from sender P_q (and therefore requiring the use of $\mathsf{ObfFunc}_q$).}
for ciphertext indices $q \in [1, \ldots, l]$ **do**
for receivers $i \in \mathcal{R}$ **do**
Let idx be the index of $\mathsf{EG}.pk_i$ in a lexicographic ordering of $\overrightarrow{pk}$
$c_{i,q} \leftarrow \mathsf{ObfFunc}_q(\overrightarrow{pk}, \mathsf{idx}, c_q))$
$c_i^* = \mathsf{EG.Eval}(\mathsf{params}_{\mathsf{EG}}, \mathsf{EG}.pk_i, [c_{i,1}, \ldots, c_{i,l}], +)$
return $c^* = \{c_i^*\}_{i \in \mathcal{R}}$

Construction 5: Obfuscation-Based HATE

Theorem 13 (Restated from Theorem 5). *The modified obfuscation-based ATE (Construction 5) is (n, t)-super-statically secure (Definition 12) for any polynomial n, t, as long as* iO *is a secure indistinguishability obfuscator,* PPRF *is a secure puncturable* PRF *with range $\mathbb{Z}_p$,* SIG *is a constrained signature scheme, and* EG *is a secure public-key encryption scheme. Moreover, it is additively server-aided homomorphic for a polynomial-size message space.*

Proof. In order to prove Theorem 4, we must show the modified obfuscation-based Homomorphic Ad Hoc Threshold Encryption construction is super-statically semantically secure and super-statically partial decryption simulatable. We prove super-static semantic security below, by showing a sequence of indistinguishable games starting at $\mathsf{EXP}(\mathcal{A}, \lambda, n, t, R)$ and ending at a message-independent game. The proof of partial decryption simulatability is the same as for Theorem 4. The server-aided homomorphism follows from the homomorphism of ElGamal encryption and exponential Shamir secret sharing.

Game 1 This is the same as Game 4 in the proof of Theorem 4. That is, this is $\mathsf{EXP}(\mathcal{A}, \lambda, n, t, R)$ as described in Definition 13 and Figure 13.

Game 2 This is the same as Game 5 in the proof of Theorem 4. That is, in this game, when creating the sender's public key (specifically, the sender's signature verification key $\mathsf{SIG}.pk$), the challenger Chal generates a constrained verification key instead of a regular one. The constraint is designed to ensure that the challenge nonce nonce^* on which the PPRFs are called can only ever be associated with the challenge recipient set.

Game 2 is indistinguishable from Game 1 by the security of constrained signatures.

Game 3 In this game, the challenger Chal changes the way encryption takes place in the obfuscated program, as shown in Algorithm 6. In particular, instead of computing $c = \mathsf{EG.Enc}(\mathsf{EG}.pk_i, [m]_i; r_i)$, the program now computes $(c_1, c_2) = \mathsf{EG.Enc}(\mathsf{EG}.pk_i, [-w]_i; r_i)$, followed by $c = (c_1, c_2 \times g^e)$. Note that this gives the exact same ciphertext as encrypting $[m]_i$ directly; in fact, the only reason we did not use these instructions explicitly to begin with is clarity.

Algorithm 6 $f^{\mathsf{Game}\ 3}_{k_{\mathsf{Dec}},k_{\mathsf{Share}},k_{\mathsf{Enc}},\mathsf{SIG}.pk}(\overrightarrow{pk}, \mathsf{idx}, c)$

