

Revisiting the Cryptographic Hardness of Finding a Nash Equilibrium

Sanjam Garg¹, Omkant Pandey², and Akshayaram Srinivasan¹(✉)

¹ University of California, Berkeley, USA
{sanjam, akshayaram}@berkeley.edu

² Stony Brook University, Brookhaven, USA
omkant@gmail.com

Abstract. The exact hardness of computing a Nash equilibrium is a fundamental open question in algorithmic game theory. This problem is complete for the complexity class PPAD. It is well known that problems in PPAD cannot be NP-complete unless $NP = \text{coNP}$. Therefore, a natural direction is to reduce the hardness of PPAD to the hardness of problems used in cryptography.

Bitansky, Paneth, and Rosen [FOCS 2015] prove the hardness of PPAD assuming the existence of quasi-polynomially hard indistinguishability obfuscation and sub-exponentially hard one-way functions. This leaves open the possibility of basing PPAD hardness on simpler, polynomially hard, computational assumptions.

We make further progress in this direction and reduce PPAD hardness directly to polynomially hard assumptions. Our first result proves hardness of PPAD assuming the existence of *polynomially hard* indistinguishability obfuscation ($i\mathcal{O}$) and one-way permutations. While this improves upon Bitansky et al.'s work, it does not give us a reduction to simpler, polynomially hard computational assumption because constructions of $i\mathcal{O}$ inherently seems to require assumptions with sub-exponential hardness. In contrast, *public key functional encryption* is a much simpler primitive and does not suffer from this drawback. Our second result shows that PPAD hardness can be based on *polynomially hard* compact public key functional encryption and one-way permutations. Our results further demonstrate the power of polynomially hard compact public key functional encryption which is believed to be weaker than indistinguishability obfuscation. Our techniques are general and we expect them to have various applications.

1 Introduction

The problem of computing a *Nash equilibrium* is fundamental to algorithmic game theory. The hardness of this problem has attracted significant attention. Since a mixed Nash equilibrium is guaranteed to exist for every game [Nas51], the problem belongs to the complexity class TFNP [MP91]. In a series of works, originating from Papadimitriou [Pap94], the problem was established to be complete for the complexity class PPAD [DGP09, CDT09]. PPAD is a subclass of

TFNP containing problems that reduce (in polynomial time) to a special problem called as END-OF-LINE (or EOL in short). Informally, EOL instance includes a “succinct” description of an exponential sized directed graph where each node has in-degree and out-degree at most 1 and a source node having in-degree 0 and out-degree 1. The goal is to find another source or a sink (having in-degree 1 and out-degree 0). It is easy to observe that such a node is guaranteed to exist by a simple parity argument.

The exact hardness of this problem, however, is still not fully understood. Since the class PPAD is *total*, it is unlikely to contain NP-complete problems unless polynomial hierarchy collapses to the first level [MP91, Pap94]. This is similar to the status of hardness assumptions in cryptography which are not believed to be NP-complete, but nevertheless, hard. Due to this similarity, cryptographic problems were suggested as natural candidates in [Pap94] for studying the hardness of PPAD. Indeed, the hardness of some total super-classes of PPAD, such as PPA and PPP, can already be reduced to “standard” cryptographic problems like factoring and collision-resistant hashing [Jer12]. However, such a reduction is not known for PPAD.

A natural extension of this idea is to consider cryptographic problems with a richer and more powerful structure. One of the richest cryptographic structure is *program obfuscation* as formulated by Barak et al. [BGI+12]. It is a compiler to transform any computer program into an “unintelligible one” while preserving its functionality. Ideally, the obfuscation of a program should be a “virtual black-box” (VBB), i.e., access to the obfuscated program should be no better than access to a black-box implementing the program [BGI+12]. Abbot et al. [AKV04] show that PPAD-hardness can be based on VBB obfuscation of a natural pseudo random function. Unfortunately, VBB obfuscation is impossible in general [BGI+12], and there are strong limitations to obfuscating pseudorandom functions [GK05, BCC+14], including the one in [AKV04].

A natural relaxation of VBB obfuscation is *indistinguishability obfuscation* ($i\mathcal{O}$) [BGI+12]. Informally, $i\mathcal{O}$ guarantees that the obfuscation of a circuit looks indistinguishable from the obfuscation of any other, functionally equivalent, circuit of same size. Starting from the work of Garg et al. [GGH+13b], several candidate constructions [BR14, BGK+14, PST14, GLSW15, Zim15, AB15, GMS16] for $i\mathcal{O}$ have been suggested based on various assumptions on multilinear maps [GGH13a] and public key functional encryption [AJ15, BV15a, AJS15].

Motivated by the progress on obfuscation, Bitansky et al. [BPR15] revisit the hardness of PPAD and provide an elegant reduction to the hardness of $i\mathcal{O}$. This is the first reduction of its kind which reduces PPAD-hardness to the security of a concrete and plausible cryptographic primitive. This, together with the progress on $i\mathcal{O}$, gives hope to the possibility of basing PPAD-hardness on simpler, more standard cryptographic primitives.

