Encoding Functions with Constant Online Rate
or
How to Compress Garbled Circuit Keys*

Benny Applebaum† Yuval Ishai‡ Eyal Kushilevitz§ Brent Waters¶

August 27, 2014

Abstract

Randomized encodings of functions can be used to replace a "complex" function $f(x)$ by a "simpler" randomized mapping $\hat{f}(x;r)$ whose output distribution on an input x encodes the value of $f(x)$ and hides any other information about x. One desirable feature of randomized encodings is low *online complexity*. That is, the goal is to obtain a randomized encoding $\hat{f}$ of f in which most of the output can be precomputed and published before seeing the input x. When the input x is available, it remains to publish only a short string $\hat{x}$, where the online complexity of computing $\hat{x}$ is independent of (and is typically much smaller than) the complexity of computing f. Yao's garbled circuit construction gives rise to such randomized encodings in which the online part $\hat{x}$ consists of n encryption keys of length κ each, where $n = |x|$ and κ is a security parameter. Thus, the *online rate* $|\hat{x}|/|x|$ of this encoding is proportional to the security parameter κ.

In this paper, we show that the online rate can be dramatically improved. Specifically, we show how to encode any polynomial-time computable function $f : \{0,1\}^n \to \{0,1\}^{m(n)}$ with online rate of $1+o(1)$ and with nearly linear online computation. More concretely, the online part $\hat{x}$ consists of an n-bit string and a single encryption key. These constructions can be based on the decisional Diffie-Hellman assumption (DDH), the Learning with Errors assumption (LWE), or the RSA assumption. We also present a variant of this result which applies to *arithmetic formulas*, where the encoding only makes use of arithmetic operations, as well as several negative results which complement our positive results.

Our positive results can lead to efficiency improvements in most contexts where randomized encodings of functions are used. We demonstrate this by presenting several concrete applications. These include protocols for secure multiparty computation and for non-interactive verifiable computation in the preprocessing model which achieve, for the first time, an optimal online communication complexity, as well as non-interactive zero-knowledge proofs which simultaneously minimize the online communication and the prover's online computation.

*Preliminary versions of this paper appeared in [AIKW12, AIKW13].

†School of Electrical Engineering, Tel-Aviv University, bennyp@post.tau.ac.il.

‡Department of Computer Science, Technion, yuvali@cs.technion.ac.il.

§Department of Computer Science, Technion, eyalk@cs.technion.ac.il.

¶Department of Computer Science, University of Texas, bwaters@cs.utexas.edu.

Contents

1 Introduction

Suppose that we want to perform some cryptographic task which involves computation and communication on n-bit data. In many scenarios, it is beneficial to minimize the online complexity (i.e., the resources spent after seeing the data) and shift the expensive computation and communication to an offline phase. This setting has been extensively studied in many contexts including signatures [EGM96, ST01], verifiable computation (delegation) [GGP10, AIK10, CKV10], and secure computation [Bea95, IPS08, BDOZ11, DPSZ12, IKM$^+$13]. The goal of the present paper is to further explore the question of minimizing the online complexity of cryptography.

Let us first consider the following concrete example from [AIK11]. Imagine a scenario of sending a weak device U to the field in order to perform some expensive computation f on sensitive data x. The computation is too complex for U to quickly perform it on its own and, since the input x is sensitive, U cannot just send the entire input out. Ideally, we would like to have a *non-interactive* solution of the following form: In an offline phase, before sent to the field, U picks a short random secret key sk and publishes a (potentially long) related public key pk. Once it observes the input x, the device U applies some cheap computation to sk and x and sends out the result $\hat{x}$, a short "encrypted" version of x. The rest of the world should be able, at this point, to recover $f(x)$ and nothing else.

Abstracting the above, the computation of U can be described as a randomized function $\hat{f} : (x; \mathsf{sk}) \mapsto (\mathsf{pk}, \hat{x})$ that *encodes* the value $f(x)$ in the sense that $(\mathsf{pk}, \hat{x})$ reveals $f(x)$ but nothing else. Using the terminology of [AIK04], the function $\hat{f}$ is referred to as a *randomized encoding* (RE) of f. The general motivation for using REs is the hope to make $\hat{f}$ in some sense "simpler" than f, where different applications dictate different notions of simplicity. The earliest uses of REs in cryptography were in the area of secure computation [Yao86, Kil88, FKN94, IK00]. Along the years, REs have found a diverse range of other applications to problems such as computing on encrypted data [SYY99, CCKM00], parallel cryptography [AIK04, AIK06], verifiable computation [GGP10, AIK10], software protection [GKR08b, GIS$^+$10, BHR12a], functional encryption [SS10, GVW12], key-dependent message security [BHHI10, App11], and others. We refer the reader to [BHR12b] for a finer-grained treatment of REs under the term "garbling schemes".

In the online/offline setting considered here, we would like to minimize the online computation and communication resources required for computing and distributing $\hat{x}$. That is, we would like the online time complexity of computing $\hat{x}$ to be much smaller than the time required for computing f, and the length of $\hat{x}$ to be not much bigger than that of x.

The best known general constructions of online-efficient REs are based on Yao's garbled circuit technique [Yao86]. In this case, the output of $f(x)$ is encoded by an offline part pk which consists of a big "garbled circuit" and an online part $\hat{x}$ which consists of n keys $K_1, \ldots, K_n$ of size κ each, where n is the bit length of x and κ is a security parameter. (Under a standard asymptotic security convention in which n serves both as an input length parameter and a security parameter, κ can be thought of as n^ε, for some small constant $\varepsilon > 0$.) Each key K_i is selected from a pair of keys $(K_{i,0}, K_{i,1})$ according to the i-th input bit x_i. Hence, the online computation and communication complexity are both $O(n\kappa)$. An appealing feature is that the online computation complexity is nearly linear in the input length, independently of the complexity of f. However, an undesirable feature is that the *online rate* of the construction — i.e., the ratio between the bit length of $\hat{x}$ and the bit length of x — grows linearly with the security parameter κ. Hence, we ask:

Is it possible to obtain a *constant* online rate or even rate of $1 + o(1)$ (e.g., $|\hat{x}| =$

$n + \mathrm{poly}(\kappa)$) while keeping the online computation independent of the complexity of f?

1.1 Our Contribution

We answer the above question in the affirmative by constructing, under a variety of standard intractability assumptions, an online-efficient RE with rate $1 + o(1)$ for every polynomial-time computable function.

Theorem 1.1. *(Informal) Under the Decisional Diffie-Hellman Assumption (DDH), the RSA Assumption, or the Learning-with-Errors Assumption (LWE), every polynomial-time computable function $f : \{0,1\}^n \to \{0,1\}^{m(n)}$ admits an RE with online rate $1 + o(1)$ and with $O(n^{1+\varepsilon})$ online computation, for any $\varepsilon > 0$.*

In more concrete terms, our constructions efficiently compile any boolean circuit C into a corresponding RE with succinct and efficiently computable online part. These constructions can be viewed as analogues of the garbled circuit construction in which the n keys determined by x are compressed into a shorter string $\hat{x}$ whose length is very close to that of x. This comes at the cost of a slight increase in the online computation complexity, which still remains nearly linear in n. An additional (related) difference is that in contrast to the standard garbled circuit construction, where each bit of $\hat{x}$ depends only on a single bit of x, in our constructions there are bits of $\hat{x}$ which depend on many bits of x. We prove that this is inherent for REs with constant or even logarithmic online rate. In particular, it is impossible to obtain a direct generalization of the garbled circuit construction in which each input bit x_i selects between a pair of keys $(K_{i,0}, K_{i,1})$ which have constant size.

The DDH and LWE based constructions are *affine* in the sense that after the private randomness is fixed in the offline phase, the remaining computation can be described as an affine function of the inputs x (over some ring R, e.g., $\mathsf{R} = \mathbb{Z}_p$ where p is the size of a DDH group). This captures a strong form of algebraic simplicity which is useful for some of the motivating applications (e.g., secure computation).

Motivated by the concrete efficiency of encoding *arithmetic* computations, we also present an LWE-based arithmetic variant of the above result that applies to arithmetic *formulas* (i.e., circuits of fan-out 1) over large finite fields, where the encoding is restricted to applying arithmetic operations to the inputs. Specifically, we obtain an affine randomized encoding (ARE, for short) with optimal online rate (i.e., $1 + o(1)$) for arithmetic mod-p formulas, assuming that elements of $\mathbb{Z}_p$ can be viewed as elements of $\mathbb{Z}_q$ for some $q \gg p$. If we insist on working in the more restricted model of [AIK11], where the encoding should be affine over the integers, then we get a constant-rate encoding.

It should be mentioned that the online *computational* overhead of our constructions is still polynomial in the security parameter. Whether this overhead can be improved remains an interesting open question.

Lower bounds. We further explore the complexity of REs in the online/offline setting by proving several lower bounds on the online and offline rate of REs which complement our positive results. Among other results, we study the minimal achievable online rate. The online rate is clearly lower-bounded by 1 for some functions with long outputs (this is the case, for instance, for the identity function). This leaves open the possibility of achieving a strictly better rate for boolean functions.

We show that even in the case of boolean functions, the online rate of *affine* REs (satisfying the algebraic simplicity condition discussed above) cannot generally be smaller than 1. Thus, achieving rate $1 + o(1)$ is essentially optimal for affine REs. While we cannot unconditionally prove a similar result for non-affine REs with, say, quadratic online computation, such a negative result follows from the conjecture that for any $c > c'$, an input for a time-(n^c) computation cannot generally be "compressed" by a time-$(n^{c'})$ algorithm into a shorter string which contains sufficient information to recover the output. See [HN06, DI06] for related conjectures. Finally, an online rate of 1 can be easily seen to be optimal if one requires a separate encoding of the function description (e.g., a circuit) and the input. This is the case for most definitions of garbling schemes from the literature, such as the one from [BHR12b].

Adaptive security. Informally, an offline/online RE is adaptively secure if $\hat{f}(x; r) = (\mathsf{pk}, \hat{x})$ remains private even if the online input x is adaptively chosen based on the offline part of the encoding, pk. Similarly to all other known implementations of garbled circuits with short keys, our constructions cannot be proved to satisfy this stronger notion of security unless analyzed in the (programmable) random oracle model. We prove that this is inherent to some extent: in any RE whose adaptive security holds in the plain model, the length of the online part $\hat{x}$ should grow with the output length of f. (This negative result is similar in spirit to negative results for non-committing encryption [Nie02] or functional encryption [BSW11].) In contrast, our constructions in the non-adaptive setting (or the adaptive setting with random oracles) have online rate of $1 + o(1)$, independently of the output length of f. Adaptive security of garbled circuits has recently been considered in the work of Bellare et al. [BHR12a]. The above negative result partially settles a question left open by [BHR12a].

On concrete efficiency. In concrete terms, our offline/online REs reduce the online communication of Yao's garbled circuit construction by a factor of $\kappa \approx 100$ at the expense of introducing "public-key" computations. This is not always a good tradeoff in practice. For instance, communicating 100 bits is typically less expensive than a single modular exponentiation. Luckily, our REs are also very cheap in online computation. For instance, the online encoding in the DDH-based construction involves at most one mod-p *addition* per input bit, where p is the order of the DDH group. Since a mod-p addition is typically much cheaper than the amortized cost of communicating a bit (let alone 100 bits), we improve the overall concrete online complexity by roughly a factor of 100. This is contrasted with most applications of public-key cryptography towards improving communication complexity, where the additional computational cost outweighs the savings in communication (cf. [SC07]). While our REs do increase the complexity of the offline encoding and online decoding, the additional overhead is insignificant when the circuit complexity of f is much bigger than its input size. Thus, our offline/online REs seem to have a true practical potential in secure computation or delegation scenarios in which a weak client (who performs the offline and online encoding) interacts with a powerful server (who performs the online decoding).

Table 1 summarizes the concrete efficiency of our basic constructions. (We mention that one can obtain a smooth tradeoff between the offline part and the online part, as shown in Section 3.5 and Theorem 4.2.)

Assumption	Online Comm.	Online Comp.	Offline Comm.	Offline Comp.	Decoding
DDH over $\mathbb{G}$ of order p	n bits + 1 element in $\mathbb{Z}_p$	n add's	n^2 $\mathbb{G}$ elements	n^2 exp's	n^2 group mul's + single exp.
LWE dimension k over $\mathbb{Z}_q$	n bits + k elements in $\mathbb{Z}_q$	nk add's	$n^2k + nk^2$ $\mathbb{Z}_q$ elements	n^2k^2 mul's	n^2k add's + k^2 mul's
RSA over modulus N	n bits + 1 element in $\mathbb{Z}_N$	n mul's + single exp.	$2n$ $\mathbb{Z}_N$ elements	n exp's	n mul's +single exp.

Table 1: The concrete complexity of our encodings. This table summarizes the concrete online complexity (in terms of communication and computation) of encoding a Boolean function over n bits. The last three columns describe the *overhead* (in terms of offline communication, offline computation, and decoding complexity) that our construction introduce on top of an existing encoding (e.g., based on Yao's garbled circuit). We abbreviate *modular additions*, *modular multiplications*, and *modular exponentiations* by *add's*, *mul's* and *exp's*. The RSA parameters are based on the optimization described in Remark 3.6. We mention that one can obtain a smooth tradeoff between the offline part and the online part. In particular, it is possible to avoid the quadratic blowup in the offline complexity of the DDH/LWE constructions at the expense of increasing the online communication (see Section 3.5 and Theorem 4.2).

1.2 Applications

Our positive results can lead to efficiency improvements in most contexts in which randomized encodings of functions are used. We focus on three representative applications.

Secure Multiparty Computation (MPC). In the online/offline model (or preprocessing model) for MPC, there are t players who wish to securely compute some fixed public function f. In the offline phase, before the inputs "arrive", the parties are allowed to invoke some (relatively expensive) protocol; later, in the online phase, the parties get their inputs and apply an online (hopefully cheap) protocol. The close connection of REs to MPC [IK00] allows to translate our results into highly efficient MPC protocols in the offline/online setting. In Section 7.1, we further extend and optimize these reductions (exploiting the affinity property and the information-theoretic techniques from [BDOZ11]). This leads to general MPC protocols in which the online phase only requires each party to broadcast a message of the same length as its input along with a message of size $\text{poly}(\kappa)$, where κ is a security parameter. Again, this is information-theoretic optimal, and it beats, in terms of online communication complexity, all previously known results even in the simplest case of two semi-honest parties. We note, however, that our protocols do not offer provable security against malicious parties which adaptively choose their inputs based on the information they receive in the offline phase, except in the random oracle model or under nonstandard assumptions. See Section 7.1 for further discussion.

It is instructive to compare the efficiency of our RE-based protocols to protocols which are based on fully homomorphic encryption (FHE). The following discussion is restricted to the preprocessing model, which does not seem to significantly improve the complexity of FHE-based protocols. In FHE based protocols (as well as all other general MPC protocols from the literature) the communication complexity grows at least linearly with the total input and output length $n+m$. In contrast, the online communication complexity of our protocol does *not* depend on the output length. This

is particularly useful when securely computing functionalities that have a short online secret input (say, shares of a signature key) and a long output (say, signatures on many predetermined messages using the shared signature key). Furthermore, our protocols can be made completely non-interactive in certain scenarios, e.g., when part of the secret input is known offline and the online part is known in its entirety to one of the parties. This is impossible to get using FHE.[1] On the other hand, our protocols are incomparable to FHE-based protocols in terms of their online computational complexity. In the case of computing a complex function f which takes inputs from Alice and Bob and delivers an output to Alice, our approach yields two-message protocols in which Bob's online computation is very efficient (nearly linear in its input), whereas FHE provides similar protocols in which Alice's computation is very efficient (quasilinear in the input and output). From a concrete efficiency point of view, the online phase of our protocols is much "lighter" (e.g., Bob only needs to add a subset of $\mathbb{Z}_p$ elements corresponding to its input) and they can also be based on a wider variety of assumptions.

Finally, in the non-interactive model for secure computation of Feige, Kilian, and Naor [FKN94], our results give a computationally secure protocol for any polynomial-time computable function f with the following efficiency feature: The length of each party's message is equal to its input length plus a security parameter, and in addition there is a message (of polynomial length) that depends only on the common randomness. This additional message can be sent by one of the parties before the inputs are known. Note that even without any security requirements, for most nontrivial functions f the message length of each party (in a perfectly correct protocol) is lower bounded by the input length. Thus, the online communication complexity of the above protocol is nearly optimal.

Verifiable Computation. In an online/offline protocol for *verifiable computation* (VC), a computationally weak client with an input x delegates a complex computation f to an untrusted server in a two phase manner. In the offline phase the client sends to the server a possibly long and computationally expensive message pk, and at the online phase (when the input x arrives) the client sends a message $\hat{x}$ to the server, and receives back the result of the computation y together with a certificate for correctness. This setting was studied in several works (e.g., [Mic94, GKR08a, KR09, GGP10, CKV10, AIK10]). Specifically, in [GGP10] Yao's garbled circuit technique was used to achieve efficient VC in the online/offline model. (The security of the construction follows from standard assumptions only when the input x is picked by the client independently of pk [BHR12a].) This connection was generalized and optimized in [AIK10]. By plugging our encodings in these protocols, we get *communication optimal* VC protocols, where the bit length of the up-stream (online) message from the client to the server is $n+\kappa$ and the bit length of the down-stream message (from server to client) is $m+\kappa$, where n is the input length, m is the output length and κ is the security parameter. Information-theoretically, $n+m$ bits are necessary even if the server is fully trusted. To the best of our knowledge, all previous protocols, including ones which are based on fully homomorphic encryption, have a *multiplicative* overhead of κ, either with respect to n or to m. (See Section 7.3 for details.)