if $\mathsf{SIG.Verify}(\mathsf{SIG}.pk_{\mathsf{Sndr}}, (\overrightarrow{pk}, \mathsf{nonce}), \sigma)$ **then**
  $w \leftarrow \mathsf{PPRF}_{k_{\mathsf{Dec}}}(\mathsf{nonce})$
  for $j \in [1, \ldots, t]$ **do**
    $\mathsf{coef}_j = \mathsf{PPRF}_{k_{\mathsf{Share},j}}(\mathsf{nonce})$
  $[-w]_{\mathsf{idx}} = g^{(\sum_{j\in[1,\ldots,t]} \mathsf{coef}_j \mathsf{idx}^j) - w}$ {This gives the idxth Shamir share of $-w$}
  $r_{\mathsf{idx}} = \mathsf{PPRF}_{k_{\mathsf{Enc},\mathsf{idx}}}(\mathsf{nonce})$
  $(c_{\mathsf{idx},1}, c_{\mathsf{idx},2}) = \mathsf{EG.Enc}(\mathsf{params}_{\mathsf{EG}}, \overrightarrow{pk}_{\mathsf{idx}}, [-w]_{\mathsf{idx}}; r_{\mathsf{idx}}) = (g^{r_{\mathsf{idx}}}, \overrightarrow{pk}^{r_{\mathsf{idx}}}_{\mathsf{idx}}[-w]_{\mathsf{idx}})$
  {This returns an ElGamal encryption of the ith exponential Shamir share of $-w$. Encryption uses randomness r_i.}
  $c = (c_{i,1}, c_{i,2}g^e)$ {This returns an ElGamal encryption of the ith exponential Shamir share of $m = (e - w) \bmod p$.}
  return c

Game 3 is indistinguishable from Game 2 by the security of indistinguishability obfuscation; the programs have identical input-output behavior.

Game 4 This is the same as Game 8 in the proof of Theorem 4, except that in addition to puncturing the PPRF keys k_{Dec} and $k_{\mathsf{Share},j}$ for all $j \in [1, \ldots, t]$, the challenger Chal also punctures $k_{\mathsf{Enc},i}$ for all $i \in \mathcal{R}^*$. To preserve the input-output behavior of the program, Chal computes $r^*_i = \mathsf{PPRF}_{k_{\mathsf{Enc},i}}(\mathsf{nonce}^*)$, and modifies the program to set $r_i = r^*_i$ when $\mathsf{nonce} = \mathsf{nonce}^*$, as shown in Algorithm 7.

Algorithm 7 $f^{\mathsf{Game}\ 4}_{k_{\mathsf{Dec}}\{\mathsf{nonce}^*\},k_{\mathsf{Share}}\{\mathsf{nonce}^*\},k_{\mathsf{Enc}}\{\mathsf{nonce}^*\},\mathsf{SIG}.pk,\mathsf{nonce}^*,w^*,\mathsf{coef}^*_1,\ldots,\mathsf{coef}^*_t,r^*_1,\ldots,r^*_n}(\overrightarrow{pk}, i, c)$

if $\mathsf{SIG.Verify}(\mathsf{SIG}.pk_{\mathsf{Sndr}}, (\overrightarrow{pk}, \mathsf{nonce}), \sigma)$ **then**
  if $\mathsf{nonce} = \mathsf{nonce}^*$ **then**
    $w = w^*$
    for $j \in [1, \ldots, t]$ **do**
      $\mathsf{coef}_j = \mathsf{coef}^*_j$
    $r_{\mathsf{idx}} = r^*_{\mathsf{idx}}$
  else
    $w \leftarrow \mathsf{PPRF}_{k_{\mathsf{Dec}}\{\mathsf{nonce}^*\}}(\mathsf{nonce})$
    for $j \in [1, \ldots, t]$ **do**
      $\mathsf{coef}_j = \mathsf{PPRF}_{k_{\mathsf{Share},j}\{\mathsf{nonce}^*\}}(\mathsf{nonce})$
    $r_{\mathsf{idx}} = \mathsf{PPRF}_{k_{\mathsf{Enc},\mathsf{idx}}\{\mathsf{nonce}^*\}}(\mathsf{nonce})$
  $[-w]_{\mathsf{idx}} = g^{(\sum_{j\in[1,\ldots,t]} \mathsf{coef}_j \mathsf{idx}^j) - w}$ {This gives the idxth Shamir share of $-w$}
  $(c_{\mathsf{idx},1}, c_{\mathsf{idx},2}) = \mathsf{EG.Enc}(\mathsf{params}_{\mathsf{EG}}, \overrightarrow{pk}_{\mathsf{idx}}, [-w]_{\mathsf{idx}}; r_{\mathsf{idx}}) = (g^{r_{\mathsf{idx}}}, \overrightarrow{pk}^{r_{\mathsf{idx}}}_{\mathsf{idx}}[-w]_{\mathsf{idx}})$
  {This returns an ElGamal encryption of the idxth exponential Shamir share of $-w$. Encryption uses randomness r_{idx}.}
  $c = (c_{\mathsf{idx},1}, c_{\mathsf{idx},2}g^e)$ {This returns an ElGamal encryption of the idxth exponential Shamir share of $m = (e - w) \bmod p$.}
  return c

Game 4 is indistinguishable from Game 3 by the security of indistinguishability obfuscation; the programs have identical input-output behavior.