1.1 Our Contribution

In this work, we revisit the problem of reducing PPAD-hardness to rich and expressive cryptographic systems. We build upon the work of [BPR15] with two specific goals:

- **Rely on polynomial-hardness of $i\mathcal{O}$:** One drawback of the BPR reduction is that it requires $i\mathcal{O}$ schemes with at least quasi-polynomial security. It is not clear if such a large loss in the reduction is necessary. Our first goal is to obtain an improved, polynomial time reduction.
- **Rely on simpler, polynomially hard, assumptions:** While tremendous progress has been made on justifying the security of current $i\mathcal{O}$ schemes, ultimately the security of the resulting constructions still either relies on an exponential number of assumptions (basically, one per pair of circuits), or a polynomial set of assumptions with exponential loss in the reduction. Our second goal is thus to completely get rid of $i\mathcal{O}$ or any other component with non-polynomial time flavor, and reduce PPAD-hardness to simpler, polynomially hard, assumptions.

With respect to our first goal, we prove the following theorem:

Theorem 1. *Assuming the existence of polynomially hard one-way permutations and indistinguishability obfuscation for P/poly, the END-OF-LINE problem is hard for polynomial-time algorithms.*

This polynomially reduces the hardness of PPAD to $i\mathcal{O}$ since PPAD is the class of problems that are reducible to the END-OF-LINE problem.

With respect to our second goal, we show that PPAD-hardness can be reduced to the security of *compact public-key functional encryption* (\mathcal{FE}) in polynomial time. We note that polynomially hard public key functional encryption is a polynomially falsifiable assumption [Nao03].

A public key functional encryption (\mathcal{FE}) scheme for general circuits [BSW11, O’N10] is similar to an ordinary (public-key) encryption scheme with the crucial difference that there are many decryption keys, each of which has an associated function f ; when an encryption of a message m is decrypted with a key for function f , it decrypts to the value $f(m)$. The intuitive security guarantee is that given the secret key corresponding to f and a ciphertext encrypting m , an adversary would not be able to get any information about m except $f(m)$. Our second result proves the following theorem:

Theorem 2. *Assuming the existence of polynomially-hard one-way permutations and compact public key functional encryption for general circuits, the END-OF-LINE problem is hard for polynomial-time algorithms.*

Compact functional encryption, as demonstrated by the recent results of Bitansky and Vaikuntanathan [BV15b] and Ananth et al. [AJS15], can be generically constructed from the so called “collusion-resistant function encryption with collusion-succinct ciphertexts”, which in turn can be constructed from simpler

polynomial hardness assumptions over multi-linear maps, as shown by Garg et al. [GGHZ16]. This is in sharp contrast to $i\mathcal{O}$ where all constructions still inherently seem to require exponential loss in the security reduction¹. Combined with the results of [GGHZ16, BV15b, AJS15], Theorem 2 bases PPAD-hardness on simpler polynomial hardness assumptions. It is interesting to note that compact public key functional encryption implies indistinguishability obfuscators [AJ15, BV15a] but with sub-exponential security loss.

1.2 Our Techniques

We now present a technical overview of our approach. Building upon the work of [BPR15], it suffices to show a sampling procedure that samples hard instances of SINK-OF-VERIFIABLE-LINE problem. We will first show how to generate such instances using polynomially-hard $i\mathcal{O}$ and then discuss how to do the same using polynomially-hard \mathcal{FE} .

PPAD Hardness from Indistinguishability Obfuscation. Let us start by recalling the definition of PPAD. The class PPAD is defined to be the set of all *total* search problems that are polynomial time reducible to the END-OF-LINE (EOL) problem. Intuitively, an EOL instance includes a *succinct* description of an exponential sized directed graph with each node having in-degree and out-degree at most 1. Given a source node (which has in-degree 0 and out-degree 1), the goal is to find another source or a sink (which has in-degree 1 and out-degree 0). By a simple parity argument one can observe that such a node is guaranteed to exist.

The hardness of PPAD was proven in [BPR15] by considering a different problem, proposed in [AKV04], called SINK-OF-VERIFIABLE-LINE problem (SVL) in [BPR15]. It was shown that SVL reduces to the EOL problem [AKV04, BPR15], and therefore hardness of SVL implies hardness of EOL and PPAD.

An instance of the SVL problem is specified by a tuple $(x_s, \text{Succ}, \text{Ver}, T)$ where x_s is called the source node, **Succ** and **Ver** are called successor and verification circuits respectively, and T is a target index. **Succ** succinctly defines an (exponential sized) directed *line graph* starting from the source node x_s . That is, a node x is connected to a node y in the graph through an outgoing edge if and only if $y = \text{Succ}(x)$. **Ver** is used to verify whether a given node is the i^{th} node (starting from the source node x_s) on the path defined by **Succ**. To be more precise, $\text{Ver}(x, i) = 1$ if and only if $x = \text{Succ}^{i-1}(x_s)$. The goal, given the instance, is to find the T -th node (Target) on the path. We want to construct an efficiently samplable distribution over instances of SVL for which no polynomial time algorithm can find the T -th node with non-negligible probability.