Non-Interactive Zero-Knowledge (NIZK). The complexity of NIZK has received much attention. The length of traditional NIZK proofs for NP grows linearly with the size of a circuit

[1] Similarly, FHE does not yield a non-interactive solution to the motivating problem described in the beginning of the introduction.

$R(x, w)$ which verifies that w is a legal witness for the statement $x \in L$. Using FHE, these traditional NIZKs can be converted into ones whose length is only $|w| + \text{poly}(\kappa)$ bits [Gen09, Gro11]. The proof consists of an FHE encryption c of w, along with a traditional NIZK proving that the ciphertext resulting from evaluating the verification algorithm on c encrypts the result of a correct verification. Thus, the prover's computation grows linearly with the time required for verifying $R(x, w)$, which can be an arbitrary polynomial in $|w|$. Moreover, there seems to be no obvious way to reduce this computational cost using offline preprocessing. Our results yield offline/online NIZK proofs with online proof length of $|w| + \text{poly}(\kappa)$ bits as before, but where the prover's online computation is nearly linear in $|w| + |x|$. This is done as follows. The common reference string of the NIZK defines a function f which maps w (along with a short seed which generates the prover's secret randomness) into a NIZK proof π. Applying our offline/online REs to this f yields the desired result. (See Section 7.2.) We note that while the length of NIZK *arguments* can be made sublinear in $|w|$ (under nonstandard but plausible assumptions), breaking this barrier in the case of *proofs* seems highly unlikely [GVW02].

1.3 Techniques

We briefly sketch some of the ideas used to prove Theorem 1.1. Our starting point is a standard garbled-circuit based encoding, such as the one from [AIK06]. In the offline phase of this encoding, we garble the circuit f and prepare, for each input i, a pair of random secret keys (K_i^0, K_i^1). In the online phase, for each i, we use the i-th bit of x to select a key $K_i^{x_i}$ and output the selected keys. In order to reduce the online complexity of the encoding, we would like to have a compact way to reveal the selected keys. Let us consider the following "riddle" which is a slightly simpler version of this problem. In the offline phase, Alice has n vectors $M_1, \ldots, M_n \in \{0,1\}^k$. She is allowed to send Bob a long encrypted version of these vectors. Later, in the online phase, she receives a bit vector $x \in \{0,1\}^n$. Her goal is to let Bob learn only the vectors which are indexed by x, i.e., $\{M_i\}_{i:x_i=1}$ while sending only a single message of length $O(n)$ bits (or even $n + \kappa$ bits).[2]

Before solving the riddle, let us further reduce it to an algebraic version in which Alice wants to reveal a 0-1 linear combination of the vectors which are indexed by x. Observe that if we can solve the new riddle with respect to nk-bit vectors $T = (T_1, \ldots, T_n)$, then we can solve the original riddle with k-bit vectors $(M_1, \ldots, M_n)$. This is done by placing the M_i's in the diagonal of T, i.e., T_i is partitioned to k-size blocks with M_i in the i-th block and zero elsewhere. In this case, Tx simply "packs" the vectors $\{M_i\}_{i:x_i=1}$.

It turns out that the linear version of the riddle can be efficiently solved via the use of a symmetric-key encryption scheme with some (additive) homomorphic properties. Specifically, let (E, D) be a symmetric encryption scheme with both key homomorphism and message homomorphism as follows: A pair of ciphertexts $\mathsf{E}_k(x)$ and $\mathsf{E}_{k'}(x')$ can be mapped (without any knowledge of the secret keys) to a new cipheretxt of the form $\mathsf{E}_{k+k'}(x+x')$. Given such a primitive the answer to the riddle is easy: Alice encrypts each vector under a fresh key K_i and publishes the ciphertexts C_i. At the online phase Alice sends the sum of keys $K_x = \sum K_i x_i$ together with the indicator vector x. Now Bob can easily construct $C = \mathsf{E}_{K_x}(Mx)$ by combining the ciphertexts indexed by x and, since K_x is known, Bob can decrypt the result. Intuitively, Bob learns nothing about a column M_j which is not indexed by x as the online key K_x is independent of the j-th key. Our

[2] The main difference between the riddle and the garbled-circuit problem is that in the latter case, the vector x itself should remain hidden; this gap is bridged by permuting the pairs and randomizing the vector x; see Section 4.

DDH and LWE based solutions are based on (approximate) implementations of this primitive. (A somewhat different approach is used in the RSA-based construction.)

The arithmetic setting is more challenging. Here, instead of computing the selection function, we should compute an affine function $Mx + v$ over the integers or over $\mathbb{Z}_p$, for some large integer p (not necessarily a prime). While it is possible to solve this via a similar encryption scheme with (stronger) additive homomorphism, there are several technical problems. Typically, all (or most) of the coordinates of x are non-zero and so we should argue that given K_x the secrecy of the key K_i was not compromised, despite the fact that K_i may participate in the linear combination K_x. This translates to some form of security under Related-Key attacks. In addition, it is harder to achieve homomorphism for integers or over $\mathbb{Z}_p$ directly, and so one should somehow embed this domain in a larger, less "friendly", message space. Still, it turns out that a variant of this gadget can be implemented based on the LWE assumption. Specifically, we use the following variant of the key-shrinking gadget of [AIK11] (which was originally introduced as a tool for garbling arithmetic circuits). Intuitively, we create a noisy version $\hat{M}$ and $\hat{v}$ of the matrix M and the vector v, and then plant them in a random linear space W of a low dimension κ over $\mathbb{Z}_q$ (where $q \gg p$). The space W is made public. Now every linear combination of $\hat{M}$ and $\hat{v}$ lies in W, and so it can be succinctly described by its coefficients with respect to W. In particular, to reveal the output $Mx + v$, it suffices for the encoding to reveal the coefficients of its representation $\hat{M}x + \hat{v}$. The security of the construction follows from the LWE assumption. See Section 5 for details.

Concurrent and subsequent works. The recent works [GKP+13b, GKP+13a] give the first *reusable* construction of garbled circuits. This implies REs in which a single offline computation can support an arbitrary polynomial number of efficient online computations. Following our work, [BGG+14] used multilinear maps to construct reusable REs with online communication complexity of $n + \text{poly}(\kappa, d)$ where d is the depth of the encoded function. The question of achieving an optimal online communication of $n + \text{poly}(\kappa)$ for reusable garbled circuits remains open. On a different front, improvements in the size of garbled circuits for uniform Turing Machine or RAM computations were recently given in [LO13, GKP+13a]. These lead to REs with succinct offline outputs. Our construction can be applied on top of these constructions, yielding REs with an online output of size $n + o(n)$, nearly linear online computation, and offline outputs that are only longer by an additive term of $O(n^{\varepsilon} \cdot T)$ than those in [LO13, GKP+13a], where T is the online computational complexity of the original constructions.

Organization. Section 2 gives the necessary background on randomized encodings (with some additional material in Appendix A). In Section 3, we present several constructions of succinct randomized encodings for a concrete boolean function called the subset function (SF). Later, in Section 4, we use these encodings as a building block and obtain succinct encodings for general boolean functions. The arithmetic case appears in Section 5. In Section 6, we deal with some lower bounds (Section 6.1) and the issue of adaptivity (Section 6.2). In Section 7, we sketch the application of succinct randomized encodings to secure multiparty computation (MPC), non-interactive zero-knowledge proofs (NIZK), and verifiable computation (VC) in the preprocessing model.

2 Randomized Encoding of Functions

Intuitively, a randomized encoding of a function $f(x)$ is a randomized mapping $\hat{f}(x;r)$ whose output distribution depends only on the output of f. We formalize this intuition via the notion of *computationally-private perfectly-correct randomized encoding* (in short RE) from [AIK06]. In the following, we assume that f is defined over $\mathbb{Z}_p^n$ for some integer p (by default $p=2$), and allow the encoding $\hat{f}$ be defined over a possibly larger alphabet $\mathbb{Z}_q^n$ for $p \leq q$ under the convention that a vector $x \in \mathbb{Z}_p^n$ can be naturally identified with a vector $x \in \mathbb{Z}_q^n$.

Definition 2.1 (Randomized Encoding (RE))**.** *Let $p = p(n), q = q(n)$ where $p(n) \leq q(n) \leq 2^{\mathrm{poly}(n)}$ and $\ell = \ell(n), m = m(n), s = s(n) = \mathrm{poly}(n)$ be integer valued functions. We naturally view $\mathbb{Z}_p$ as a subset of $\mathbb{Z}_q$. Let $f : \mathbb{Z}_p^n \to \mathbb{Z}_p^\ell$ be an efficiently computable function. We say that an efficiently computable randomized function $\hat{f} : \mathbb{Z}_q^n \times \{0,1\}^m \to \mathbb{Z}_q^s$ is a* perfectly-correct computationally-private randomized encoding *of f (in short, RE), if there exist an efficient decoder algorithm* Dec *and an efficient simulator* Sim *that satisfy the following conditions:*

- **Perfect correctness.** *For every $x \in \mathbb{Z}_p^n$, $\Pr_r[\mathsf{Dec}(1^n, \hat{f}(x;r)) \neq f(x)] = 0$.*
- (t, ε) **privacy.** *For every sequence $\{x_n\}_n$, where $x_n \in \mathbb{Z}_p^n$, and every $t(n)$-size circuit $\mathcal{A}$*

$$\left|\Pr[\mathcal{A}(\hat{f}(x_n;r)) = 1] - \Pr[\mathcal{A}(\mathsf{Sim}(1^n, f(x_n))) = 1]\right| \leq \varepsilon(n).$$

By default, we require (n^c, n^{-c})-privacy for every constant c, i.e., the distributions are computationally indistinguishable *(denoted by $\stackrel{c}{\equiv}$). The encoding is* statistically *secure if t is unbounded and* perfectly *secure if, in addition, $\varepsilon = 0$.*

Remarks.

- (Security parameter.) The above definition uses n both as an input length parameter and as a cryptographic "security parameter" quantifying computational privacy. When describing our constructions, it will be convenient to use a separate parameter κ for the latter, where computational privacy will be guaranteed as long as $\kappa \geq n^\varepsilon$ for some constant $\varepsilon > 0$.
- (Collections) Let $\mathcal{F}$ be a collection of functions with an associated representation (by default, a boolean or arithmetic circuit). We say that a class of randomized functions $\hat{\mathcal{F}}$ is an RE of $\mathcal{F}$ if there exists an efficient algorithm (compiler) which gets as an input a function $f \in \mathcal{F}$ and outputs (in time polynomial in the representation length $|f|$) three circuits $(\hat{f} \in \hat{\mathcal{F}}, \mathsf{Dec}, \mathsf{Sim})$ which form a $(t = n^{\omega(1)}, \varepsilon = n^{-\omega(1)})$-RE of f.

2.1 Efficiency Measures

So far the notion of RE can be trivially satisfied by taking $\hat{f} = f$ and letting the simulator and decoder be the identity functions. To make the definition non-trivial, we should impose some efficiency constraint. In this work, our main measure of efficiency is online complexity.

Online/Offline Complexity. We would like to measure separately the complexity of the outputs of $\hat{f}$ which depend solely on r (*offline* part) from the ones which depend both on x and r (*online* part). Without loss of generality, we assume that $\hat{f}$ can be written as $\hat{f}(x;r) = (\hat{f}_{\mathsf{off}}(r), \hat{f}_{\mathsf{on}}(x;r))$, where $\hat{f}_{\mathsf{off}}(r)$ does not depend on x at all. The *online communication complexity* (resp., *online computational complexity*) of $\hat{f}$ is the bit length (resp., the time complexity) of $\hat{f}_{\mathsf{on}}(x;r)$. Similarly, the *offline communication complexity* (resp., *offline computational complexity*) of $\hat{f}$ is the bit length (resp., the time complexity) of $\hat{f}_{\mathsf{off}}(r)$. The *rate* of $\hat{f}$ is ρ if the online communication complexity is at most ρ-times larger than the bit length $n \log p$ of the input of the encoded function f.

Efficient online encodings. Let $\hat{\mathcal{F}}$ be an encoding of the collection $\mathcal{F}$. We say that $\hat{\mathcal{F}}$ is *online-efficient* if for every function $f \in \mathcal{F}$, the online computational complexity of the encoding $\hat{f}$ is *independent* of the computational complexity (i.e., circuit size) of the encoded function f (but grows with the bit length of the input of f). The encoding is *online-succinct* (or simply succinct) if, in addition to being online efficient, every $f \in \mathcal{F}$ is encoded by a $1 + o(1)$-rate encoding.

Remark 2.2 (**Online inputs**). *In some applications, it is natural to think of the encoded function f as having online inputs x_{on} and offline inputs x_{off}. In this case, we measure the online commuincation/computational complexity of the encoding $\hat{f}$ with respect to the outputs that depend on x_{on}. By default, we simply assume that all the input x is an online input and there is no offline part.*

Some of the applications of REs further require some form of algebraic simplicity; this is captured by the notion of affinity.

Affine RE. We say that an encoding $\hat{f} : \mathbb{Z}_q^n \times \{0,1\}^m \to \mathbb{Z}_q^s$ is an *affine randomized encoding* (ARE) if, for every fixing of the randomness r, the online part of the encoding $\hat{f}_{\mathsf{on}}(x;r)$ becomes an affine function over the ring $\mathbb{Z}_q$, i.e., $\hat{f}_{\mathsf{on}}(x;r) = M_r \cdot x + v_r$, where M_r (resp., v_r) is a matrix (resp., vector) that depends on the randomness r.[3] It will sometimes be the case that certain outputs of $\hat{f}$ are restricted to an interval $[0, q']$ in $\mathbb{Z}_q$. Each such entry will only contribute $\lceil \log_2 q' \rceil$ towards computing the rate.

Remark 2.3 (**ARE vs. DARE**). *Previous works considered a stronger form of affinity called* decomposable affine randomized encoding *(DARE).*[4] *Decomposability requires that each output of $\hat{f}$ depends on a single deterministic input x_i. Hence, a decomposable affine randomized encoding can be written as $\hat{f}(x;r) = (\hat{f}_{\mathsf{off}}(r), \hat{f}_1(x_1;r), \ldots, \hat{f}_n(x_n;r))$ where each function $\hat{f}_i$ is affine with respect to x_i. It is known how to convert an ARE to DARE, however, the known transformation introduces a non-constant ($O(n)$) multiplicative blow-up in the online communication complexity. In Section 6.1, we show that this is inherent and decomposability* cannot *be achieved with constant rate.*

Remark 2.4 (**On Adaptive Security**). *In the online/offline model, it is natural to ask if the encoding can be* adaptively *secure, namely, if security holds when the online input x is chosen based on the offline part of the encoding (See Definition 6.5). We will show (Lemma 6.4) that, in the standard model, adaptively secure REs cannot be online-efficient, let alone have constant rate*

[3] We may assume WLOG that the "affine" representation of the encoding is given explicitly, as one can always "learn", for every fixed r, the matrix/vector M_r, v_r by solving a system of linear equations over $\mathbb{Z}_q$.

[4] In fact, in the conference version of [AIK11] the term ARE was used to denote DARE.

(assuming the existence of one-way functions). On the other hand, it turns out that this barrier can be bypassed via the use of a (programmable) random oracle (Lemma 6.7).

It is well known that REs can be manipulated via composition and concatenation [AIK04]. These standard properties (and others) are deferred to Section A.

3 Succinct AREs for the Subset Function

In order to succinctly encode boolean circuits, we will need a succinct encoding for the following concrete function g, called the *Subset Function*. It has length parameter n and message size κ and is defined by

$$g(M, x) = ((M_i)_{i \in x}, x),$$

where $M = (M_1, \ldots, M_n) \in (\{0,1\}^\kappa)^n$ is a vector of n "messages", and $x \in \{0,1\}^n$ is a selection vector which is viewed as the set $\{i : x_i = 1\}$. (The latter convention will be implicit through the whole section.) Our goal is to encode g by an RE of the form $\hat{g}(M, x; r) = (\hat{g}_{\mathsf{off}}(M; r), x, K(x; r))$ where $K(x; r)$ is of bit length κ^c for some universal constant c. Security will hold as long as n is bounded by some arbitrary polynomial in κ whose degree may be *independent* of the constant c. We will construct such an encoding based on several assumptions. Specifically, we will show (Section 3.1) that such an encoding can be based on a special form of symmetric-key encryption with additive homomorphism which, in turn, can be constructed under the DDH assumption (Section 3.2) or the LWE assumption (Section 3.3). We also present a direct encoding (which does not go through the additive homomorphism) under the RSA assumption (Section 3.4).