Game 5 This game is the same as Game 9 in the proof of Theorem 4. That is, in this game, the challenger Chal chooses coef_j^* truly at random for $j \in [1, \ldots, t]$. (Note that this game really requires t hybrids that we are skipping in the interest of brevity.)

Game 5 is indistinguishable from Game 4 by the security of puncturable PRFs.

Game 6 This game is the same as Game 10 in the proof of Theorem 4. That is, in this game, the challenger Chal chooses w^* truly at random.

Game 6 is indistinguishable from Game 5 by the security of puncturable PRFs.

Game 7.idx **for** $\mathsf{idx} \in [1, \ldots, n]$

In this game, the challenger Chal modifies the obfuscated program to choose the encryption randomness uniformly at random for the idxth party. That is, Chal sets r^*_{idx} at random (instead of setting it to $k_{\mathsf{Enc},\mathsf{idx}}(\mathsf{nonce}^*)$).

Game 7.1 is indistinguishable from Game 6, and Game 7.idx is indistinguishable from Game 7.$(\mathsf{idx}-1)$ for $\mathsf{idx} \in [2, \ldots, n-t]$, by the security of puncturable PRFs.

We let Game 7 denote Game 7.n.

Game 8 In this game, the challenger Chal modifies the obfuscated program to hardcode the ciphertexts $(c^*_{\mathsf{idx},1}, c^*_{\mathsf{idx},2})$ for $\mathsf{idx} \in [1, \ldots, n]$ instead of w^*, $\mathsf{coef}_1^*, \ldots, \mathsf{coef}_t^*$ and $r_1^*, \ldots, r_n^*$, as described in Algorithm 8. To preserve the input-output behavior of the program, Chal computes $(c^*_{\mathsf{idx},1}, c^*_{\mathsf{idx},2})$ exactly as they would have been computed in Algorithm 7.

Algorithm 8 $f^{\mathsf{Game}\ 8}_{k_{\mathsf{Dec}}\{\mathsf{nonce}^*\},k_{\mathsf{Share}}\{\mathsf{nonce}^*\},k_{\mathsf{Enc}}\{\mathsf{nonce}^*\},\mathsf{SIG}.pk,\mathsf{nonce}^*,(c^*_{1,1},c^*_{1,2}),\ldots,(c^*_{n,1},c^*_{n,2})}(\overrightarrow{pk}, \mathsf{idx}, c)$

if $\mathsf{SIG.Verify}(\mathsf{SIG}.pk_{\mathsf{Sndr}}, (\{\mathsf{EG}.pk_j\}_{j\in\mathcal{R}}, \mathsf{nonce}), \sigma)$ **then**
  if $\mathsf{nonce} = \mathsf{nonce}^*$ **then**
    $(c_{\mathsf{idx},1}, c_{\mathsf{idx},2}) = (c^*_{\mathsf{idx},1}, c^*_{\mathsf{idx},2})$
  else
    $w \leftarrow \mathsf{PPRF}_{k_{\mathsf{Dec}}\{\mathsf{nonce}^*\}}(\mathsf{nonce})$
    for $j \in [1, \ldots, t]$ **do**
      $\mathsf{coef}_j = \mathsf{PPRF}_{k_{\mathsf{Share},j}\{\mathsf{nonce}^*\}}(\mathsf{nonce})$
    $r_{\mathsf{idx}} = \mathsf{PPRF}_{k_{\mathsf{Enc},\mathsf{idx}}\{\mathsf{nonce}^*\}}(\mathsf{nonce})$
    $[-w]_{\mathsf{idx}} = g^{(\sum_{j\in[1,\ldots,t]} \mathsf{coef}_j \mathsf{idx}^j) - w}$ {This gives the idxth Shamir share of $-w$}
    $(c_{\mathsf{idx},1}, c_{\mathsf{idx},2}) = \mathsf{EG.Enc}(\mathsf{params}_{\mathsf{EG}}, \overrightarrow{pk}_{\mathsf{idx}}, [-w]_{\mathsf{idx}}; r_{\mathsf{idx}}) = (g^{r_{\mathsf{idx}}}, \overrightarrow{pk}^{r_{\mathsf{idx}}}_{\mathsf{idx}}[-w]_{\mathsf{idx}})$
    {This returns an ElGamal encryption of the idxth exponential Shamir share of $-w$. Encryption uses randomness r_{idx}.}
  $c = (c_{\mathsf{idx},1}, c_{\mathsf{idx},2} g^e)$ {This returns an ElGamal encryption of the idxth exponential Shamir share of $m = (e - w) \bmod p$.}
  return c