BPR Approach. Bitansky et al., building upon [AKV04], consider a line graph where the i -th node is defined by the output of pseudorandom function (PRF) on i , i.e., the i -th node is (i, σ) such that $\sigma = \text{PRF}_S(i)$ for a randomly chosen

¹ An informal explanation of this observation appears in [GLSW15].

key S . Intuitively, σ is a signature on i . The successor circuit of the hard SVL instance, Succ , is then defined by *obfuscating* a “verify and sign” circuit, VS_S , using general purpose $i\mathcal{O}$; VS_S simply outputs the next point $(i+1, \text{PRF}_S(i+1))$ if the input is a valid point (i, σ) and rejects otherwise. The verification circuit Ver simply tests that a given input will not be rejected by the successor circuit. The source node is given by $(1, \text{PRF}_S(1))$ and the target index T is set to a super-polynomial value in the security parameter.

Intuitively, the hardness of the above instance relies on the fact that it is impossible to obtain a signature on a node before obtaining the signature on the previous node in the path. Since T is super-polynomial in the security parameter, it follows that no polynomial time algorithm can obtain a signature on T . While the underlying idea of this reduction is intuitive, reducing its hardness to $i\mathcal{O}$ is more involved. This is shown by first changing the obfuscated circuit Succ so that it does not behave correctly on a randomly chosen point u , and simply outputs \perp . One can think of the Succ circuit being “punctured” at point u . This would also imply that the “punctured” circuit does not output a signature on $u+1$ unlike the original circuit. The next step uses *this fact* to “puncture” the circuit at the point $u+1$. This step is realized through the “punctured” programming approach of Sahai and Waters [SW14]. At a high level, this process is then repeated for the next point $u+2$, and then for $u+3$, and so on, until the circuit does not have the ability to sign on any point in the interval $[u, T]$. Once the circuit is “punctured” at T , it can be observed that no algorithm can find the T^{th} node with non-zero probability. Performing these changes however, requires more care since the number of points in $[u, T]$ is not polynomial. In hindsight, the primary reason for sub-exponential loss in this approach is because it is not possible to “puncture” a larger interval in a “single shot.” In particular, to be able to use the security of $i\mathcal{O}$, this approach must increase the “punctured” interval by one point at a time.

Our Approach: Many Chains of Varying Length. Our main idea is to introduce a richer structure to the nodes in the graph, that avoids the need to increase the “punctured” interval by one point at a time. Instead, we want to make longer “jumps”, sometimes of exponential length, in the proof strategy. Specifically, we aim to make only polynomially many jumps in total to travel from u to T .

In particular, instead of considering one signature per node, we consider κ signatures for every node where 2^κ is the total number of nodes on the line. That is, a node in our graph is of the form $(i, \sigma_1, \dots, \sigma_\kappa)$ where σ_j is a signature on the first j bits of i computed using a key S_j (different for each index) for every $j \in [\kappa]$. The successor circuit is obfuscation of a program which simply checks each signature on appropriate prefixes of i , and if so, it signs all κ prefixes of $i+1$ using appropriate keys. The verification circuit is as before, the source node is simply the signatures on the first node, i.e., $(0^\kappa, \text{PRF}_{S_1}(0), \dots, \text{PRF}_{S_\kappa}(0^\kappa))$, and $T = 2^\kappa - 1$. Observe that the BPR reduction is equivalent to having only σ_κ .

We now explain how this structure on the nodes helps us in achieving a polynomial loss in the reduction. As before, we start by “puncturing” the successor circuit on a random point u . To illustrate the main idea, let us assume

that the binary representation of u has k trailing 1s, i.e., u is of the form: $u_1 \cdots u_{\kappa-k-1} \| 01^k$ where $1 \leq k \leq \kappa$. Then, $u + 1 = u_1 \cdots u_{\kappa-k-1} \| 10^k$, i.e., it has k trailing 0s. Observe that:

1. The first $\kappa - k$ prefix bits of $u + 1$ are *identical* to the first $\kappa - k$ prefix bits of *all points* in the interval $[u + 1, u + 2^k]$.
2. Signature $\sigma_{\kappa-k}$ (corresponding to the prefix of length $\kappa - k$) for the node $u + 1$ is not needed (for checking and signing) anywhere else on the line graph except for nodes in the interval $[u + 1, u + 2^k]$.