3.1 ARE for the Subset Function via Additive Homomorphic Encryption

Definition 3.1 (Additive Homomorphic Encryption (AHE))**.** *An additive homomorphic Encryption is a triple of efficient algorithms* $(\mathsf{Setup}, \mathsf{E}, \mathsf{D})$ *for which the following hold:*

- ***Syntax:*** *The randomized algorithm* Setup *takes a length parameter* 1^κ *and outputs a string* param *which specifies four (additive) groups: key-space* $\mathcal{K}$*, message-space* $\mathcal{M}$*, ciphertext-space* $\mathcal{C}$ *and public randomness space* $\mathcal{W}$*. We assume that* κ*-bit strings can be efficiently embedded in* $\mathcal{M}$ *and denote the identity element of* $\mathcal{M}$ *by* $\mathbf{0}$*. The inputs to the encryption and decryption algorithms consist of a message/ciphertext, a key* K*, some private randomness, and some public randomness* $W \xleftarrow{R} \mathcal{W}$ *which is selected during the encryption. Both algorithms also depend on the string* param*. (We make this dependency implicit to simplify notation.)*

- ***Semantic security:*** *Let* $\mathsf{param} = (\mathcal{K}, \mathcal{M}, \mathcal{C}, \mathcal{W}) \xleftarrow{R} \mathsf{Setup}(1^\kappa)$*. For every* $n = \mathrm{poly}(\kappa)$ *and every* n*-tuple of messages* $M_1, \ldots, M_n \in \mathcal{M}$*, we have that*

$$\left(\mathsf{param}, (W_i, \mathsf{E}_K(M_i; W_i))_{i \in [n]}\right) \stackrel{c}{\equiv} \left(\mathsf{param}, (W_i, \mathsf{E}_K(\mathbf{0}; W_i))_{i \in [n]}\right),$$

  where $W_i \xleftarrow{R} \mathcal{W}$*,* $K \xleftarrow{R} \mathcal{K}$*, and indistinguishability is parameterized by* κ*.*

- ***Additive Homomorphism:*** *For every* $n = \mathrm{poly}(\kappa)$ *and every* n*-tuple of keys* $K_1, \ldots, K_n \in \mathcal{K}$*,* n*-tuple of messages* $M_1, \ldots, M_n \in \mathcal{M}$*, and public randomness* $W \in \mathcal{W}$*, we have that*

$$\mathsf{D}_{\sum_i K_i}\left(\sum_i E_{K_i}(M_i; W); W\right) = \sum_i M_i,$$

where sums are computed over the corresponding groups. In fact, it suffices to have a relaxed *form of additive homomorphism which holds in the special case where all messages, except for one, equal to* $\mathbf{0} \in \mathcal{M}$.

The definition implies that the key size is *independent* of the homomorphism parameter n. This will be crucial for our applications. We show how to encode the subset function $g(M, x)$ with length n and message size κ based on AHE.

Lemma 3.2. *Assume that AHE exists. Then the Subset Function* $g(M, x)$*, where* $M \in (\{0,1\}^\kappa)^n, x \in \{0,1\}^n$*, has an encoding*

$$\hat{g}(M, x; r) = (\hat{g}_{\mathsf{off}}(M; r), x, \sum_{i \in x} K_i(r)),$$

where $\hat{g}_{\mathsf{off}}$ *outputs* $O(n^2)$ *ciphertexts in* $\mathcal{C}$*, the functions* K_i *output an element in* $\mathcal{K}$*, and the sum is computed over the key-space* $\mathcal{K}$.

Proof. At the offline phase, we invoke $\mathsf{Setup}(1^\kappa)$ and obtain a specification param of $\mathcal{K}$, $\mathcal{M}$, $\mathcal{C}$ and $\mathcal{W}$. We encode each entry of the offline input $M = (M_1, \ldots, M_n)$ by an element of $\mathcal{M}$, and from now on identify M_i with its encoding. We define a diagonal $n \times n$ matrix $\{M_{i,j}\}$ whose diagonal equals to the message vector M, i.e., $M_{i,i} = M_i, \forall i \in [n]$ and $M_{i,j} = \mathbf{0}, \forall i \neq j$. Next, we select a tuple of public random elements $W = (W_1, \ldots, W_n) \stackrel{R}{\leftarrow} \mathcal{W}^n$, a tuple of random keys $K = (K_1, \ldots, K_n) \stackrel{R}{\leftarrow} \mathcal{K}^n$ and compute a matrix of "ciphertexts" $C = (C_{i,j}) \in \mathcal{C}^{n \times n}$, where $C_{i,j} = \mathsf{E}_{K_i}(M_{i,j}; W_j)$. The output of $\hat{g}_{\mathsf{off}}$ consists of the tuple (param, W, C) and the online part $\hat{g}_{\mathsf{on}}$ consists of the pair $(x, K_x = \sum_{i \in x} K_i)$.

Decoding. Given $(\mathsf{param}, W, C, x, K_x)$, we decode $(M_i)_{i \in x}$ by exploiting the homomorphism property of the above encryption. Namely, for each $j \in x$ we compute

$$Y_j = \sum_{i \in x} C_{i,j} = \sum_{i \in x} \mathsf{E}_{K_i}(M_{i,j}; W_j),$$

and output the value $\mathsf{D}_{K_x}(Y_j; W_j)$.

Simulation. For $\ell = 0, \ldots, n$ define the hybrid $H_\ell(M, x)$ exactly as in $\hat{g}$ except that

$$M_{i,i} = \begin{cases} M_i & \text{if } i < \ell \text{ or } i \in x, \\ \mathbf{0} & \text{otherwise} \end{cases}$$

The first hybrid H_0 can be sampled based on $((M_i)_{i \in x}, x)$, and so it is being used as the simulator. The last hybrid H_n corresponds to the distribution of the encoding $\hat{g}$. Hence, by a standard argument, it suffices to show that each pair of neighboring hybrids is computationally indistinguishable. Assume, towards a contradiction, that $\mathcal{A}$ distinguishes the hybrid $H_{\ell-1}$ from H_ℓ with non-negligible advantage δ. Observe that in this case $x_\ell = 0$, as otherwise the two hybrids are identically distributed. We construct a new adversary $\mathcal{B}$ that breaks the semantic security of the scheme. Given a challenge $(\mathsf{param}, \vec{w}, \vec{c})$ where $\mathsf{param} \stackrel{R}{\leftarrow} \mathsf{Setup}(1^\kappa)$ and $\vec{w} = (w_1, \ldots, w_n) \stackrel{R}{\leftarrow} \mathcal{W}^n$, the adversary $\mathcal{B}$ distinguishes between

$$\vec{c} \stackrel{R}{\leftarrow} (\mathsf{E}_K(\mathbf{0}; w_1), \ldots, \mathsf{E}_K(\mathbf{0}; w_n)) \quad \text{and} \quad \vec{c} \stackrel{R}{\leftarrow} (\mathsf{E}_K(\mathbf{0}; w_1), \ldots, \mathsf{E}_K(M_\ell; w_\ell), \ldots, \mathsf{E}_K(\mathbf{0}; w_n)))$$

as follows. Use param to compute the hybrid $H_{\ell-1}$ where the public randomness $W_1, \ldots, W_n$ is set to $\vec{w}$, and the ℓ-th row of the ciphertext matrix C takes the value $\vec{c}$. It is not hard to verify that the resulting distribution is identical to $H_{\ell-1}$ if $\vec{c} \xleftarrow{R} (\mathsf{E}_K(\mathbf{0}; w_1), \ldots, \mathsf{E}_K(\mathbf{0}; w_n))$, and to H_ℓ if $\vec{c} \xleftarrow{R} (\mathsf{E}_K(\mathbf{0}; w_1), \ldots, \mathsf{E}_K(M_\ell; w_\ell), \ldots, \mathsf{E}_K(\mathbf{0}; w_n)))$, and the claim follows. □

Complexity. To encode the online part, one has to compute n additions (over the key space) and send x together with a single key element. The cost of the offline part is n^2 encryptions/ciphertexts and n public randomizers. One can obtain a smooth tradeoff between the offline part and the online part by partitioning the inputs to blocks (see Section 3.5). Also note that decoding costs n^2 additions over the key space (which can be reduced via the previous optimization) and n decryption operations. Finally, we mention that in our RSA-based solution (Section 3.4) the offline complexity is only linear in n but quadratic in κ. (The latter can be improved assuming sub-exponential hardness of RSA.)

3.2 AHE based on DDH

A DDH problem generator is a randomized algorithm which given a security parameter 1^κ outputs a specification param of a cyclic (multiplicative) group $\mathbb{G} = \langle \alpha \rangle$ of order p where p is κ-bit long prime and α is a group generator. The group order p is explicitly included in param. We say that the DDH assumption holds (with respect to DDH) if a random DDH tuple $(\mathsf{param}, \alpha, \alpha^a, \alpha^b, \alpha^{ab})$ is computationally indistinguishable from a random tuple $(\mathsf{param}, \alpha, \alpha^a, \alpha^b, \alpha^c)$ where $a, b, c \xleftarrow{R} \mathbb{Z}_p$. A DDH-based AHE can be constructed via the following symmetric-key version of ElGamal encryption. The algorithms $(\mathsf{Setup}, \mathsf{E}, \mathsf{D})$ are defined via

$$\mathsf{Setup}(1^\kappa) = (\mathcal{K} = \mathbb{Z}_p, \mathcal{M} = \mathcal{C} = \mathcal{W} = \mathbb{G}) \qquad \text{where } (\mathbb{G}, \mathbb{Z}_p) \leftarrow \mathsf{DDH}(1^\kappa),$$

and

$$\mathsf{E}_K(M; W) = W^K \cdot M, \qquad \mathsf{D}_K(C; W) = C / W^K.$$

(Note that for $\mathbb{G}$ we use *multiplicative* notation as opposed to the additive notation used in Definition 3.1.) The security of the scheme easily follows from the security of ElGamal public-key encryption.

Claim 3.3. *Under the DDH assumption, for any polynomial $n(\kappa)$, and any pair of n-tuple messages $(M_i)_{i\in[n]}$ and $(M'_i)_{i\in[n]}$*

$$\left(\mathsf{param}, (W_i, W_i^K \cdot M_i)_{i\in[n]}\right) \stackrel{c}{\equiv} \left(\mathsf{param}, W_i, W_i^K \cdot M'_i)_{i\in[n]}\right), \quad \textit{where } K \xleftarrow{R} \mathbb{Z}_p, W_i \xleftarrow{R} \mathbb{G}. \tag{1}$$

Proof. By the semantic security of ElGamal, Eq. 1 holds even when α^K is added to both ensembles. Since public-key ElGamal is secure under the DDH assumption, the claim follows. □

It is not hard to see that the scheme satisfies relaxed homomorphism. Indeed, for a vector of n messages $(M_1, \ldots, M_n)$ where all but a single message M_j are equal to the identity element $M_i = 1, \forall i \neq j$ we have

$$\mathsf{D}_{\sum_i K_i}\left(\prod_i \mathsf{E}_{K_i}(M_i; W); W\right) = \left(W^{\sum_i K_i} \cdot \prod_i M_i\right) / W^{\sum_i K_i} = M_j,$$

as needed.

Complexity. The key, the ciphertext, and the public randomizer are all of bit length κ. Hence, the encoding of Lemma 3.2 has online complexity of $n + \kappa$ and offline complexity of $O(n^2\kappa)$ based on DDH. (See Section 3.5 for a generic optimization.)

3.3 AHE based on LWE

LWE. The Learning With Errors (LWE) assumption of [Reg05] generalizes the Learning Parity with noise problem (LPN) and asserts that it is hard to solve a random system of noisy linear equations. Formally, let LWE be a problem generator which given a security parameter 1^κ outputs a modulus q, an integer μ and a noise-sampling circuit χ_μ which samples integers of absolute value bounded by μ. We say that the (decisional) LWE problem is *hard* if for every polynomial $t = t(\kappa)$, it holds that

$$(\mathsf{param}, W, Wk + e) \stackrel{c}{\equiv} (\mathsf{param}, W, z),$$

where

$$\mathsf{param} = (q, \mu, \chi_\mu) \stackrel{R}{\leftarrow} \mathsf{LWE}(1^\kappa), W \stackrel{R}{\leftarrow} \mathbb{Z}_q^{t\times\kappa}, k \stackrel{R}{\leftarrow} \mathbb{Z}_q^\kappa, e \stackrel{R}{\leftarrow} \chi_\mu^t, \text{ and } z \stackrel{R}{\leftarrow} \mathbb{Z}_q^t.$$

We will assume that the problem is hard for q which is super-polynomial in κ, e.g., $O(\kappa^{\log\kappa})$ and for some noise distribution χ_{q^α} where $\alpha \in (0, 1)$ is a constant. One can define such a problem generator (e.g., by letting χ_μ be "truncated" discrete gaussian) for which the hardness of LWE follows from the worst-case hardness of approximating shortest-vector problems in a lattice of dimension κ to within a quasi-polynomial ratio $2^{\mathrm{polylog}(\kappa)}$ [Reg05, Pei09]. (See also discussion in [AIK11].)

An LWE-based AHE can be constructed via the following LWE-based symmetric-key encryption which generalizes the LPN-based construction of [GRS08]. (See also [AHI11].) The algorithms $(\mathsf{Setup}, \mathsf{E}, \mathsf{D})$ are defined via

$$\mathsf{Setup}(1^\kappa) = (\mathcal{K} = \mathbb{Z}_q^\kappa, \mathcal{M} = \mathbb{Z}_2^\kappa, \mathcal{C} = \mathbb{Z}_q^\kappa, \mathcal{W} = \mathbb{Z}_q^{\kappa\times\kappa}) \qquad \text{where } (q, \mu, \chi_\mu) \stackrel{R}{\leftarrow} \mathsf{LWE}(1^\kappa),$$

and

$$\mathsf{E}_K(M; W, E) = WK + e + \Delta M, \qquad \mathsf{D}_K(C; W) = \lfloor (C - WK)/\Delta \rceil,$$

where $e \stackrel{R}{\leftarrow} \chi_\mu^\kappa$, $\Delta = \lfloor q/2 \rceil$ and the operator $\lfloor\cdot\rceil$ denotes rounding to the closest integer. The semantic security of the scheme follows immediately from the LWE assumption (cf. [GRS08]). Furthermore, it is not hard to see that the scheme satisfies homomorphism. For $n = \mathrm{poly}(\kappa)$ messages, the "merged" ciphertext

$$\sum_{i=1}^n \mathsf{E}_{K_i}(M_i; W) = W \sum_{i=1}^n K_i + \Delta \sum_{i=1}^n M_i + \sum_{i=1}^n e_i$$

contains noise whose magnitude is bounded by $n \cdot q^\alpha < \Delta/2$ and therefore decryption succeeds.

Complexity. The bit length of the key and ciphertext is $\tilde{O}(\kappa)$ and the bit length of the public randomizer is $\tilde{O}(\kappa^2)$. Hence, if we instantiate the encoding of Lemma 3.2 with LWE-based AHE, we obtain an online complexity of $n + \tilde{O}(\kappa)$ and offline complexity of $n^2\tilde{O}(\kappa) + n\tilde{O}(\kappa^2)$ (the first term comes from the matrix of ciphertexts and the second term comes from the array of public randomizers). Again, the reader is referred to Section 3.5 for a generic optimization.

3.4 Encoding SF based on RSA

The RSA assumption [RSA78] asserts that for every efficient adversary $\mathcal{A}$ of complexity $\text{poly}(\kappa)$

$$\Pr[\mathcal{A}(N, e, \alpha^e) = \alpha] \leq \text{neg}(\kappa),$$

where N is a random κ-bit RSA modulus (i.e., product of a pair of random primes p, q) e is a randomly chosen prime of length κ which is co-prime to $\varphi(N)$, and $\alpha \xleftarrow{R} \mathbb{Z}_N$. (More generally, p, q, e can be chosen according to some other distribution specified by some efficient problem generator $\mathsf{Gen}(1^\kappa)$.) We present an RSA based encoding for the subset function. We begin with an encoding for the subset function with length n and block size of 1. In this simple case, the input consists n single bit messages $m = (m_1, \ldots, m_n)$ and a selection vector $x \in \{0,1\}^n$ and it outputs the messages $(m_i)_{i \in x}$ chosen by x, together with the selection vector x. We will later show (Lemma 3.5) that such an encoding can be upgraded to encode the SF with κ-bit messages via simple concatenation.

Lemma 3.4. *Under the RSA assumption, the simplified SF $h(m, x)$ where $x \in \{0,1\}^n$ and $m \in (\{0,1\})^n$, has an encoding of the form $\hat{h}(x, m; r) = (\hat{h}_{\mathsf{off}}(m; r), x, K(x; r))$ where $\hat{h}_{\mathsf{off}}$ is of bit length $n\kappa$ and $K(x; r)$ is of length κ.*

Proof. The encoding $\hat{h}$ relies on a symmetric variant of RSA encryption. At the offline phase, generate a random RSA modulus N together with its factorization p and q, choose a random $u \xleftarrow{R} \mathbb{Z}_N$ and n random primes $e_1, \ldots, e_n$ of length κ which are all co-prime to $\varphi(N)$, and a string $r \xleftarrow{R} \mathbb{Z}_2^\kappa$. For every $i \in [n]$ compute

$$y_i = u^{1/e_i}, \quad C_i = m_i \oplus \text{hc}(r, y_i),$$

where hc is the Goldreich-Levin hardcore predicate (i.e., inner-product over $\mathbb{Z}_2$). The offline part of the encoding is $(N, u, (e_1, \ldots, e_n), r, (C_1, \ldots, C_n))$, and the online part is the pair $(x, v = u^{\prod_{i \in x} 1/e_i})$.

Decoding. For $i \in x$, recover y_i by computing $v^{\prod_{j : \in (x \setminus i)} e_j} = u^{e_i}$ and let m_i be $C_i \oplus \text{hc}(r, y_i)$.