Game 8 is indistinguishable from Game 7 by the security of indistinguishability obfuscation; the programs have identical input-output behavior.

Game 9.i for $i \in [1, \ldots, n-t]$
In this game, the challenger Chal modifies the obfuscated program to hard-code encryptions of zero for honest parties. Let idx_i be the index of the ith honest public key in a lexicographic ordering of $\overrightarrow{pk}^*$. Chal sets $(c^*_{1,\mathsf{idx}_i}, c^*_{2,\mathsf{idx}_i}) = \mathsf{EG.Enc}(\overrightarrow{pk}^*_{\mathsf{idx}_i}, 0; r^*_{\mathsf{idx}})$.
Game 9.1 is indistinguishable from Game 8, and Game 9.i is indistinguishable from Game 9.$(i-1)$ for $i \in [2, \ldots, n-t]$, by the semantic security of EG.
We let Game 9 denote Game 9.$(n-t)$.

Game 10 This game is the same as Game 12 in the proof of Theorem 4. That is, in this game, the challenger Chal encrypts random values for corrupt parties. Let idx_i be the index of the ith corrupt public key in a lexicographic ordering of $\overrightarrow{pk}^*$. Instead of $\{[-w^*]_{\mathsf{idx}_i}\}$ to obtain $\{(c^*_{\mathsf{idx}_i,1}, c^*_{\mathsf{idx}_i,2})\}$, Chal encrypts uniformly random values.
Game 10 is indistinguishable from Game 9 by the security of indistinguishability obfuscation.

Game 11 In this game, the challenger switches to using a random message m^*.
Game 11 is indistinguishable from Game 10 because the distributions do not change at all; e^* is still uniformly random, and the obfuscated program, which no longer contains any information about w^*, is unaffected.

The rest of the games are what we did before, but in reverse, with m_L instead of m_R.

Game	Justification	$\mathsf{SIG}.pk$	Obfuscated Program	c^*	m^*
1		real	real	real	m_R
2	Constrained Signatures	constrained to only verify on nonce^* when $\{pv_j\}_{j\in\mathcal{R}} = \{pv_j\}_{j\in\mathcal{R}^*}$			
3	iO		semantic changes		
4	iO		puncture k_{Dec}, k_{Share} and k_{Enc} at nonce^*; hardcode correct values w^*, $\{\mathsf{coef}_j^*\}_{j\in[1,\ldots,t]}$ and $\{r^*_{\mathsf{idx}}\}_{\mathsf{idx}\in[1,\ldots,n]}$		
5	PPRF		hardcode random $\{\mathsf{coef}_j^*\}_{j\in[1,\ldots,t]}$		
6	PPRF		hardcode random mask w^*	compute e^* using the random mask	
7	PPRF		hardcode random $\{r^*_{\mathsf{idx}}\}_{\mathsf{idx}\in[1,\ldots,n]}$		
8	iO		hardcode $\{(c^*_{1,\mathsf{idx}}, c^*_{2,\mathsf{idx}})\}_{\mathsf{idx}\in[1,\ldots,n]}$ instead of w^*, $\{\mathsf{coef}_j^*\}_{j\in[1,\ldots,t]}$ and $\{r^*_{\mathsf{idx}}\}_{\mathsf{idx}\in[1,\ldots,n]}$		
9	semantic security		hardcode encryptions of 0 instead of encryptions of shares of $-w^*$ for honest parties		
10	iO		hardcode encryptions of random values instead of encryptions of shares of $-w^*$ for corrupt parties		
11	identical distributions				random

Fig. 16: Summary of Hybrids in Proof of Theorem 5