As before, suppose that we have punctured the successor circuit at a random node u . Then, the fact that the punctured circuit does not output *any* signature on $u + 1$ means that it does not output the signature $\sigma_{\kappa-k}$ on the first $\kappa - k$ bits of $u + 1$; consequently, and most importantly, this means that it does not output this signature on the first $\kappa - k$ bits of *any point in the interval* $[u + 1, u + 2^k]$. This allows us to increase the interval from $[u + 1, u + 2^k]$ by considering only a constant number of hybrids. We then repeat this process by considering $u + 2^k$ as our next point and iterate until we reach T .

Metaphorically, the signatures can be thought of as “virtual chains” emanating from each node and connecting to other nodes. The first chain coming out of a node i is connected to i ’s immediate neighbor which is $i + 1$. The second chain is connected to a node two hops away from i and the j -th chain is connected to a node 2^j hops away from i and so on. The number of chains coming out from a node i is one more than the number of trailing ones in the binary representation of i . Equivalently, the number of chains coming out of i is the number of bits that change from i to $i + 1$. Puncturing the circuit is viewed as cutting chains of appropriate lengths between points. While BPR strategy always cuts a chain of length 1, our proof strategy cuts the longest possible chain it can and then iterates the process again until it reaches the target T . See Fig. 1 for an illustration.

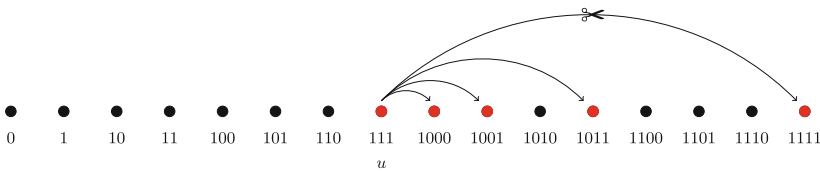


Fig. 1. Illustration of cutting a chain for $u = 0111$

While implementing the above idea we face the difficulty that for a random u the number of chains coming out of u could be very small (as small as 1). We get over this difficulty by initially cutting “smaller” length chains until we have the ability to cut “larger” length chains. Intuitively, this is made possible since the number of trailing 1s in $u + 2^k$ is strictly larger than the number of trailing 1s (given by k) in u . We show that we need to cut no more than a linear (in

the security parameter κ) number of chains to reach T and hence our reduction suffers only a polynomial (in fact linear) loss in the security parameter.

PPAD Hardness from Functional Encryption. We now give a technical overview of our hardness result for PPAD from *compact* functional encryption with *polynomial loss*. As noted earlier, although $i\mathcal{O}$ can be reduced to compact \mathcal{FE} [AJ15, BV15a], we cannot directly rely on this reduction since it suffers sub-exponential security loss. Instead, we try to directly reduce PPAD-hardness to compact \mathcal{FE} .

To directly reduce PPAD-hardness to \mathcal{FE} , we follow the same approach as before, and generate hard on average instances of SVL using functional encryption. To demonstrate the technical challenges while proving the result from \mathcal{FE} we will be considering a single PRF key, as in BPR [BPR15], instead of our idea of using κ keys to implement “multiple chains of varying length”. The scenario with a single PRF key already captures the main technical challenges while keeping the exposition simple. Later, we will explain how to combine the two ideas together to obtain a direct polynomial reduction to \mathcal{FE} .

The line graph implicitly defined by this successor circuit will be similar to the BPR reduction as before. The successor circuit encodes a pseudo random function $\text{PRF}_S : \{0, 1\}^\kappa \rightarrow \{0, 1\}^\kappa$ in its description. The source node is given by $(0^\kappa, \text{PRF}_S(0^\kappa))$. A node (x, σ) is present on the line graph if and only if $\sigma = \text{PRF}_S(x)$. The successor circuit takes as input (x, σ) , checks the validity of the node and if the node is valid outputs $(x + 1, \text{PRF}_S(x + 1))$. The target index is given by $2^\kappa - 1$.

Our goal is to produce an “obfuscated” (or encrypted) version of this successor circuit using \mathcal{FE} . To do this, we will rely on the “binary tree construction” idea of [AJ15, BV15a] for constructing $i\mathcal{O}$ from \mathcal{FE} . Note that though this reduction suffers from sub-exponential loss and we tailor the construction of our successor circuit so that it suffers only from a polynomial loss.

Binary Tree Based Evaluation [AJ15, BV15a]. Let us first recall the main ideas of [AJ15, BV15a] for constructing $i\mathcal{O}$ from \mathcal{FE} . We present an “over-simplified” version of their construction which is actually sufficient for our purposes but is not sufficient for achieving $i\mathcal{O}$ security.

An “obfuscation” for a circuit $C : \{0, 1\}^\kappa \rightarrow \{0, 1\}^*$ is a sequence of $\kappa + 1$ functional keys $\text{FSK}_1, \dots, \text{FSK}_{\kappa+1}$ generated using independently sampled master secret keys $\text{MSK}_1, \dots, \text{MSK}_{\kappa+1}$ along with a ciphertext c_ϕ encrypting the empty string under public-key PK_1 (corresponding to MSK_1). The first κ function keys implement the “bit-extension” functionality. That is, the i^{th} function key corresponds to a function that takes in an $(i - 1)$ -bit string $y \in \{0, 1\}^{i-1}$ and outputs functional encryptions of $y\|0$ and $y\|1$ under PK_{i+1} ². The function key $\text{FSK}_{\kappa+1}$ corresponds to the circuit C .