Simulation. Fix some $x \in \{0,1\}^n$ and $m = (m_1, \ldots, m_n) \in \{0,1\}^n$. Given $(x, (m_i)_{i \in x})$ we simulate the distribution $\hat{h}(m, x; r)$ by sampling $\hat{h}(m', x; r)$ where m'_i equals to m_i if $i \in x$, and $m'_i \xleftarrow{R} \{0,1\}$ otherwise. We prove that $\hat{h}(m, x; r)$ is computationally indistinguishable from $\hat{h}(m', x; r)$ via a hybrid argument as follows.

Security. For $i \in [n]$ define the hybrid distribution H_i by $\hat{h}((m_{1:i} | m'_{i+1:n}), x; r)$. Clearly, H_0 corresponds to the simulated distribution while H_n corresponds to the distribution of the real encoding. Suppose there exists a distinguisher $\mathcal{A}$ that distinguishes $H_{\ell-1}$ from H_ℓ with non-negligible advantage ε. Observe that it must be the case that $\ell \notin x$; otherwise the distributions $H_{\ell-1}$ and H_ℓ are the identical. We use $\mathcal{A}$ to solve an RSA challenge $(N, e, z = \alpha^e)$ as follows. First, for every $i \neq \ell$ we choose a random κ-bit prime e_i and let

$$u = z^{\prod_{i \neq \ell} e_i}, \quad v = z^{\prod_{i \neq \ell, i \notin x} e_i}, \qquad y_i = z^{\prod_{j \neq \ell, i} e_j}.$$

Letting $e_\ell = e$ we can write

$$u = \alpha^{\prod_i e_i}, \quad v = \alpha^{\prod_{i \notin x} e_i}, \qquad y_i = \alpha^{\prod_{j \neq i} e_j}.$$

Let us condition on the event that all the e_i's are co-prime to $\varphi(N)$ (which happens with all but negligible probability). In this case, the e_i's are distributed exactly as in the real encoding, and u is uniformly and independently distributed in $\mathbb{Z}_N$ (since $\alpha \stackrel{R}{\leftarrow} \mathbb{Z}_N$ and exponentiation to the power of $\prod_i e_i$ induces a permutation over $\mathbb{Z}_N$). Furthermore, v equals to $u^{\prod_{i\in x}}$ and y_i equals to u^{1/e_i} for $i \neq \ell$. Hence, the joint distribution of N, u, the e_i's, and the y_i's is exactly as in the real encoding. The only missing information is y_ℓ which should take the value $u^{1/e_\ell} = \alpha^{\prod_{i\neq\ell} e_i}$. We will use $\mathcal{A}$ to recover u^{1/e_ℓ} and then use it to find α.

By the Goldreich-Levin theorem, in order to recover u^{1/e_ℓ} it suffices to construct an algorithm $\mathcal{B}$ that distinguishes the pair $(r \stackrel{R}{\leftarrow} \mathbb{Z}_2^\kappa, \sigma = \mathrm{hc}(u^{1/e_\ell}, r))$ from the pair $(r \stackrel{R}{\leftarrow} \mathbb{Z}_2^\kappa, \sigma \stackrel{R}{\leftarrow} \{0,1\})$ with noticeable advantage. To achieve this we let $\mathcal{B}(r,\sigma)$ be the outcome of $\mathcal{A}(N, u, (e_i)_{i\in[n]}, r, (C_i)_{i\in[n]}, x, v)$ where

$$C_i = \begin{cases} m_i \oplus \mathrm{hc}(y_i, r) & \text{if } i < \ell \text{ or } i \in x, \\ m_i \oplus \sigma & \text{if } i = \ell, \\ R_i \stackrel{R}{\leftarrow} \{0,1\} & \text{if } i > \ell \text{ and } i \notin x \end{cases}.$$

It is not hard to verify that, when $(N, z, e_1, \ldots, e_n, r)$ are uniformly chosen, the resulting distribution corresponds to H_ℓ if $\sigma = \mathrm{hc}(u^{1/e_\ell}, r)$, and to $H_{\ell-1}$ if $\sigma \stackrel{R}{\leftarrow} \{0,1\}$. Hence, by Markov's inequality, with probability at least $\varepsilon/2$ the tuple $(N, z, e_1, \ldots, e_n)$ is *good* in the sense that $\mathcal{B}$ has distinguishing advantage of $\varepsilon/2$. Therefore, we recover u^{1/e_ℓ} with noticeable probability. Finally, we employ Shamir's algorithm [Sha83] which, given $X, Y \in \mathbb{Z}_N$ and relatively prime integers a, b for which $X^a = Y^b$, efficiently computes $Y^{1/a}$. Letting $X = u^{1/e_\ell}, Y = z$ and $a = e_\ell, b = \prod_{j\neq\ell} e_j$, we recover the RSA solution z^{1/e_ℓ}. □

The above encoding can be easily used to encode the subset function with κ-bit messages.

Lemma 3.5. *Under the RSA assumption, the SF $g(M,x)$ where $x \in \{0,1\}^n$ and $M \in (\{0,1\}^\kappa)^n$, has an encoding of the form $\hat{g}(x, M; r) = (\hat{g}_{\mathsf{off}}(M; r), x, K(x; r))$ where $\hat{g}_{\mathsf{off}}$ is of bit length $n\kappa^2$ and $K(x; r)$ is of length κ^2.*

Proof. Observe that $g(M,x)$ is (deterministically) encoded by $(h(x, m^i))_{i=1}^\kappa$ where

$$m^i = (M_{1,i}, \ldots, M_{n,i}).$$

By the concatenation lemma, the latter function can be encoded by concatenating the encodings $\hat{h}(x, m^i; r^i)$ for $i \in [\kappa]$ from Lemma 3.4. By the composition lemma, the resulting function also encodes g. This encoding almost satisfies the lemma except that there are κ copies of x. These multiple copies can be replaced by a single copy (formally, think of x as a deterministic encoding of its copies and invoke the composition lemma) leading to an encoding that satisfies the lemma. □

Remark 3.6 (Optimization)**.** *Suppose that RSA over modulus of bit length $\tau = \tau(\kappa)$ is secure against $2^{\Theta(\kappa)}$-time adversaries. In this case, one can extract roughly κ independent pseudorandom bits by using κ independent copies of the Goldreich-Levin hardcore predicate (or, in practice, via a good hash function). Applying this optimization to Lemma 3.4 (i.e., replacing the single hardcore bit with κ hardcore bits) allows to encode the SF with respect to a message of length κ, with offline complexity (bit length of $\hat{g}_{\mathsf{off}}$) of $n\tau$ and online complexity (bit length of $K(x; r)$) of τ. Note that, from a "concrete security" point of view, we should use a τ-bit modulus anyway in order to guarantee an overall security of κ bits. Hence, this optimization is essentially for free.*

3.5 Reducing the Offline Complexity

Given a succinct encoding for SF, one can obtain a new encoding with a smooth tradeoff between the online and the offline complexity as follows. Let $h(M', x')$ be the subset function with length N and message size κ, i.e., $M' \in (\{0,1\}^\kappa)^N, x' \in \{0,1\}^N$, and assume that we have an encoding $\hat{h}(M', x') = (\hat{h}_{\mathsf{off}}(M'; r), x', K'(x'; r))$ where $\hat{h}_{\mathsf{off}}$ has of computational complexity of $N^a \kappa^{a'}$ and K' is of length κ^b and computational complexity of $N\kappa^b$ for some constants a, a', b.[5] In order to encode the SF $g(M, x)$ with input length n, partition the input to n/N blocks of size N each, i.e., let $M^i = (M_{iN+1} \dots M_{(i+1)N})$ and $x^i = (x_{iN+1} \dots x_{(i+1)N})$, and think of $g(M, x)$ as the concatenation of $h(M^i, x^i)$ for $i = 1, \dots, n/N$. Now encode g by the encoding

$$\hat{g}(M, x; (r^1, \dots, r^{n/N})) = (\hat{h}(M^i, x^i; r^i))_{i=1}^{n/N}$$

reordering the outputs and letting $r = (r^1, \dots, r^{n/N})$, we can write $\hat{g}(M, x; r)$ as $\hat{g}_{\mathsf{off}}(M; r), x, K(x; r))$. Let $\mathsf{Comp}(f)$ denote the computational complexity (circuit size) of a function f, and $\mathsf{Len}(f)$ denotes the output length of f (in bits). Then, the complexity of the new encoding satisfies

$$\begin{aligned}
\mathsf{Comp}(\hat{g}_{\mathsf{off}}) &= \frac{n}{N} \cdot \mathsf{Comp}(\hat{h}_{\mathsf{off}}) = nN^{a-1}\kappa^{a'} \\
\mathsf{Len}(K) &= \frac{n}{N} \cdot \mathsf{Len}(K') = \frac{n}{N}\kappa^b \\
\mathsf{Comp}(K) &= \frac{n}{N} \cdot \mathsf{Comp}(K') = n\kappa^b
\end{aligned}$$

Hence, a larger value of N reduces the online complexity, while a smaller value reduces the offline complexity. By letting $N > \omega(\kappa^b)$, the online communication remains $n + o(n)$ (as κ is polynomially related to n). Combining the above with Lemma 3.5 and Lemma 3.2 and the LWE/DDH constructions from Sections 3.2 and 3.3, we derive the following lemma:

Lemma 3.7 (Encoding SF)**.** *Assume that the DDH assumption, LWE assumption or the RSA assumption holds. There exists a universal constant C such that for every n and κ for which $n^{\Omega(1)} < \kappa < o(n)$ the Subset Function $g(M, x)$ with length n and message size κ has an RE of the form $\hat{g}(x, M; r) = (\hat{g}_{\mathsf{off}}(M; r), x, K(x; r))$ where:*

- *$K(x; r)$ is of bit length at most $o(n)$.*
- *The encoding $\hat{g}$ (including both the online and offline parts) can be computed in time $n\kappa^C$.*
- *In the case of DDH and LWE $K(x; r)$ is affine in x (for every fixed value of r).*

4 Succinct AREs for Boolean Circuits

In this section, we present a succinct encoding for any polynomial-time computable function. We begin by showing that if $F : \{0,1\}^n \to \{0,1\}^\ell$ has a decomposable affine randomized encoding (DARE) then it also has a succinct encoding. In the following, let κ be a security parameter which is polynomially related to n, i.e., $\kappa = n^\delta$ for some fixed $\delta > 0$. We will employ a succinct

[5] For example, for DDH $a = 2, a' = 1$ and $b = 1$, for LWE $a = 2, a' = 2$ and $b = 1$, and for RSA $a = 1, a' = 2$ and $b = 2$.

encoding for the subset function $g(M, \hat{x})$ with length $N = 2n$ and message size κ. We will also make use of the following simple observation: if a $\kappa \times 2n$ matrix M is composed of n pairs of columns $(M_{2i-1}|M_{2i}) = (v_i^0, v_i^1)_{i\in[n]}$, then for any $x \in \{0,1\}^n$ the sub-matrix $(v_i^{x_i})_{i\in[n]}$ can be written as $(M_i)_{i\in \mathrm{pad}(x)}$, where $\mathrm{pad}(x)$ maps an n-bit vector x to the $2n$-bit vector $(1-x_1, x_1, \ldots, 1-x_n, x_n)$, and $i \in \mathrm{pad}(x)$ if $\mathrm{pad}(x)_i = 1$.

Lemma 4.1. *Let $F : \{0,1\}^n \to \{0,1\}^\ell$ be an efficiently computable function having a decomposable ARE $f(x;\rho) = (f_{\mathsf{off}}(\rho), f_1(x_1;\rho), \ldots, f_n(x_n;\rho))$, where the output length of each f_i is κ bits. Also, assume that the subset function $g(M, \hat{x})$ with length $2n$ and message size κ has an RE of the form $\hat{g}(M, \hat{x}; r) = (\hat{g}_{\mathsf{off}}(M; r), \hat{x}, K(\hat{x}; r))$. Then, F is encoded by the randomized function*

$$\hat{F}(x;\rho,s,r) = (f_{\mathsf{off}}(\rho), \hat{g}_{\mathsf{off}}(M;r), x \oplus s, K(\mathrm{pad}(x \oplus s); r)),$$

where

$$M = (f_1(s_1;\rho)|f_1(s_1 \oplus 1;\rho)| \cdots |f_n(s_n;\rho)|f_n(s_n \oplus 1;\rho)) \in \{0,1\}^{\kappa \times 2n}.$$

Proof. It will be useful to start by encoding the n-wise one-out-of-two selection function H which maps an online input $x \in \{0,1\}^n$ and an offline matrix of pairs $V = (v_1^0|v_1^1|\ldots|v_n^0|v_n^1) \in \{0,1\}^{\kappa\times 2n}$ to the tuple $(v_i^{x_i})_{i\in[n]}$. Observe that the output of H is essentially the value of the subset function g applied to the matrix V and the vector $\mathrm{pad}(x) \in \{0,1\}^{2n}$, except that H hides x whereas g reveals it. Nevertheless one can easily randomize x and then employ the subset function. Specifically, select a random mask $s \xleftarrow{R} \{0,1\}^n$, let $\hat{x} \in \{0,1\}^{2n}$ be the vector $\mathrm{pad}(x \oplus s)$, and construct the $\kappa \times 2n$ matrix $M = (v_1^{s_1}|v_1^{s_1\oplus 1}|\ldots|v_n^{s_n}|v_n^{s_n\oplus 1})$. It is not hard to show that the randomized mapping $h(V, x; s) \mapsto g(M, \hat{x})$ is an encoding of H. Indeed, the output distribution of $g(M, \hat{x})$ consists of the matrix $(M_i)_{i\in\hat{x}}$ and the vector $\hat{x}$ — the former simply equals to $(v_i^{x_i})_{i\in[n]}$ and the latter is just a sequence of n pairs of a random bit and its complement.

Next, let us view h as a deterministic function of V, x and s. Since h can be written as $g(M_{V,s}, \hat{x}_{x,s})$, we can apply the substitution lemma (Fact A.1) and encode h by the mapping $\hat{g}(M_{V,s}, \hat{x}_{x,s}; r)$. By the composition lemma (Fact A.3), the latter encoding also encodes H. Overall, our encoding for $H(V, x)$ is defined as follows:

$$(V, x; s, r) \mapsto (\hat{g}_{\mathsf{off}}(M_{V,s}; r), \mathrm{pad}(x \oplus s), K(\mathrm{pad}(x \oplus s); r)).$$

To improve the online complexity, we replace the redundant value $\mathrm{pad}(x \oplus s)$, which is sent in the clear, with $x \oplus s$. The encoding is still valid as $x \oplus s$ is a (deterministic) encoding of $x \oplus s$.

We can now prove the lemma. Let us view ρ as a deterministic input and encode the deterministic function $f(x, \rho)$. Since f is decomposable, we can write it as

$$(f_{\mathsf{off}}(\rho), H(V_\rho, x)), \quad \text{where } V_\rho = (f_1(0;\rho)|\ldots|f_n(0;\rho)|f_1(1;\rho)|\ldots|f_n(1;\rho))$$

and H is the n-wise one-out-of-two selection function. Therefore, by the substitution and concatenation lemmas (Facts A.1 and A.2), f can be encoded by $(f_{\mathsf{off}}(\rho), \hat{h}(V, x; s, r))$, where $\hat{h}$ encodes H. Plugging in our (improved) encoding of H, we obtain an encoding of the form

$$\hat{f}(x, \rho; s, r) = (f_{\mathsf{off}}(\rho), \hat{g}_{\mathsf{off}}(M_{s,\rho}; r), x \oplus s, K(\mathrm{pad}(x \oplus s); r)).$$

By the composition lemma (Fact A.3), the function $\hat{f}(x; \rho, s, r)$ encodes $F(x)$ and the lemma follows. □

It follows that F has an encoding with online complexity of $n + \mathsf{Len}(K)$, online computational complexity of $O(n + \mathsf{Comp}(K))$, and offline computational complexity of $\mathsf{Comp}(f_{\mathsf{off}}) + \mathsf{Comp}(\hat{g}_{\mathsf{off}})$, where $\mathsf{Comp}(\cdot)$ and $\mathsf{Len}(\cdot)$ measure the computational complexity (circuit size), and the output length (in bits) of a given function. Furthermore, observe that for every fixed randomness s each bit of the term $\mathrm{pad}(x \oplus s)$ can be written as x_i or as $1 - x_i$ and so if $K(\hat{x}; r)$ is affine (over some ring) then so is $\hat{F}_{\mathsf{on}}$.

In [AIK06] it is shown that, assuming the existence of one-way functions, any efficiently computable function $F(x)$ can be encoded by a decomposable ARE $f(x; \rho) = (f_{\mathsf{off}}(\rho), f_1(x_1; \rho), \ldots, f_n(x_n; \rho))$, where the output length of the f_i's is κ bits, and the computational complexity of f_{off} is $\kappa \cdot \mathsf{Comp}(f)$. Combining this with Lemma 4.1 and our encodings for the Subset Function, we derive succinct encodings for general boolean functions. By using the optimized encoding of Lemma 3.7, we can do this while keeping the online computational complexity asymptotically "almost linear", as in the following theorem.

Theorem 4.2 (Theorem 1.1 restated)**.** *Assume that the DDH assumption, or LWE assumption or the RSA assumption holds. Let $\varepsilon > 0$ be an arbitrary constant. Then, every efficiently computable function $F : \{0,1\}^n \to \{0,1\}^{\ell(n)}$ has an encoding $\hat{F}$ with the following properties:*

- *The online communication is $n + o(n)$ and the online computational complexity is $O(n^{1+\varepsilon})$.*
- *The offline computational/communication complexity is $O(n^{\varepsilon}\mathsf{Comp}(F))$.*
- *In the case of LWE and DDH the encoding is affine.*

Proof. Let $\kappa = n^{\delta}$ for sufficiently small constant δ whose value will be determined later. By Lemma 3.7, we obtain an encoding for the $\hat{g}$ that satisfies the properties of Lemma 4.1. Furthermore, $K(x; r)$ is of bit length at most $o(n)$ and the encoding is computable in time $O(n\kappa^C)$ for some universal constant C. By [AIK06] F has a DARE f and so, by applying Lemma 4.1 we obtain an encoding with online communication of $n+o(n)$, online complexity of $O(n\kappa^C)$ and offline complexity of $O(|F|\kappa) + O(n\kappa^C)$. The theorem now follows by letting $\delta = \varepsilon/(2C)$. □

Remarks.

- (Reduction) The proof of Lemma 4.1 shows that the task of succinctly encoding a function F that admits an online efficient DARE reduces (information-theoretically) to the task of succinctly encoding the subset function g.
- (General compiler) Theorem 4.2 is constructive, i.e., it describes a compiler that given a circuit for F outputs a description of the encoding $\hat{F}$, its decoder and its simulator.
- (Online inputs) The theorem generalizes to the case where the function f has some online inputs x_{on} and offline inputs x_{off} as in Remark 2.2. Namely, the part of the encoding $\hat{f}$ which depends on x_{on} is of length $|x_{\mathsf{on}}| + o(|x_{\mathsf{on}}|)$.

5 Succinct AREs for Arithmetic Formulas

In this section we construct a succinct ARE for arithmetic formulas over subexponentially large modulus p (i.e., $p = 2^{o(n)}$) as follows.

Theorem 5.1. *Let $F : \mathbb{Z}_p^n \to \mathbb{Z}_p^\ell$ be an efficiently-computable arithmetic formula where $p = \Theta(2^{n^\delta})$ for some $\delta \in (0,1)$. Then, assuming LWE, F has a succinct ARE $\hat{F}$ over $\mathbb{Z}_q$ where $q = \Theta(2^{n^{\delta'}})$ for some $\delta' \in (\delta, 1)$ of the following form:*

$$\hat{F}(x; r, s) = (\hat{F}_{\mathsf{off}}(r), \hat{x} = x + s \pmod p, K(\hat{x}, r)),$$

where $K(\hat{x}, r)$ is short (of length $o(n \log p)$) and affine over $\mathbb{Z}_q$.

We note that by lifting mod-q computations to the integers, we can get an encoding over the integers with constant rate. The theorem will be proven in two steps. In Section 5.1, we show that it suffices to obtain a succinct encoding for the mod-q *universal affine function* (AF), and in Section 5.2, we construct such an encoding under the LWE assumption.

5.1 Reduction to the Affine Function

Let $F : \mathbb{Z}_p^n \to \mathbb{Z}_p^\ell$ be an efficiently-computable arithmetic formula where $p = \Theta(2^{n^\delta})$ for some $\delta \in (0,1)$. Let $m = m(n)$ be some polynomial (whose value will be related to the size of F). The *universal affine function* (AF) $g = g_{n,m,p}$ over $\mathbb{Z}_p$ is defined by

$$g(M, v, \hat{x}) \mapsto (M\hat{x} + v, \hat{x}), \quad \text{where } M \in \mathbb{Z}_p^{m \times n}, v \in \mathbb{Z}_p^m, \hat{x} \in \mathbb{Z}_p^n.$$

Lemma 5.2 (Succinct ARE for mod-p formulas)**.** *The function F has a succinct ARE $\hat{F}$, assuming that the function g has an ARE $\hat{g}$ of the form*

$$\hat{g}(M, v, \hat{x}; r) = (\hat{g}_{\mathsf{off}}(M, v; r), \hat{x}, K(\hat{x}; r))$$

where $K(\hat{x}; r) \in \mathbb{Z}_q^\kappa$ is an affine function in $\hat{x}$ and $\kappa \log q = o(n \log p)$.

Proof. Let $F : \mathbb{Z}_p^n \to \mathbb{Z}_p^\ell$ be a function computable by $s = \mathrm{poly}(n)$-size arithmetic formula. In [IK02, AIK04] (see also [CFIK03]) the function F is information-theoretically encoded via an ARE $f(x; \rho) = (f_{\mathsf{off}}(\rho), f_{\mathsf{on}}(x; \rho))$ over $\mathbb{Z}_p$. Namely, $f_{\mathsf{on}}(x; \rho) = M_\rho x + v_\rho$ where $M_\rho \in \mathbb{Z}_p^{m \times n}$ and $v_\rho \in \mathbb{Z}_p^m$ are computed based on the randomness ρ, and $m = O(\ell s^2)$.

By the composition and concatenation lemmas (Facts A.2 and A.3), it suffices to encode the online-part $f_{\mathsf{on}}(x; \rho)$ (viewed as a single argument function) by a succinct encoding $\hat{f}_{\mathsf{on}}(x, \rho; R)$. Indeed, in this case $F(x)$ can be encoded by $\hat{F}(x; (\rho, R)) = (f_{\mathsf{off}}(\rho), \hat{f}_{\mathsf{on}}(x; (\rho, R)))$. Furthermore, by the substitution lemma (Fact A.1), we may simply encode the affine function $L : (M, v, x) \to Mx + v$. Observe that unlike g the function L does not reveal x. However, one can succinctly encode L by g as follows.

Claim 5.3. *The function $L(M, v, x)$ is perfectly encoded via the encoding*

$$h : (M, v, x; s) \mapsto g(M, v', \hat{x}), \qquad \text{where } s \xleftarrow{R} \mathbb{Z}_p^n, \quad v' = v - Ms, \quad \hat{x} = x + s.$$

Proof. Decoding is trivial as $M\hat{x} + v' = Mx + v$. Given an output y of L we perfectly simulate $h(M, v, x; s)$ via (y, z) where $z \xleftarrow{R} \mathbb{Z}_p^n$. □

We can now complete the proof of the lemma. Consider the encoding

$$\hat{h} : ((M, v, x, s); r)) \mapsto \hat{g}(M, v' = v - Ms, \hat{x} = x + s; r) = (\hat{g}_{\mathsf{off}}(M, v'; r), \hat{x}, K(\hat{x}; r)).$$

Then, by the substitution lemma (Fact A.1), $\hat{h}$ encodes the function h, and so, by the composition lemma (Fact A.3), it also encodes L. Combining everything together, we encode the function F via an ARE

$$\hat{F}(x; (\rho, s, r)) = (f_{\mathsf{off}}(\rho), \hat{g}_{\mathsf{off}}(M_\rho, v_\rho - M_\rho s; r), x + s, K(x + s; r))$$

with optimal online rate, as promised. □

5.2 ARE for the Universal Affine Function

Our goal is to encode the universal affine function $g_{n,m,p}$. We will make use of the following fact which "lifts" g to the integers.

Fact 5.4. *The function $g_{n,m,p}(M, v, x)$ can be encoded via the function*

$$(M, v, x; R) \mapsto (Mx + M_0, x)$$

where $R \xleftarrow{R} [0, p^2]^m$, $M_0 = v + pR$, and both addition and multiplication are computed over the integers. The encoding has statistical privacy error of $O(mnp^2/p^3)$ which is negligible in n.

The fact follows from [AIK11, Lemma 5.2]. Hence, it suffices to construct a succinct encoding for the function $g'(M, M_0, x) = (Mx + M_0, x)$ computed over the integers where $M \in [0 : p - 1]^{m \times n}, M_0 \in [0 : 2p^3]^m$ and $x \in [0 : p]^n$. We will construct such an encoding based on the LWE problem with the following parameters.

Parameters. Let $\alpha, \varepsilon_1, \varepsilon_2 \in (0, 1)$ be constants such that $\varepsilon_1\varepsilon_2 > \delta$ (recall that $p = \Theta(2^{n^\delta})$). We will base our encoding on the assumption that LWE is hard for dimension $\kappa = n^{\varepsilon_1}$, modulus $q = \Theta(2^{\kappa^{\varepsilon_2}})$, and noise distribution χ_{q^α} which samples integers bounded by q^α. In fact, we will need a family of distributions χ_μ which are almost invariant under small shifts. Namely, for every integer B the statistical distance between χ_μ and $\chi_\mu + B$ is bounded by $\mathrm{poly}(|B|/\mu)$. Standard noise distributions (e.g., discrete gaussian, or uniform over μ-size interval) have this property (cf. [Reg05, AIK11]). Overall, for proper choice of parameters our assumption is implied by the worst-case hardness of approximating the shortest vector in a κ-dimensional lattice to within a subexponential ratio.

The encoding. We will use a variant of the LWE-based gadget from [AIK11]. Let $\beta \in (\alpha, 1)$ and $\Delta = 3npq^\beta$. At the offline phase, select a random public matrix $W \xleftarrow{R} \mathbb{Z}_q^{m \times \kappa}$, a random secret key matrix $K \xleftarrow{R} \mathbb{Z}_p^{\kappa \times n}$ and a noise matrix $E \xleftarrow{R} \chi_{q^\alpha}^{m \times n}$. Then compute an $m \times n$ "padding" matrix $Y = WK + E \pmod q$ and an $m \times (n + 1)$ "ciphertext" matrix $C = Y + \Delta \cdot M \pmod q$. In addition, let $K_0 \xleftarrow{R} \mathbb{Z}_q^\kappa$ and $E_0 \xleftarrow{R} \chi_{q^\beta}^\kappa$ and compute $C_0 = WK_0 + E_0 + \Delta \cdot M_0$. The output is

$$\hat{g}_{\mathsf{off}}(M, M_0) = (W, C, C_0), \qquad \hat{g}_{\mathsf{on}}(x) = (x, \hat{K} = K \cdot x + K_0 \pmod q)$$

Decoding. Given $(W, C, C_0, x, \hat{K})$ we decode the integer vector $Mx + M_0$ as follows: (1) compute $M' = Cx + C_0 - W\hat{K}$ over the integers; (2) To "clean" the noise divide each entry of M' by Δ and round to the closest integer. Output the resulting vector $\lfloor M'/\Delta \rceil$.

Claim 5.5. *The outcome of the decoder equals to $Mx + M_0$ over the integers.*

Proof. First observe that over $\mathbb{Z}_q$ the matrix M' equals to

$$Cx+C_0-W\hat{K} = WKx+Ex+WK_0+E_0+\Delta(Mx+M_0)-WKx-WK_0 = Ex+E_0+\Delta(Mx+M_0).$$

Moreover, this equality also holds over the integers since $q > \Omega(2^{n^{\delta'}})$ where $\delta' > \delta$, while the absolute value of each entry in $Ex + E_0 + \Delta(Mx + M_0)$ (computed over the integers) is smaller than $O(npq^\alpha + q^\beta + \Delta(np^2 + p^3)) = o(q)$. Hence, q is large enough to ensure that there is no wraparound and the outcome of the first step is $Ex + E_0 + \Delta(Mx + M_0)$ over the integers. Now observe that Δ is large enough to ensure that there is no rounding error in the second step. To see this, note that each entry of $Ex + E_0$ is bounded by $npq^\alpha + q^\beta < \Delta/2$. We conclude that decoding succeeds with probability 1. □

Simulation. Given $(x, z = Mx + M_0)$ we simulate $\hat{g}$ as follows. Choose $W \xleftarrow{R} \mathbb{Z}_q^{m\times\kappa}$, $C \xleftarrow{R} \mathbb{Z}_q^{m\times n}$ and $\hat{K} \xleftarrow{R} \mathbb{Z}_q^\kappa$ and let

$$C_0 = W\hat{K} + E_0 + \Delta z - Cx \pmod q,$$

where $E_0 \xleftarrow{R} \chi_\beta^\kappa$. Output $(W, C, C_0, x, \hat{K})$.

Claim 5.6. *For every M, M_0 and x, the simulated distribution $\mathsf{Sim}(x, Mx+M_0)$ is computationally indistinguishable from the encoding $\hat{g}(M, M_0, x)$.*

Proof. It is not hard to show that, under the LWE assumption, the uniform distribution ($W \xleftarrow{R} \mathbb{Z}_q^{m\times\kappa}, Y \xleftarrow{R} \mathbb{Z}_q^{m\times n}$) is computationally indistinguishable from the following "product"-LWE distribution

$$(W \xleftarrow{R} \mathbb{Z}_q^{m\times\kappa}, Y = WK + E) \text{ where } K \xleftarrow{R} \mathbb{Z}_q^{\kappa\times n}, E \xleftarrow{R} \chi_{q^\alpha}^{m\times n}.$$

This can be proved via a standard hybrid argument, cf. [ACPS09]. We will show that if the claim does not hold then the above distributions can be distinguished.

Fix some (M, M_0, x) and assume, towards a contradiction, that there exists an adversary $\mathcal{A}$ that distinguishes $\mathsf{Sim}(x, Mx + M_0)$ from $\hat{g}(M, M_0, x)$ with advantage ε. We will use $\mathcal{A}$ to distinguish between the product LWE distribution and the uniform distribution with advantage $\varepsilon - \mathrm{neg}(n)$. Given a challenge (W, Y) compute $C = Y + Mx \pmod q$, $\hat{K} \xleftarrow{R} \mathbb{Z}_q^\kappa$ and $C_0 = W\hat{K} + E_0 + \Delta z - Cx \pmod q$ and output $\mathcal{A}(W, C, C_0, x, \hat{K})$.

It is not hard to verify that when the input (W, Y) is uniform in $\mathbb{Z}_q^{m\times\kappa} \times \mathbb{Z}_q^{m\times n}$ the tuple $(W, C, C_0, x, \hat{K})$ is distributed identically to $\mathsf{Sim}(x, Mx + M_0)$. Now assume the input is sampled according to the LWE distribution, i.e., $W \xleftarrow{R} \mathbb{Z}_q^{m\times\kappa}$ and $Y = WK + E$ where $K \xleftarrow{R} \mathbb{Z}_q^{\kappa\times n}$ and $E \xleftarrow{R} \chi_{q^\alpha}^{m\times n}$. We claim that the tuple $(W, C, C_0, x, \hat{K})$ is statistically close to $\hat{g}(M, M_0, x)$.

First observe that the joint distribution of (W, K, E, C) is statistically close in both experiments. (The two are not identical due to the fact that the encoding samples the error matrix E conditioned

on being small – however, this increases the statistical distance by a negligible quantity.) Fix (W, K, E, C) and observe that in both experiments, $\hat{K}$ is uniformly distributed. (Recall that in the encoding $\hat{K} = Kx + K_0$ where $K_0 \xleftarrow{R} \mathbb{Z}_q^\kappa$). Fix K_0 as well. Finally, since $K_0 = \hat{K} - Kx$ the value of C_0 in the encoding can be written as

$$C_0 = WK_0 + E_0 + \Delta \cdot M_0 = W(\hat{K} - Kx) + E_0 + \Delta(M_0 + Mx) - (WK + E + \Delta \cdot M)x + WKx + Ex, \pmod q$$

rearranging and substituting $z = Mx + M_0$ we get

$$C_0 = W\hat{K} + E_0 + Ex + \Delta z - Cx \pmod q.$$

Since each entry of Ex is bounded by $B = npq^\alpha$ and since $E_0 \xleftarrow{R} \chi_{q^\beta}$, the statistical distance between $E_0 + Ex$ and E_0 is at most $m \cdot \mathrm{poly}(B/q^\beta) = \mathrm{neg}(n)$. Hence, C_0 as computed by the algorithm is statistically-close to the distribution of C_0 in the encoding and the claim follows. □

6 More on Online/Offline Encodings

6.1 Some Lower Bounds

Lemma 6.1 (**DARE have super-constant online-rate**). *For every even integer n, there exists a function $f : \{0,1\}^n \to \{0,1\}^{n/2}$ such that any $(t, 1/3)$ decomposable encoding $\hat{f}$ of f has online rate of at least $\min(\log(t/s)/2, n/4)$ where s is the circuit size of the decoder.*

Specifically, super-polynomial security implies lower-bound of $\omega(\log n)$, and sub-exponential security 2^{n^ε} implies polynomial rate of n^ε which matches the construction of [AIK06]. We also mention that the proof actually holds for a stronger statement as it rules out even an extremely poor distinguishing advantage ε that tends to 1 exponentially fast (e.g., $\varepsilon = 1 - 2^{-n/3}$ cannot be achieved with rate smaller than $\min(\log(t/s)/2, n/12)$).

We note that the dependence of the rate bound on the complexity of the decoder seems inherent. Indeed, if the decoder is more powerful than the distinguisher, then the encoder can just encrypt the input using an encryption scheme which the decoder can break but the distinguisher cannot. Such an encryption scheme can be of the form $E(x) = (C(r), x + r)$ where r is a random mask of the same length as x and C is a perfectly binding non-interactive commitment scheme. With a sufficiently powerful decoder, E can be used as a DRE with online rate 1 for an arbitrarily f.

We now prove Lemma 6.1.