To evaluate the obfuscated circuit on an input $x \in \{0, 1\}^\kappa$, one does the following: decrypt c_ϕ under FSK_1 to obtain encryptions of 0 and 1. Depending

² The randomness needed for generating the encryptions is obtained using a PRF.

on the bit x_1 , choose either the left or right encryption and decrypt it using FSK_2 and so on. Thus, in κ steps one can obtain an encryption of x under $\text{PK}_{\kappa+1}$ which can be used to compute $C(x)$ using $\text{FSK}_{\kappa+1}$. One can think of the construction as having a binary tree structure where evaluating the circuit on an input x corresponds to traversing along the path labeled x .

Sub-exponential Loss. An intuitive reason for why this construction requires sub-exponential loss to achieve $i\mathcal{O}$ is that the behavior of the obfuscated circuit should be changed on all κ -bit inputs which are 2^κ in number. The key insight in our reduction is that we can achieve our goals by changing the behavior of the obfuscated circuit at only polynomial many inputs and thus incurring only a polynomial security loss.

Our Construction. We will motivate our construction through a series of attempts and fixes.

First Attempt. Our first attempt was to mimic the construction of [AJ15, BV15a]. We generate $2\kappa + 1$ functional keys $\text{FSK}_1, \dots, \text{FSK}_{2\kappa+1}$ where the first 2κ of them correspond to the bit-extension function used for encrypting (x, σ) under $\text{PK}_{2\kappa+1}$ and $\text{FSK}_{2\kappa+1}$ corresponds to the circuit `Next` that checks the validity of the node (x, σ) and outputs the next node in the graph if (x, σ) is valid. The main question with this approach is: How does the circuit `Next` check the validity of the input node and output the next node in the path? The circuit `Next` must somehow have access to the PRF key S but this access should not be “visible” to the outside world.

We definitely cannot hardwire the PRF key S in the circuit as the current constructions of public key functional encryption schemes do not provide any meaningful notions of “function-privacy”. One possible approach is to “propagate” the key S along the entire tree. That is, encrypt the key S in the ciphertext c_ϕ and the bit extension functions output encryptions that also includes S . Though this approach sounds promising, we are unable to use the “punctured” programming techniques of Sahai and Waters that were crucial in the reduction of PPAD hardness to $i\mathcal{O}$. In particular, to puncture the key S at a point x we need to puncture the key along every path thus incurring a sub-exponential loss that we wanted to avoid. To fix this issue, we develop “fine-grained” puncturing techniques.

Second Attempt: “Prefix Puncturing.” To solve the problem explained earlier, we develop techniques to “surgically” puncture the PRF key S along a path x without affecting the distribution on rest of the paths. We now explain the details.

Every string $y \in \{0, 1\}^{\leq \kappa}$ has a natural association with a node in the binary tree where the root is associated with the empty string ϕ . At a high level, we want the set of keys K_y appearing in node y to have the following properties:

- The keys derived from K_y can be used for checking the validity of every node in the subtree rooted at y . This translates to be able to compute the PRF

value at x for every (x, σ) that appears in the subtree rooted at y . We denote this property as *prefix puncturability*.

- The keys derived from K_y can be used for computing the next node for every node in the subtree rooted at y . This would translate to the ability to compute the PRF value at $x + 1$ for every (x, σ) appearing at the subtree rooted at y .

A pseudorandom function that has a natural binary tree structure and has the prefix-puncturable property is the construction due to Goldreich et al. [GGM86]. We exploit this property in the GGM construction to propagate the “prefix-punctured” keys along the binary tree.

At every node $y \in \{0, 1\}^{\leq \kappa}$, we propagate two keys S_y, S_{y+1} where S_y denotes the key S prefix-punctured at string y . Intuitively, S_y is the key used for checking the input node is valid and S_{y+1} is used for generating the next node on the path³. The bit extension function generates $S_{y\|0}, S_{y\|0+1}$ and $S_{y\|1}, S_{y\|1+1}$ from S_y, S_{y+1} and propagates these values along with $y\|0$ and $y\|1$ respectively. The circuit Next receives S_x, S_{x+1} where $x \in \{0, 1\}^\kappa$ and checks the validity of the input signature using S_x and generates the next node in the path if the input is valid using S_{x+1} .

Note that the puncturing of the keys does not happen after the level κ as by this time we have parsed the x which completely determines the key S_x, S_{x+1} . Therefore, we need to propagate S_x, S_{x+1} along the entire subtree rooted at x where we parse σ . This creates the following problem: consider a scenario where the successor circuit already outputs \perp on the point x and we are trying to extend the interval to include $x + 1$. Recall that the crucial idea behind the ability to increase the interval is that S_{x+1} does not occur anywhere else in the computation of the circuit. We observe that S_{x+1} gets propagated along the entire subtree (of exponential size) rooted at x where the input σ is parsed. Hence, to “remove all traces” of S_{x+1} along the subtree rooted at x , we need to incur a sub-exponential loss.

Final Construction: “Encrypt the Next Signature.” We solve the above problem by “implicitly” checking whether the given node is valid. This implicit checking is facilitated by encrypting the signature on the next node by using the signature on the current node. Intuitively, an evaluator can obtain the signature on the next node if and only if he holds a valid signature on the current node.

Instead of propagating the keys S_x, S_{x+1} in clear in the subtree parsing σ , we “cut-short” the tree at level where x is parsed. Once x is parsed (and hence we have the values S_x and S_{x+1}), we apply a length doubling injective pseudo random generator PRG on the signature S_x to obtain two halves $\text{PRG}_0(S_x)$ and $\text{PRG}_1(S_x)$. We encrypt S_{x+1} under $\text{PRG}_1(S_x)$ and output the encryption along with $\text{PRG}_0(S_x)$. The Next circuit takes $\sigma, \text{PRG}_0(S_x)$ and the encrypted version

³ Note that instead of S_{y+1} it is enough to propagate $S_{y+1\|0^{\kappa-|y|}}$. It is in fact crucial for our reduction that we propagate $S_{y+1\|0^{\kappa-|y|}}$ instead of S_{y+1} . But we will use S_{y+1} for ease of notation and exposition.

of S_{x+1} and checks whether $\text{PRG}_0(\sigma) = \text{PRG}_0(S_x)$ ⁴ and if yes it decrypts using $\text{PRG}_1(\sigma)$ to obtain S_{x+1} . Notice that now we don't run into the same problem while trying to increase the interval to include S_{x+1} . This is because we can first change S_x to a random string by relying on pseudo randomness at punctured point property of GGM PRF and then relying on semantic security of secret key encryption we can change the encryption under $\text{PRG}_1(S_x)$ to some junk value. Implementing these two steps is non-trivial and we rely on "hidden trapdoor" technique of Ananth et al. [ABSV15] while generating the function keys to achieve this.

Note that we still haven't explained how the successor circuit is "punctured" at a random point in the first place. To this end, we "artificially" change the honest execution of the circuit to have a hardwired random value v and the circuit checks if $\text{PRG}(x) = v$ and if so outputs \perp . The honest execution does not output \perp for any input x with overwhelming probability since PRG has sparse images. We then change this random v to $\text{PRG}(u)$ for a random u relying on the security of the PRG. A consequence of this fix is that even our honest evaluation of the successor circuit looks somewhat "artificial". This seems necessary to circumvent the sub-exponential loss incurred while constructing obfuscation from functional encryption.

Putting it All Together. To show hardness of PPAD from \mathcal{FE} by incurring polynomial loss in the security reduction we need to combine the above ideas with that of "multiple-chains of varying length". As explained in the chain-cutting technique we generate κ GGM keys S_1, \dots, S_κ . We propagate the "prefix-punctured" keys corresponding to every index $i \in [\kappa]$ along every node in the binary tree. A careful reader might have noticed that though it is necessary to check the validity of the input signatures for every prefix, it is actually sufficient to generate signatures on the next node on the path only for those bit positions that change when incrementing by 1. This is because for the rest of the bit positions that share the same prefix with the input node and we can just output those input signatures along with those newly computed ones, provided the input is valid. This observation is in fact crucial to prove the security of our construction. We need to ensure that the Next circuit must have the ability to check the validity of every signature but it has access only to those prefix punctured keys corresponding to the bit positions that change when incrementing by 1.

We satisfy these two "conflicting" properties by decoupling the process of checking the input signatures and the process of generating the next node on the path. In order to check the input signatures we propagate $\text{PRG}_0(S_{i,x})$ for every $i \in [\kappa]$ and to generate the signatures on the next node on the path we propagate an encrypted version of $S_{j,x+1}$ under $\text{PRG}_1(S_{j,x})$ only for those bits j that change when incrementing x .

⁴ We need this explicit check for the verification circuit to decide if a particular node is an i^{th} node or not. Also, we need a stronger property on pseudo random generator called as left half injectivity for this check to be correct always.

1.3 Subsequent Work

Garg et al. in [GPSZ16] extended our techniques to base Trapdoor Permutations on polynomial hardness of compact Functional Encryption. In the same work, they also showed how to base Non-Interactive Key Exchange (NIKE) for unbounded parties from polynomially hard compact Functional Encryption. Recently, Garg and Srinivasan [GS16] extended our techniques to construct adaptively secure Functional Encryption against unbounded collusions from single-key, selectively secure Functional encryption with weakly compact ciphertexts.