Proof. Let $f(x, y) = (x_1 + x_2 + \cdots + x_{n/2}) \cdot y$ where $x, y \in \mathbb{F}_2^{n/2}$ and multiplication is understood as a multiplication of a vector y by the scalar $\sum x_i$ over $\mathbb{F}_2$. Let $\hat{f}$ be a $(t, \frac{1}{2})$ decomposable encoding of f with decoding complexity of s, and assume, towards a contradiction, that $\hat{f}$ has online complexity of kn where $k < \min(\log(t/s)/2, n/4)$. Since $\hat{f}$ is decomposable, there must exist an input x_i that affects a set S of at most $2k$ outputs. We will exploit this property to break the security of $\hat{f}$.

Consider a uniformly chosen y and $x = 0^n$. We show how to distinguish the distribution $\hat{f}((x, y); r)$ from $\mathsf{Sim}(f(x, y))$. Given z (sampled from one of the above distributions), enumerate all 2^{2k} strings z' which differ from z only with respect to the indices which are influenced by x_i, and apply the decoder Dec to each of the modified strings. If the outcome in at least one of these

tests corresponds to y the distinguisher outputs "pass" and otherwise it outputs "fail". (Recall that distinguishing should be hard even if the inputs x, y are known.)

We claim that when z is the outcome of the simulator the test passes with negligible probability. Indeed, the input to the simulator is independent of y (the simulator gets $f(x,y)=0$), and therefore the probability of the above process correctly guessing y in 2^{2k} attempts is bounded by $2^{2k}/2^{n/2}$ which, by our assumption on k, is upper bounded by $\frac{1}{2}$.

On the other hand, if z is the outcome of $\hat{f}((x,y);r)$ then: (1) the string $z' = \hat{f}((e_i,y);r)$, where e_i is the i-th unit vector, differs from z only on (subset of) the coordinates which are influenced by x_i; and (2) $\mathsf{Dec}(z') = f(x') = y$ and so the test passes with probability 1. Since the complexity of the test is $2^{2k} \cdot s < t$ and since it has a distinguishing advantage of $1 - 2^{2k-n/2} \geq \frac{1}{2}$ the lemma follows. □

We now prove a lower bound on the online rate of AREs. It is clear that the online rate needs to be at least 1 for functions with a long output (such as the identity function). We show that this is also the case for some boolean functions.

Lemma 6.2 (**ARE have online rate of at least 1**). *There exists a boolean function $f : \{0,1\}^n \to \{0,1\}$ such that the online rate of any ARE of f (over an arbitrary ring) is at least $1 - o(1)$.*

Proof. Let $n = k + \log k$. We will show that the "universal" function $f(x,i) = x_i$, where x is in $\{0,1\}^k$ and i is in $\{0,1\}^{\log k}$, cannot be encoded by an ARE $\hat{f}(x,i;r)$ with online communication complexity smaller than k. Indeed, fix some randomness r and let $\hat{f}(x,i;r) = (F, A_r \cdot x + A'_r \cdot i + b)$ be a perfectly correct ARE of f over $\mathbb{Z}_q$ where F is the offline part. Then it is possible to fully recover x from $A_r x$ (by computing $A'_r i + b$ for $i = 1, \ldots, n$), from which it follows that the bit length of $A_r x$ is at least k. □

While we cannot unconditionally prove a similar result for non-affine REs with, say, quadratic online computation, such a negative result follows from the (conjectured) impossibility of compression. Formally, we say that a function f is *compressed* by a function g if: (1) (lossless recovery) there exists a recovery function h such that for every input x, $h(g(x)) = f(x)$; and (2) (g is shrinking) for every x the length of $g(x)$ is shorter than the length of x. We conjecture that for every positive constants $c > c'$, there exists a function computable in time n^c that cannot be compressed by a time-$(n^{c'})$ function g. (See [HN06, DI06] for related conjectures.)

Lemma 6.3 (**Incompressibility ⇒ RE have online rate of at least 1**). *Suppose that the above incompressibility conjecture holds. Then, the class of efficiently computable functions does not have an online efficient encoding with online rate smaller than* 1.

Proof. Assume, towards a contradiction, that there exists a compiler C that for every polynomial-time computable function f outputs an encoding $\hat{f}$ whose online communication is smaller than n and its online computational complexity is $n^{c'}$ for some universal constant c'. Then, any efficiently computable function f can be compressed by the $n^{c'}$-time computable function $\hat{f}_{\mathsf{on}}(x;r)$ where r is some fixed string, e.g., the all-zero string. Indeed, it is possible to (efficiently) recover the value of $f(x)$ by computing $z = \hat{f}_{\mathsf{off}}(r)$ and applying the decoder to $(z, \hat{f}_{\mathsf{on}}(x;r))$. This contradicts the incompressibility assumption. □

In the case of functions with many outputs (and assuming the existence of one-way functions) we can lower-bound the *offline* complexity by the output length ℓ where ℓ can be polynomially

larger than n. To this end, we will prove that the output length of the simulator (or even the online-complexity of the simulator) is lower-bounded by ℓ.

Lemma 6.4 (**Communication complexity of RE is larger than the output length**). *Assuming one-way functions, for every constant c there exists a function $f : \{0,1\}^n \to \{0,1\}^{n^c}$ such that every $(n^{\omega(1)}, 1/3)$-private RE of f has communication complexity of at least n^c bits. Furthermore, there are at least n^c bits in the output of the simulator $\mathsf{Sim}(y)$ that depend on the input y (as opposed to the randomness).*

Note that the existence of one-way functions is necessary, as otherwise one can obtain an encoding with total complexity n, by letting $\hat{f}(x;r) = x'$ where x' is random sibling of x under f, and take the decoder to be $\mathsf{Dec}(x') = f(x')$, and the simulator $\mathsf{Sim}(y) = x'$ where x' is a random preimage of y. If one-way functions do not exist then this encoding can be implemented efficiently.

Proof. Fix some constant c, and let $f : \{0,1\}^n \to \{0,1\}^\ell$ be a pseudorandom generator with output length $\ell = n^c$. (The existence of such a pseudorandom generator follows from the existence of one-way functions [HILL99].) It suffices to prove the "furthermore" part as the online-complexity of the simulator lower-bounds the communication complexity of the encoding. Let $\hat{f}(x;r)$ be an RE of f with decoder Dec and simulator Sim such that the number of bits of $\mathsf{Sim}(y)$ that depend on y is smaller than ℓ. Then, we distinguish the output of f from a truly random string via the following test: Given a string $y \in \{0,1\}^\ell$, we accept if and only if the outcome of $\mathsf{Dec}(\mathsf{Sim}(y))$ is equal to y.

First we claim that when y is random the test accepts with probability at most $\frac{1}{2}$. Indeed, fix some value r for the randomness of the simulator and some value d for the randomness of the decoder. Then the image of $\mathsf{Sim}(y;r) = (z_r, \mathsf{Sim}_{\mathsf{on}}(y;r))$ can take at most $2^{\ell-1}$ values, and therefore the decoder $\mathsf{Dec}(\cdot;s)$ recovers y successfully for at most half of all y's in $\{0,1\}^\ell$.

On the other hand, if y is in the image of f, the test accepts with probability at least $2/3 - \mathrm{neg}(n)$. Indeed, let x be a preimage of y, then by definition $\mathsf{Dec}(\hat{f}(x;r))$ outputs $y = f(x)$ with probability 1. Since $\hat{f}(x;r)$ is $(t, 1/3)$ indistinguishable from $\mathsf{Sim}(f(x))$, it follows that $\mathsf{Dec}(\mathsf{Sim}(y)) = y$ with probability at least $2/3 - \mathrm{neg}(n)$. □

6.2 On Adaptive Security

The standard security definition of REs can be captured by the following game: (1) The challenger secretly tosses a random coin $b \xleftarrow{R} \{0,1\}$; (2) the adversary chooses an input x submits it to the challenger and gets as a result the string $\hat{y}$ which, based on the secret bit b, is either sampled from the encoding $\hat{f}(x;r)$ or from the simulator $\mathsf{Sim}(f(x))$. At the end, the adversary outputs his guess b' for the bit b. The security of REs says that the t-bounded adversaries cannot win the game (guess b) with probability better than $\frac{1}{2} + \varepsilon$. In the online/offline setting it is natural to consider an *adaptive* version of this game in which the adversary chooses its input x based on the offline part of the encoding. Syntactically, this requires an online/offline simulator $\mathsf{Sim}(y;r) = (\mathsf{Sim}_{\mathsf{off}}(r); \mathsf{Sim}_{\mathsf{on}}(x;r))$ whose offline part does not depend on its input $f(x)$, and has the same length as the offline part of the encoding. Formally,

Definition 6.5 (**Adaptively-secure RE**). *Let f be a function and $\hat{f}(x;r) = (\hat{f}_{\mathsf{off}}(r), \hat{f}_{\mathsf{on}}(x;r))$ be a perfectly-correct RE with decoder Dec and online/offline simulator $\mathsf{Sim}(y;r) = (\mathsf{Sim}_{\mathsf{off}}(r), \mathsf{Sim}_{\mathsf{on}}(y;r))$. We say that $\hat{f}$ is* (t,ε) adaptively private *if every t-bounded adversary $\mathcal{A}$ wins the following game*

with probability at most $\frac{1}{2}+\varepsilon$: (1) The challenger secretly tosses a random coin $b \xleftarrow{R} \{0,1\}$, chooses randomness r and outputs

$$\hat{y}_{\mathsf{off}} = \begin{cases} \hat{f}_{\mathsf{off}}(r) & \text{if } b = 1, \\ \mathsf{Sim}_{\mathsf{off}}(r) & \text{if } b = 0. \end{cases}$$

(2) Based on $\hat{y}_{\mathsf{off}}$ the adversary $\mathcal{A}$ chooses an input x, submits it to the challenger and gets as a result the string

$$\hat{y}_{\mathsf{on}} = \begin{cases} \hat{f}_{\mathsf{on}}(x;r) & \text{if } b = 1, \\ \mathsf{Sim}_{\mathsf{on}}(f(x);r) & \text{if } b = 0. \end{cases}$$

At the end, the adversary outputs his guess b' and wins if $b' = b$.

As follows from Lemma 6.4, in the standard model adaptively secure REs cannot be online-efficient let alone have constant rate (assuming the existence of one-way functions).

Corollary 6.6 (**Adaptive security requires long online-communication**). *Assuming one-way function, for every constant c there exists a function $f : \{0,1\}^n \to \{0,1\}^{n^c}$ such that every RE of f has online communication complexity of at least n^c bits.*

Proof. By Lemma 6.4, there exists a function f for which the online part of the simulator must longer than n^c. Privacy ensures that the online communication complexity of $\hat{f}$ satisfies the same bound. □

On the other hand, it turns out that this barrier can be bypassed via the use of a (programmable) random oracle as shown in the following lemma.

Lemma 6.7. *Suppose that $f : \{0,1\}^n \to \{0,1\}^\ell$ has a (t,ε)-private RE $\hat{f}$ with online communication αn. Then, for every $\beta > 0$, f has a $(t' = \Omega(t), \varepsilon' = \varepsilon + t/2^{\beta n})$ adaptively-private RE g with online communication of $(\alpha+\beta)n$ in the random oracle model. Furthermore, if $\hat{f}$ is affine then so is g.*

Proof. Let $\hat{f}(x;r) = (\hat{f}_{\mathsf{on}}(x;r), \hat{f}_{\mathsf{off}}(r))$ be a (t,ε) (affine) RE of $f(x)$ over $\mathbb{Z}_q$ with decoder Dec, and simulator $\mathsf{Sim} = (\mathsf{Sim}_1, \mathsf{Sim}_2)$, where Sim_1 denotes the first αn bits of Sim and $\mathsf{Sim}_2, \hat{f}_{\mathsf{off}}(r)$ output s-long vectors in $\mathbb{Z}_q$. Let $H : \{0,1\}^{\beta n} \to \mathbb{Z}_q^s$ be a random oracle. Consider the encoding

$$g_{\mathsf{off}}(r,k) := \hat{f}_{\mathsf{off}}(r) + H(k) \pmod q \qquad\qquad g_{\mathsf{on}}(x;r,k) := (\hat{f}_{\mathsf{on}}(x;r), k)$$

where $k \xleftarrow{R} \{0,1\}^{\beta n}$. Given $g(x;(r,k)) = (\hat{f}_{\mathsf{off}}(r) + H(k), \hat{f}_{\mathsf{on}}(x;r), k)$, the decoder Dec' decodes $\hat{f}_{\mathsf{off}}(r)$ (by subtracting $H(k)$ from the first entry) and then applies the original decoder Dec to $\hat{f}(x;r)$. The simulator Sim' works as follows: at the offline phase it outputs $\rho \xleftarrow{R} \mathbb{Z}_q^s$, then at the online phase given y it outputs the pair $(\mathsf{Sim}_1(y;r), k)$, where $k \xleftarrow{R} \{0,1\}^{\beta n}$ and programs the random oracle so that $H(k) = \rho - \mathsf{Sim}_2(y;r) \pmod q$.

We claim that the resulting encoding is (t',ε') adaptively private. Assume, towards a contradiction, that we have a t'-bounded adversary $\mathcal{B}$ that wins the adaptive-RE game with probability $\frac{1}{2}+\varepsilon'$. Then, we can construct a t-bounded adversary that distinguishes $\hat{f}(x;r)$ from $\mathsf{Sim}(f(x))$ with advantage ε as follows: Choose a random string $z_1 \xleftarrow{R} \mathbb{Z}_q^s$ for the offline phase and send it to $\mathcal{B}$, let x be its response. If $\mathcal{B}$ makes a query to the random oracle, we record the query and answer

it with a random string. (Without loss of generality, $\mathcal{B}$ never asks the same query twice.) Now, we try to break the encoding $\hat{f}$ with respect to the input x. As a result, we get $\hat{y} = (\hat{y}_{\mathsf{on}}, \hat{y}_{\mathsf{off}})$ which is either sampled from $\hat{f}(x)$ or from $\mathsf{Sim}(f(x))$. Send $\mathcal{B}$ the string $z_2 = (\hat{y}_{\mathsf{on}}, k)$ where $k \stackrel{R}{\leftarrow} \{0,1\}^{\beta n}$, and set $H(k) = \rho - \hat{y}_{\mathsf{off}}$. If k happens to be a query that $\mathcal{B}$ already asked, we terminate with failure. Otherwise, we continue the emulation while answering queries k' to the random oracle randomly (if $k' \neq k$) or by $H(k)$ if $k' = k$.

To analyze the success probability of $\mathcal{A}$ first observe that the emulation fails with probability at most $t/2^{\beta n}$. Furthermore, assuming non-failure, the emulation is perfect in the sense that if $\hat{y}$ is sampled from the encoding $\hat{f}(x; r)$ (resp., from the simulator $\mathsf{Sim}(f(x))$) then the view of $\mathcal{B}$ is distributed identically to its view in the adaptive RE game when the challenge bit $b = 1$ (resp. $b = 0$). Hence, the lemma follows. □

7 Applications

7.1 MPC with Optimal Online Communication

In this section, we describe the application of succinct randomized encodings to secure multiparty computation (MPC) in the preprocessing model. We start with the two-party case, and later generalize to the multiparty case. For concreteness, we focus on distributing the DDH-based encoding obtained by combining Lemmas 3.2 and 4.1 with the DDH-based AHE. Similar protocols can be obtained based on any succinct *affine* RE. We do not know how to get similar results from general (non-affine) succinct REs.

Let F be a deterministic two-party functionality which takes an input $a \in \{0,1\}^{n_a}$ from Alice and an input $b \in \{0,1\}^{n_b}$ from Bob, and delivers an output c to Alice.[6] The DDH-based encoding of F can be written as

$$\hat{F}(a, b; R) = (\hat{F}_{\mathsf{off}}(R),\ a \oplus r^a,\ b \oplus r^b,\ \sum_{i=1}^{n_a} K^A_{i,a_i \oplus r^a_i} + \sum_{i=1}^{n_b} K^B_{i,b_i \oplus r^b_i} \mod p), \tag{2}$$

where the "masks" $r^a \in \{0,1\}^{n_a}$, $r^b \in \{0,1\}^{n_b}$, and the "keys" $K^A_{i,\sigma}, K^B_{i,\sigma} \in \mathbb{Z}_p$ are random and independent of a, b (these values are given as part of R).

In the semi-honest model, the protocol is straightforward. In the offline phase, a trusted party samples R and sends the value $\hat{F}_{\mathsf{off}}(R)$ together with the mask r^a to Alice, and the mask r^b along with the $2n_a + 2n_b$ keys $K^A_{i,\sigma}, K^B_{i,\sigma}$ to Bob. (In the real world, this step can be implemented via the use of any off-the-shelf secure two-party protocol.) In the online phase, Alice sends to Bob $a \oplus r^a$ and Bob replies with $b \oplus r^b$ and $\sum_{i=1}^{n_a} K^A_{i,a_i \oplus r^a_i} + \sum_{i=1}^{n_b} K^B_{i,b_i \oplus r^b_i} \mod p$. Alice computes the output using the decoder of $\hat{F}$. Note that the view of Bob is completely random, whereas the view of Alice contains the output of $\hat{F}$ which can be simulated given $F(a, b)$. This proves the following:

Theorem 7.1. *Suppose that the DDH assumption holds in a prime order group of size $p = p(\kappa)$. Let $F(a, b)$ be a polynomial-time computable functionality which delivers its output to Alice. Assume trusted preprocessing which does not depend on the inputs. Then, F can be securely realized in the semi-honest model by a protocol in which Alice sends a message of length $|a|$ and Bob sends a message of length $|b| + \lceil \log p \rceil$, independently of the length of the output or the complexity of F.*

[6] The case of general two-party functionalities reduces to this case via a standard reduction, cf. [Gol04].