Rosen et al. [RSS16] investigated the possibility of basing average-case PPAD hardness on standard cryptographic assumptions. They showed that average-case PPAD hardness does not imply one-way functions in a black-box manner and average-case SVL hardness cannot be based on injective trapdoor functions in a black-box manner. An implication of this work is that it might be possible to base PPAD hardness on one-way functions but such a result has to use techniques that significantly deviate from Bitansky et al. [BPR15] and our work.

Hubáček and Yogev [HY16] extended our result to base hardness of a complexity class CLS on compact Functional Encryption. CLS is a sub-class of PPAD and captures Continuous Local Search problems. They showed a reduction between the SVL problem and a problem called as END-OF-METERED-LINE which is contained in CLS. This allowed them to base hardness of CLS on polynomially hard compact Functional Encryption.

2 PPAD

A large part of this section is taken verbatim from [BPR15]. A search problem is given by a tuple (I, R) . I defines the set of instances and R is an NP relation. Given $x \in I$, the goal is to find a witness w (if it exists) such that $R(x, w) = 1$. We say that a search problem (I_1, R_1) polynomial time reduces to another search problem (I_2, R_2) if there exists polynomial time algorithms P, Q such that for every $x_1 \in I_1$, $P(x_1) \in I_2$ and given w_2 such that $(P(x_1), w_2) \in R_2$, $R_1(x_1, Q(w_2)) = 1$.

A search problem is said to be *total* if for any $x \in \{0, 1\}^*$, there exists a polynomial time procedure to test whether $x \in I$ and for all $x \in I$, the set of witnesses w such that $R(x, w) = 1$ is non-empty. The class of total search problems is denoted by TFNP. PPAD [Pap94] is a subset of TFNP and is defined by its complete problem called as END-OF-LINE (abbreviated as EOL).

Definition 1 [Pap94]. $\text{EOL} = \{I_{\text{EOL}}, R_{\text{EOL}}\}$ where $I_{\text{EOL}} = \{(x_s, \text{Succ}, \text{Pred}) : \text{Succ}(x_s) \neq x_s = \text{Pred}(x_s)\}$ and $R_{\text{EOL}}((x_s, \text{Succ}, \text{Pred}), w) = 1$ iff $(\text{Pred}(\text{Succ}(w)) \neq w) \vee (\text{Succ}(\text{Pred}(w)) \neq w \wedge w \neq x_s)$.

Definition 2 [Pap94]. The complexity class PPAD is the set of all search problems (I, R) such that $(I, R) \in \text{TFNP}$ and (I, R) polynomial time reduces to EOL.

A related problem to EOL is the SINK-OF-VERIFIABLE-LINE (abbreviated as SVL) which is defined as follows:

Definition 3 [AKV04,BPR15]. $\text{SVL} = \{I_{\text{SVL}}, R_{\text{SVL}}\}$ where $I_{\text{SVL}} = \{(x_s, \text{Succ}, \text{Ver}, T)\}$ and $R_{\text{SVL}}((x_s, \text{Succ}, \text{Ver}, T), w) = 1$ iff $(\text{Ver}(w, T) = 1)$.

SVL instance defines a single directed path with the source being x_s . Succ is the successor circuit and there is a directed edge between u and v if and only if $\text{Succ}(u) = v$. Ver is the *verification* circuit and is used to test whether a given node is the i^{th} node from x_s . That is, $\text{Ver}(x, i) = 1$ iff $x = \text{Succ}^{i-1}(x_s)$. The goal is to find the T^{th} node in the path. It is easy to observe that for every *valid* SVL instance the set of witness w is not empty. But SVL may not be total since there is no known efficient procedure to test whether the instance is valid or not. But it was shown in [AKV04,BPR15] that SVL polynomial time reduces to EOL.

Lemma 1 [AKV04,BPR15]. SVL *polynomial time reduces to EOL*.

3 Preliminaries

κ denotes the security parameter. A function $\mu(\cdot) : \mathbb{N} \rightarrow \mathbb{R}^+$ is said to be negligible if for all polynomials $\text{poly}(\cdot)$, $\mu(\kappa) < \frac{1}{\text{poly}(\kappa)}$ for large enough κ . For a probabilistic algorithm \mathcal{A} , we denote by $\mathcal{A}(x; r)$ the output of \mathcal{A} on input x with the content of the random tape being r . We will omit r when it is implicit from the context. We denote $y \leftarrow \mathcal{A}(x)$ as the process of sampling y from the output distribution of $\mathcal{A}(x)$ with a uniform random tape. For a finite set S , we denote $x \xleftarrow{\$} S$ as the process of sampling x uniformly from the set S . We model non-uniform adversaries $\mathcal{A} = \{\mathcal{A}_\kappa\}$ as circuits such that for all κ , \mathcal{A}_κ is of size $p(\kappa)$ where $p(\cdot)$ is a polynomial. We will drop the subscript κ from the adversary's description when it is clear from the context. We will also assume that all algorithms are given the unary representation of security parameter 1^κ as input and will not mention this explicitly when it is clear from the context. We will use PPT to denote Probabilistic Polynomial Time algorithm. We denote $[k]$ to be the set $\{1, \dots, k\}$. We will use $\text{negl}(\cdot)$ to denote an unspecified negligible function and $\text{poly}(\cdot)$ to denote an unspecified polynomial.