The malicious model. Security in the malicious model is handled via a "homomorphic MAC" over $\mathbb{Z}_p$ which allows Alice to verify that Bob sent her the correct linear combination of his keys, namely the one defined by $a \oplus r^a$ and $b \oplus r^b$. This approach has been used by Bendlin et al. for securely computing arithmetic circuits in the preprocessing model [BDOZ11].[7] Note that there is no room for cheating in sending $a \oplus r^a, b \oplus r^b$: Any choice of these messages uniquely defines the input, which can be easily extracted by the simulator.

The homomorphic MAC construction proceeds as follows. Suppose for simplicity that in the offline phase Bob is given n secret keys $K_1, \ldots, K_n \in \mathbb{Z}_p$ and in the online phase he is expected to reveal some publicly known linear combination $\sum_{i=1}^n \mu_i K_i$ of these keys. The goal is to provide Alice with a mechanism for checking the correctness of Bob's online message without revealing additional information about the keys. This is done by giving to Bob, in the offline phase, independent random field elements $r_1, \ldots, r_n \in \mathbb{Z}_p$, and giving to Alice a random $\alpha \in \mathbb{Z}_p$ along with the n values $\beta_i = \alpha K_i + r_i$. Note that Alice's offline information gives her no information about Bob's keys. In the online phase, when the coefficients μ_i are revealed (in our case $\mu_i \in \{0, 1\}$), Bob sends the values $(M = \sum_{i=1}^n \mu_i K_i, T = \sum_{i=1}^n \mu_i r_i)$. Alice accepts M only if $\alpha M + T = \sum_{i=1}^n \mu_i \beta_i$. Since Alice can compute T from M and α, she does not learn anything except the output. On the other hand, a forging attack by a malicious Bob can be used to guess α. See [BDOZ11] for more details.

The multiparty case. In the case of a general number of parties we can proceed similarly, except that in this case the offline phase additively secret-shares each key between the parties over $\mathbb{Z}_p$. (This is needed in order to prevent the adversary from learning both the output of the encoding and the keys.) As before, assume without loss of generality that only one party Alice gets the output. The protocol proceeds in two rounds: in the first round each party broadcasts its masked input, and in the second round each party sums up and sends to Alice his shares of the keys defined by the public messages of the first round. Each of these messages is verified by Alice using the homomorphic MAC, as before.

On an adaptive choice of inputs. In the presence of malicious parties, the protocol described above realizes the randomized encoding functionality $\hat{f}$ with *statistical* security in the preprocessing model. However, this functionality should be viewed as a reactive functionality which first delivers an offline part, then receives an online input from each party, and finally delivers the online output to Alice. In order for the final protocol to be secure, the encoding should be simulatable with such an adaptive choice of inputs. While we do not have a proof for the adaptive security of our constructions under standard assumptions, it may still hold heuristically when the output for f is shorter than the input, and can be made provably adaptive in the random oracle model (see Section 6.2 and [BHR12a]). As in the case of standard garbled circuits with short keys, obtaining adaptive security in the plain model under standard assumptions remains an interesting open problem. We note that this issue is not relevant in the semi-honest model, or when the online inputs are public and are generated independently of the offline phase (this is meaningful when there are secret offline inputs).

[7] In contrast to our protocols, the online communication complexity of the protocol from [BDOZ11] depends on the circuit size.

Private simultaneous messages protocols. In the non-interactive model for secure computation from [FKN94], there are k parties $P_1, \ldots, P_k$, where party P_i holds an input x_i and all parties have access to a common random string r. To evaluate a function f on their inputs, the parties simultaneously send messages to an external referee who has no access to r. From these messages the referee should be able to recover $f(x_1, \ldots, x_k)$. On the other hand, the messages should give no additional information about the inputs. In the computational security model, this is formalized by requiring that the referee's view can be efficiently simulated given the output. It is easy to see that our succinct AREs yield protocols in this model with nearly optimal online communication. For instance, in two-party case, the DDH-based encoding of Eq. (2) can be used to obtain a protocol in which P_1 on input a sends the message $(a \oplus r^a, \sum_{i=1}^{n_a} K^A_{i,a_i \oplus r^a_i} + z \mod p)$, where z is random in $\mathbb{Z}_p$, and P_2 on input b sends the message $(b \oplus r^b, \sum_{i=1}^{n_b} K^B_{i,b_i \oplus r^b_i} - z \mod p)$. Note that the information revealed by the second entries of these two messages is equivalent to the last entry in the right hand side of Eq. (2). In addition, this protocol requires an offline message $\hat{F}_{\mathsf{off}}(R)$ that can be sent by one of the parties before the inputs are known. This protocol can be easily extended to the case of a larger number of parties k. In the general case, the online communication of each party only includes its masked input together with a single group element.

7.2 Non-Interactive Zero-Knowledge Proofs

We move on to the case of non-interactive zero-knowledge proofs (NIZK). Such proof systems are similar to standard zero-knowledge protocols except that interaction is traded for the use of a public random string σ to which both the prover and the verifier have a read-only access. Formally,

Definition 7.2. *A NIZK for an* **NP** *relation $R(x, w)$ is a pair of probabilistic polynomial-time algorithms (P, V) that satisfies the following properties:*

- *(Completeness) for every $(x, w) \in R$, it holds that $\Pr[V(x, \sigma, P(x, w, \sigma; \rho_P); \rho_V) = 1] > 1 - \mathrm{neg}(|x|)$, where ρ_P is the private randomness of the prover and ρ_V is the private randomness of the verifier.*
- *(Statistical Soundness) for every $x \notin L_R$ (i.e., x such that $\forall w, (x, w) \notin R$) and every computationally unbounded prover algorithm P^* we have that $\Pr[V(x, \sigma, P^*(x, \sigma); \rho_V) = 1] < \mathrm{neg}(|x|)$;*
- *(Zero-knowledge) there exists a probabilistic polynomial-time simulator M such that for every string sequence $\{(x_n, w_n)\}$ where $(x_n, w_n) \in R$ it holds that*

$$\{(x_n, \sigma, P(x_n, w_n, \sigma; \rho_P))\} \stackrel{c}{\equiv} \{M(x_n)\},$$

where in all the above σ is uniformly distributed over $\{0,1\}^{\mathrm{poly}(|x|)}$.

In the online/offline setting we assume that the prover's message is partitioned into two parts: (1) An offline message that depends solely on the language L and the public reference string σ; and (2) An online message that depends on the public input x and the private witness w. (We also assume that the verifier gets x in the online phase.) Accordingly, for a prover $P(x, w, \sigma; \rho_P) = (P_{\mathsf{off}}(\sigma, \rho_P), P_{\mathsf{on}}(x, w, \rho_P))$ we define the online communication (resp. computational) complexity of P to be the output length (resp., the circuit size) of $P_{\mathsf{on}}(x, w, \rho_P)$.

Theorem 7.3. *Assume that the language L has a NIZK and assume that RSA, LWE, or DDH holds. Then, L has a NIZK $(\hat{P}, \hat{V})$ with online communication of $|w| + o(|x| + |w|)$ and online computation of $(|x| + |w|)^{1+\varepsilon}$ where ε is an arbitrary small constant.*

Proof. Let (P, V) be a NIZK for L. In [AIK04, AIK06] it is shown that if $P(x, w, \sigma, \rho_P)$ is encoded by $\hat{P}(x, w, \sigma, \rho_P; r)$ then $(\hat{P}, \hat{V})$ is a NIZK for L where the new verifier $\hat{V} = V(x, \sigma, \mathsf{Dec}(\hat{y}); \rho_V)$ uses the decoder Dec to translate the prover's encoded message $\hat{y}$ to the corresponding message of the original prover, and then invokes the original verifier.

To prove the theorem we compile the prover $P(x, w, \sigma, \rho_P)$ into its succinct randomized encoding $\hat{P}(x, w, \sigma, \rho_P; r_x, r_w)$ constructed in Section 4 while treating (σ, ρ_P) as offline inputs and (x, w) as online inputs. As a result the online computation is $(|x| + |w|)^{1+\varepsilon}$ and the online communication is $|x| + |w| + o(|x| + |w|)$. To further reduce the online communication observe that the online part has the form $(x \oplus r_x, w \oplus r_w, K(x, w; r))$, since x is public this is information theoretically equivalent to sending the message $(r_x, w \oplus r_w, K(x, w; r))$, however, the randomness r_x can be sent at the offline phase and so the total online communication is $|w| + o(|x| + |w|)$ as needed. □

Remark 7.4 (Adaptivity)**.** *The zero-knowledge property is proven under the assumption that the online input x is chosen independently of the offline part. When x is chosen based on the offline part, we do not know how to efficiently simulate the view of the verifier under standard assumptions. Still, security in this case may hold heuristically, and can be provably achieved in the random oracle model (see Section 6.2). It is important to note that this issue is not relevant to the* soundness *of the protocol (which holds statistically). Furthermore, the problem is completely avoided when the online input is generated independently of the offline phase (e.g., by the prover).*

7.3 Verifiable Computation

In the problem of Verifiable Computation (VC) a computationally weak client C wishes to delegate the computation of a function f on an input x to a computationally strong but untrusted server P. We consider two-message protocols in the offline/online setting. Namely, the client sends an offline message $\alpha = C_{\mathsf{off}}(\rho_C)$ before seeing the input x, then at the online phase the client sends a single message $\beta = C_{\mathsf{on}}(x, \rho_C)$ to the server (which potentially reveals x). The server responds with an answer $\gamma = P(\alpha, \beta, \rho_P)$. Based on this answer, the client applies some cheap verification process $V(x, \gamma, \rho_V)$ and either recovers the result $f(x)$ or announces an error in the case of a cheating server. Formally,

Definition 7.5. (Verifiable Computation) *A VC protocol for an efficiently computable function $f : \{0,1\}^n \to \{0,1\}^m$ is a pair of probabilistic polynomial-time algorithms, a client $C = (C_{\mathsf{on}}, C_{\mathsf{off}}, V)$ and a server P, that satisfies the following properties:*

- *(Completeness) for every $x \in \{0,1\}^n$, it holds with probability 1 that*

$$V\left(x, P(C_{\mathsf{off}}(\rho_C), C_{\mathsf{on}}(x, \rho_C); \rho_P); \rho_C\right) = f(x),$$

  where ρ_P is the private randomness of the server and ρ_C is the private randomness of the client.

- *(Soundness) for every x and every efficiently computable server P^* we have that*

$$\Pr[V\left(x, P^*(C_{\mathsf{off}}(\rho_C), C_{\mathsf{on}}(x, \rho_C); \rho_P); \rho_C\right) \notin \{f(x), \perp\}] < \mathrm{neg}(|x|).$$

We will be interested in *useful* protocols in which it is more efficient to run the client's algorithm than to compute the function itself.

Theorem 7.6. *Assume that RSA, LWE, or DDH holds and let $\varepsilon > 0$ be an arbitrary small constant. For every efficiently computable function $f : \{0,1\}^n \to \{0,1\}^m$ there exists an online/offline VC $C = (C_{\mathsf{on}}, C_{\mathsf{off}}, V)$ with the following complexity:*

- *Offline communication/computational complexity of the client is $|f| \cdot \kappa$, where $|f|$ denotes the circuit size of f and κ is a security parameter.*
- *Online communication complexity of the client is $n + o(n)$ and its online computational complexity is $n^{1+\varepsilon}$.*
- *The communication complexity of the server is $m + n^{\varepsilon}$ and its computational complexity is $|f| \cdot n^{\varepsilon}$.*
- *The verification step V has computational complexity of $O(m + n^{\varepsilon})$.*

Observe that the online communication of the protocol is essentially optimal (up to additive loss) as even if the server is fully trusted the client has to send at least n-bits to describe x and the server has to send at least m bits to describe $f(x)$.

Proof. Let $\kappa = n^{\varepsilon}$. In [AIK10] it is shown that the following protocol is VC. Let $g(x,k) = \mathtt{MAC}_k(f(x))$ where $\mathtt{MAC}_k : \{0,1\}^m \to \{0,1\}^{\kappa}$ is a one-time (information-theoretic) MAC with an error of $2^{-\kappa}$. Let $\hat{g}(x,k;r)$ be the computationally-private perfectly-correct RE for g where k is treated as an offline input. Let $\hat{g}_{\mathsf{off}}(k;r)$ be the offline message of the client, and let $(x, \hat{g}_{\mathsf{on}}(x;r))$ be the online message of the client. The server P sends the pair $\gamma_1 = f(x)$ and $\gamma_2 = \mathsf{Dec}((\hat{g}_{\mathsf{off}}(k;r), \hat{g}_{\mathsf{on}}(x;r)))$ where Dec is the decoder of the encoding. Finally, the client accepts γ_1 if $\gamma_2 = \mathtt{MAC}_k(\gamma_1)$.

To prove the theorem we employ the encoding constructed in Section 4 and instantiate the MAC with the pair-wise independent hash function from [IKOS08] whose circuit complexity is $O(m + \kappa)$. □

We remark that one can add input privacy (i.e., hide x from the server) without increasing the complexity (see [AIK10]). Also, as in [GGP10], the offline phase can be re-used (and therefore amortize) without increasing the (asymptotic) complexity by encrypting $\hat{g}_{\mathsf{off}}(k;r)$ and $\hat{g}_{\mathsf{on}}(x;r)$ under fully-homomorphic encryption and letting the server return an encryption of γ_2. Re-using the offline phase remain secure as long as the server does not learn whether the client accepted or rejected the interaction.

Remark 7.7 (Adaptivity)**.** *We do not prove that the soundness property holds when the input x is chosen adaptively based on the offline part. As usual, this can be solved with the aid of a random oracle (see Section 6.2). It is important to note that in this context non-adaptive solution is still meaningful as in the typical scenario, where the client is the one who selects which input x to delegate, the problem is completely avoided.*

Acknowledgements. The first author was supported by Alon Fellowship, ISF grant 1155/11, Israel Ministry of Science and Technology (grant 3-9094), and GIF grant 1152/2011. The second author was supported by the European Research Council as part of the ERC project CaC (grant 259426). The third author was supported by ISF grant 1361/10 and BSF grant 2008411. The fourth author was supported by NSF grants CNS-0915361 and CNS-0952692, AFOSR Grant No: FA9550-08-1-0352, DARPA through the U.S. Office of Naval Research under Contract N00014-11-1-0382, DARPA N11AP20006, Google Faculty Research award, the Alfred P. Sloan Fellowship, and Microsoft Faculty Fellowship, and Packard Foundation Fellowship. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Department of Defense or the U.S. Government.

References

[ACPS09] Benny Applebaum, David Cash, Chris Peikert, and Amit Sahai. Fast cryptographic primitives and circular-secure encryption based on hard learning problems. In Shai Halevi, editor, *CRYPTO 2009*, volume 5677 of *LNCS*, pages 595–618. Springer, August 2009.

[AHI11] Benny Applebaum, Danny Harnik, and Yuval Ishai. Semantic security under related-key attacks and applications. In Bernard Chazelle, editor, *2nd ICS*, pages 45–60. Tsinghua University Press, January 2011.

[AIK04] Benny Applebaum, Yuval Ishai, and Eyal Kushilevitz. Cryptography in NC^0. In *45th FOCS*, pages 166–175. IEEE Computer Society Press, October 2004.

[AIK06] Benny Applebaum, Yuval Ishai, and Eyal Kushilevitz. Computationally private randomizing polynomials and their applications. *Computational Complexity*, 15(2):115–162, 2006.

[AIK10] Benny Applebaum, Yuval Ishai, and Eyal Kushilevitz. From secrecy to soundness: Efficient verification via secure computation. In Samson Abramsky, Cyril Gavoille, Claude Kirchner, Friedhelm Meyer auf der Heide, and Paul G. Spirakis, editors, *ICALP 2010, Part I*, volume 6198 of *LNCS*, pages 152–163. Springer, July 2010.

[AIK11] Benny Applebaum, Yuval Ishai, and Eyal Kushilevitz. How to garble arithmetic circuits. In Rafail Ostrovsky, editor, *52nd FOCS*, pages 120–129. IEEE Computer Society Press, October 2011.

[AIKW12] Benny Applebaum, Yuval Ishai, Eyal Kushilevitz, and Brent Waters. Encoding functions with constant online rate or how to compress garbled circuits keys. Cryptology ePrint Archive, Report 2012/693, 2012. `http://eprint.iacr.org/2012/693`.

[AIKW13] Benny Applebaum, Yuval Ishai, Eyal Kushilevitz, and Brent Waters. Encoding functions with constant online rate or how to compress garbled circuits keys. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part II*, volume 8043 of *LNCS*, pages 166–184. Springer, August 2013.

[App11] Benny Applebaum. Key-dependent message security: Generic amplification and completeness. In Kenneth G. Paterson, editor, *EUROCRYPT 2011*, volume 6632 of *LNCS*, pages 527–546. Springer, May 2011.

[BDOZ11] Rikke Bendlin, Ivan Damgård, Claudio Orlandi, and Sarah Zakarias. Semi-homomorphic encryption and multiparty computation. In Kenneth G. Paterson, editor, *EUROCRYPT 2011*, volume 6632 of *LNCS*, pages 169–188. Springer, May 2011.

[Bea95] Donald Beaver. Precomputing oblivious transfer. In Don Coppersmith, editor, *CRYPTO'95*, volume 963 of *LNCS*, pages 97–109. Springer, August 1995.

[BGG+14] Dan Boneh, Craig Gentry, Sergey Gorbunov, Shai Halevi, Valeria Nikolaenko, Gil Segev, Vinod Vaikuntanathan, and Dhinakaran Vinayagamurthy. Fully key-homomorphic encryption, arithmetic circuit ABE and compact garbled circuits. In Phong Q. Nguyen and Elisabeth Oswald, editors, *EUROCRYPT 2014*, volume 8441 of *LNCS*, pages 533–556. Springer, May 2014.

[BHHI10] Boaz Barak, Iftach Haitner, Dennis Hofheinz, and Yuval Ishai. Bounded key-dependent message security. In Henri Gilbert, editor, *EUROCRYPT 2010*, volume 6110 of *LNCS*, pages 423–444. Springer, May 2010.

[BHR12a] Mihir Bellare, Viet Tung Hoang, and Phillip Rogaway. Adaptively secure garbling with applications to one-time programs and secure outsourcing. In Xiaoyun Wang and Kazue Sako, editors, *ASIACRYPT 2012*, volume 7658 of *LNCS*, pages 134–153. Springer, December 2012.

[BHR12b] Mihir Bellare, Viet Tung Hoang, and Phillip Rogaway. Foundations of garbled circuits. In Ting Yu, George Danezis, and Virgil D. Gligor, editors, *ACM CCS 12*, pages 784–796. ACM Press, October 2012.

[BSW11] Dan Boneh, Amit Sahai, and Brent Waters. Functional encryption: Definitions and challenges. In Yuval Ishai, editor, *TCC 2011*, volume 6597 of *LNCS*, pages 253–273. Springer, March 2011.

[CCKM00] Christian Cachin, Jan Camenisch, Joe Kilian, and Joy Müller. One-round secure computation and secure autonomous mobile agents. In Ugo Montanari, José D. P. Rolim, and Emo Welzl, editors, *ICALP 2000*, volume 1853 of *LNCS*, pages 512–523. Springer, July 2000.

[CFIK03] Ronald Cramer, Serge Fehr, Yuval Ishai, and Eyal Kushilevitz. Efficient multi-party computation over rings. In Eli Biham, editor, *EUROCRYPT 2003*, volume 2656 of *LNCS*, pages 596–613. Springer, May 2003.

[CKV10] Kai-Min Chung, Yael Kalai, and Salil P. Vadhan. Improved delegation of computation using fully homomorphic encryption. In Tal Rabin, editor, *CRYPTO 2010*, volume 6223 of *LNCS*, pages 483–501. Springer, August 2010.

[DI06] Ivan Damgård and Yuval Ishai. Scalable secure multiparty computation. In Cynthia Dwork, editor, *CRYPTO 2006*, volume 4117 of *LNCS*, pages 501–520. Springer, August 2006.

[DPSZ12] Ivan Damgård, Valerio Pastro, Nigel P. Smart, and Sarah Zakarias. Multiparty computation from somewhat homomorphic encryption. In Reihaneh Safavi-Naini and Ran Canetti, editors, *CRYPTO 2012*, volume 7417 of *LNCS*, pages 643–662. Springer, August 2012.

[EGM96] Shimon Even, Oded Goldreich, and Silvio Micali. On-line/off-line digital signatures. *Journal of Cryptology*, 9(1):35–67, 1996.

[FKN94] Uriel Feige, Joe Kilian, and Moni Naor. A minimal model for secure computation (extended abstract). In *26th ACM STOC*, pages 554–563. ACM Press, May 1994.

[Gen09] Craig Gentry. Fully homomorphic encryption using ideal lattices. In Michael Mitzenmacher, editor, *41st ACM STOC*, pages 169–178. ACM Press, May / June 2009.

[GGP10] Rosario Gennaro, Craig Gentry, and Bryan Parno. Non-interactive verifiable computing: Outsourcing computation to untrusted workers. In Tal Rabin, editor, *CRYPTO 2010*, volume 6223 of *LNCS*, pages 465–482. Springer, August 2010.

[GIS+10] Vipul Goyal, Yuval Ishai, Amit Sahai, Ramarathnam Venkatesan, and Akshay Wadia. Founding cryptography on tamper-proof hardware tokens. In Daniele Micciancio, editor, *TCC 2010*, volume 5978 of *LNCS*, pages 308–326. Springer, February 2010.

[GKP+13a] Shafi Goldwasser, Yael Tauman Kalai, Raluca A. Popa, Vinod Vaikuntanathan, and Nickolai Zeldovich. How to run turing machines on encrypted data. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part II*, volume 8043 of *LNCS*, pages 536–553. Springer, August 2013.

[GKP+13b] Shafi Goldwasser, Yael Tauman Kalai, Raluca A. Popa, Vinod Vaikuntanathan, and Nickolai Zeldovich. Reusable garbled circuits and succinct functional encryption. In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum, editors, *45th ACM STOC*, pages 555–564. ACM Press, June 2013.

[GKR08a] Shafi Goldwasser, Yael Tauman Kalai, and Guy N. Rothblum. Delegating computation: interactive proofs for muggles. In Richard E. Ladner and Cynthia Dwork, editors, *40th ACM STOC*, pages 113–122. ACM Press, May 2008.

[GKR08b] Shafi Goldwasser, Yael Tauman Kalai, and Guy N. Rothblum. One-time programs. In David Wagner, editor, *CRYPTO 2008*, volume 5157 of *LNCS*, pages 39–56. Springer, August 2008.

[Gol04] Oded Goldreich. *Foundations of Cryptography: Basic Applications*, volume 2. Cambridge University Press, Cambridge, UK, 2004.

[Gro11] Jens Groth. Minimizing non-interactive zero-knowledge proofs using fully homomorphic encryption. Cryptology ePrint Archive, Report 2011/012, 2011. `http://eprint.iacr.org/2011/012`.

[GRS08] Henri Gilbert, Matthew J. B. Robshaw, and Yannick Seurin. How to encrypt with the LPN problem. In Luca Aceto, Ivan Damgård, Leslie Ann Goldberg, Magnús M.

Halldórsson, Anna Ingólfsdóttir, and Igor Walukiewicz, editors, *ICALP 2008, Part II*, volume 5126 of *LNCS*, pages 679–690. Springer, July 2008.

[GVW02] Oded Goldreich, Salil P. Vadhan, and Avi Wigderson. On interactive proofs with a laconic prover. *Computational Complexity*, 11(1-2):1–53, 2002.

[GVW12] Sergey Gorbunov, Vinod Vaikuntanathan, and Hoeteck Wee. Functional encryption with bounded collusions via multi-party computation. In Reihaneh Safavi-Naini and Ran Canetti, editors, *CRYPTO 2012*, volume 7417 of *LNCS*, pages 162–179. Springer, August 2012.

[HILL99] Johan Håstad, Russell Impagliazzo, Leonid A. Levin, and Michael Luby. A pseudorandom generator from any one-way function. *SIAM Journal on Computing*, 28(4):1364–1396, 1999.

[HN06] Danny Harnik and Moni Naor. On the compressibility of NP instances and cryptographic applications. In *47th FOCS*, pages 719–728. IEEE Computer Society Press, October 2006.

[IK00] Yuval Ishai and Eyal Kushilevitz. Randomizing polynomials: A new representation with applications to round-efficient secure computation. In *41st FOCS*, pages 294–304. IEEE Computer Society Press, November 2000.

[IK02] Yuval Ishai and Eyal Kushilevitz. Perfect constant-round secure computation via perfect randomizing polynomials. In Peter Widmayer, Francisco Triguero Ruiz, Rafael Morales Bueno, Matthew Hennessy, Stephan Eidenbenz, and Ricardo Conejo, editors, *ICALP 2002*, volume 2380, pages 244–256, July 2002.

[IKM+13] Yuval Ishai, Eyal Kushilevitz, Sigurd Meldgaard, Claudio Orlandi, and Anat Paskin-Cherniavsky. On the power of correlated randomness in secure computation. In Amit Sahai, editor, *TCC 2013*, volume 7785 of *LNCS*, pages 600–620. Springer, March 2013.

[IKOS08] Yuval Ishai, Eyal Kushilevitz, Rafail Ostrovsky, and Amit Sahai. Cryptography with constant computational overhead. In Richard E. Ladner and Cynthia Dwork, editors, *40th ACM STOC*, pages 433–442. ACM Press, May 2008.

[IPS08] Yuval Ishai, Manoj Prabhakaran, and Amit Sahai. Founding cryptography on oblivious transfer - efficiently. In David Wagner, editor, *CRYPTO 2008*, volume 5157 of *LNCS*, pages 572–591. Springer, August 2008.

[Kil88] Joe Kilian. Founding cryptography on oblivious transfer. In *20th ACM STOC*, pages 20–31. ACM Press, May 1988.

[KR09] Yael Tauman Kalai and Ran Raz. Probabilistically checkable arguments. In Shai Halevi, editor, *CRYPTO 2009*, volume 5677 of *LNCS*, pages 143–159. Springer, August 2009.

[LO13] Steve Lu and Rafail Ostrovsky. How to garble RAM programs. In Thomas Johansson and Phong Q. Nguyen, editors, *EUROCRYPT 2013*, volume 7881 of *LNCS*, pages 719–734. Springer, May 2013.

[Mic94] Silvio Micali. CS proofs (extended abstracts). In *35th FOCS*, pages 436–453. IEEE Computer Society Press, November 1994.

[Nie02] Jesper Buus Nielsen. Separating random oracle proofs from complexity theoretic proofs: The non-committing encryption case. In Moti Yung, editor, *CRYPTO 2002*, volume 2442 of *LNCS*, pages 111–126. Springer, August 2002.

[Pei09] Chris Peikert. Public-key cryptosystems from the worst-case shortest vector problem: extended abstract. In Michael Mitzenmacher, editor, *41st ACM STOC*, pages 333–342. ACM Press, May / June 2009.

[Reg05] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In Harold N. Gabow and Ronald Fagin, editors, *37th ACM STOC*, pages 84–93. ACM Press, May 2005.

[RSA78] Ronald L. Rivest, Adi Shamir, and Leonard M. Adleman. A method for obtaining digital signature and public-key cryptosystems. *Communications of the Association for Computing Machinery*, 21(2):120–126, 1978.

[SC07] Radu Sion and Bogdan Carbunar. On the practicality of private information retrieval. In *NDSS 2007*. The Internet Society, February / March 2007.

[Sha83] Adi Shamir. On the generation of cryptographically strong pseudorandom sequences. *ACM Trans. Comput. Syst*, 1(1):38–44, 1983.

[SS10] Amit Sahai and Hakan Seyalioglu. Worry-free encryption: functional encryption with public keys. In Ehab Al-Shaer, Angelos D. Keromytis, and Vitaly Shmatikov, editors, *ACM CCS 10*, pages 463–472. ACM Press, October 2010.

[ST01] Adi Shamir and Yael Tauman. Improved online/offline signature schemes. In Joe Kilian, editor, *CRYPTO 2001*, volume 2139 of *LNCS*, pages 355–367. Springer, August 2001.

[SYY99] Tomas Sander, Adam Young, and Moti Yung. Non-interactive cryptocomputing for NC1. In *40th FOCS*, pages 554–567. IEEE Computer Society Press, October 1999.

[Yao86] Andrew Chi-Chih Yao. How to generate and exchange secrets (extended abstract). In *27th FOCS*, pages 162–167. IEEE Computer Society Press, October 1986.

A Useful properties of REs

Fact A.1 (Substitution). *Suppose that the function $\hat{f}(x;r)$ is a (t,ε)-encoding of $f(x)$ with the simulator and decoder $(\mathsf{Sim},\mathsf{Dec})$. Let $h(z)$ be a function of the form $f(g(z))$ where $z \in \{0,1\}^k$ and $g : \{0,1\}^k \to \{0,1\}^n$. Then, the function $\hat{h}(z;r) = \hat{f}(g(z);r)$ is a (t,ε)-encoding of h with the same simulator and the same decoder.*

Proof. Follows immediately from the definition. For correctness we have:

$$\Pr_r[\mathsf{Dec}(\hat{h}(z;r)) \neq h(z)] = \Pr_r[\mathsf{Dec}(\hat{f}(g(z);r)) \neq f(g(z))] = 0,$$

and for privacy we have

$$\mathsf{Sim}(h(z)) \equiv \mathsf{Sim}(f(g(z))) \equiv_{t,\varepsilon} \hat{f}(g(z);r) \equiv \hat{h}(z;r),$$

as required. □

Fact A.2 (Concatenation). *Suppose that $\hat{f}_i(x;r_i)$ is a (t,ε)-encoding of the function $f_i : \{0,1\}^n \to \{0,1\}^{\ell_i}$ with simulator Sim_i, decoder Dec_i and complexity at most s, for every $i \in [c]$. Then the function $\hat{f}(x;(r_1,\ldots,r_c)) = (\hat{f}_i(x;r_i))_{i=1}^c$ is a $(t-cs, c\varepsilon)$-encoding of $f(x) = (f_1(x),\ldots,f_c(x))$ with simulator $\mathsf{Sim}(y) = (\mathsf{Sim}_i(y_i))_{i=1}^c$ and decoder $\mathsf{Dec}(\hat{y}) = (\mathsf{Dec}_i(\hat{y}_i))_{i=1}^c$.*

Proof. Perfect correctness follows from $\Pr_r[\mathsf{Dec}(\hat{f}(x;r)) \neq f(x)] \leq \sum \Pr_r[\mathsf{Dec}(\hat{f}_i(x;r_i)) \neq f_i(x)] = 0$. Privacy is proved via a standard hybrid argument. Specifically, suppose, towards a contradiction, that $\mathcal{A}$ is a $(t-cs)$ size adversary that distinguishes $\hat{f}(x;r)$ from $\mathsf{Sim}(f(x);\rho)$ with advantage $c\varepsilon$. Then, by an averaging argument, for some $j \in \{1,\ldots,c\}$ the adversary $\mathcal{A}$ distinguishes with advantage at least ε between the tuple

$$(\hat{f}_1(x;r_1),\ldots,\hat{f}_{j-1}(x;r_{j-1}),\mathsf{Sim}_j(f_j(x)),\ldots,\mathsf{Sim}_c(f_c(x)))$$

and the tuple

$$(\hat{f}_1(x;r_1),\ldots,\hat{f}_j(x;r_j),\mathsf{Sim}_{j+1}(f_j(x)),\ldots,\mathsf{Sim}_c(f_c(x))).$$

Now, we can define an adversary $\mathcal{B}$ that ε-distinguishes $\hat{f}_j(x;r_j)$ from $\mathsf{Sim}_j(f_j(x))$. Given a challenge $\hat{y}_j$, the adversary $\mathcal{B}$ samples $(\hat{f}_i(x;r_i))_{i<j}$ and $(\mathsf{Sim}_i(f_i(x)))_{i>j}$ with complexity $c \cdot s$, and invokes $\mathcal{A}$ on the resulting vector with the challenge planted in the j-th position. This gives rise to a (t,ε)-adversary, contradicting our hypothesis. □

Fact A.3 (Composition). *Suppose that:*

- *$g(x;r_g)$ is a (t_1,ε_1)-encoding of $f(x)$ with decoder Dec_g and simulator Sim_g, and*
- *$h((x,r_g);r_h)$ is a (t_2,ε_2)-encoding of the function $g(x,r_g)$, viewed as a single-argument function, with decoder Dec_h, simulator Sim_h and complexity s.*

Then the function $\hat{f}(x;(r_g,r_h)) = h((x,r_g);r_h)$ is a $(\min(t_1-s,t_2),\varepsilon_1+\varepsilon_2)$-encoding of $f(x)$ where (r_g,r_h) are its random inputs and the simulator and decoder are $\mathsf{Sim}(y) = \mathsf{Sim}_h(\mathsf{Sim}_g(y))$ and $\mathsf{Dec}(\hat{y}) = \mathsf{Dec}_g(\mathsf{Dec}_h(\hat{y}))$.

Proof. To prove perfect correctness note that $\Pr_{r_g,r_h}[\mathsf{Dec}(\hat{f}(x;r_g,r_h)) \neq f(x)]$ is upper-bounded by

$$\Pr_{r_g,r_h}[\mathsf{Dec}(h(x,r_g;r_h)) \neq g(x,r_g)] + \Pr_{r_g}[\mathsf{Dec}(\hat{g}(x;r_g)) \neq f(x)] = 0.$$

We prove privacy by noting that $\mathsf{Sim}_g(f(x))$ is (t_1,ε_1)-indistinguishable from $g(x;r_g)$. Hence, $\mathsf{Sim}_h(\mathsf{Sim}_g(f(x)))$ is (t_1-s,ε_1) indistinguishable from $\mathsf{Sim}_h(g(x;r_g))$. However, the latter distribution is (t_2,ε_2)-indistinguishable from $h((x,r_g);r_h)$, and so $h(x;(r_g,r_h))$ is $(\min(t_1-s,t_2),\varepsilon_1+\varepsilon_2)$-indistinguishable from $\mathsf{Sim}_g(f(x))$. □