A binary string $x \in \{0, 1\}^\kappa$ is represented as $x_1 \cdots x_\kappa$. x_1 is the most significant (or the highest order bit) and x_κ is the least significant (or the lowest order bit). The i -bit prefix $x_1 \cdots x_i$ of the binary string x is denoted by $x_{[i]}$. We use $x||y$ to denote concatenation of binary strings x and y . We say that a binary string y is a prefix of x if and only if there exists a string $z \in \{0, 1\}^*$ such that $x = y||z$.

Injective Pseudo Random Generator. We give the definition of an injective Pseudo Random Generator PRG.

Definition 4. An injective pseudo random generator PRG is a deterministic polynomial time algorithm with the following properties:

- **Expansion:** There exists a polynomial $\ell(\cdot)$ (called as the expansion factor) such that for all κ and $x \in \{0, 1\}^\kappa$, $|\text{PRG}(x)| = \ell(\kappa)$.

- **Pseudo randomness:** For all κ and for all poly sized adversaries \mathcal{A} ,

$$|\Pr[\mathcal{A}(\text{PRG}(U_\kappa)) = 1] - \Pr[\mathcal{A}(U_{\ell(\kappa)}) = 1]| \leq \text{negl}(\kappa)$$

where U_i denotes the uniform distribution on $\{0, 1\}^i$.

- **Injectivity:** For every κ and for all $x, x' \in \{0, 1\}^\kappa$ such that $x \neq x'$, $\text{PRG}(x) \neq \text{PRG}(x')$.

We in fact need an additional property from an injective PRG. Let us consider PRG where the expansion factor (or the output length) is given by $2 \cdot \ell(\cdot)$. Let us denote the first $\ell(\cdot)$ bits of the output of the PRG by the function PRG_0 and the next $\ell(\cdot)$ bits of the output of the PRG by PRG_1 .

Definition 5. A pseudo random generator PRG is said to be left half injective if for every κ and for all $x, x' \in \{0, 1\}^\kappa$ such that $x \neq x'$. $\text{PRG}_0(x) \neq \text{PRG}_0(x')$.

Note that left half injective PRG is also an injective PRG. We note that the standard construction of pseudo random generator for arbitrary polynomial stretch from one-way permutations is left half injective. For completeness, we state the construction:

Lemma 2. Assuming the existence of one-way permutations, there exists a pseudo random generator that is left half injective.

Proof. Let $f : \{0, 1\}^\kappa \rightarrow \{0, 1\}^\kappa$ be a one-way permutation with hardcore predicate $B : \{0, 1\}^\kappa \rightarrow \{0, 1\}$ [GL89]. Let G be an algorithm defined as follows: On input $x \in \{0, 1\}^\kappa$, $G(x) = f^n(x) \| B(x) \| B(f(x)) \cdots B(f^{n-1}(x))$ where $n = 2\ell(\kappa) - \kappa$. Clearly, $|G(x)| = 2\ell(\kappa)$. The pseudo randomness property of $G(\cdot)$ follows from the security of hardcore bit. The left half injectivity property follows from the observation that f^n is a permutation.

Puncturable Pseudo Random Function. We recall the notion of puncturable pseudo random function from [SW14]. The construction of pseudo random function given in [GGM86] satisfies the following definition [BW13, KPTZ13, BGI14].

Definition 6. A puncturable pseudo random function \mathcal{PRF} is a tuple of PPT algorithms $(\text{KeyGen}_{\mathcal{PRF}}, \text{PRF}, \text{Punc})$ with the following properties:

- **Efficiently Computable:** For all κ and for all $S \leftarrow \text{KeyGen}_{\mathcal{PRF}}(1^\kappa)$, $\text{PRF}_S : \{0, 1\}^{\text{poly}(\kappa)} \rightarrow \{0, 1\}^\kappa$ is polynomial time computable.
- **Functionality is preserved under puncturing:** For all κ , for all $y \in \{0, 1\}^\kappa$ and $\forall x \neq y$,

$$\Pr[\text{PRF}_{S\{y\}}(x) = \text{PRF}_S(x)] = 1$$

where $S \leftarrow \text{KeyGen}_{\mathcal{PRF}}(1^\kappa)$ and $S\{y\} \leftarrow \text{Punc}(S, y)$.

- **Pseudo randomness at punctured points:** For all κ , for all $y \in \{0, 1\}^\kappa$, and for all poly sized adversaries \mathcal{A}

$$|\Pr[\mathcal{A}(\text{PRF}_S(y), S\{y\}) = 1] - \Pr[\mathcal{A}(U_\kappa, S\{y\}) = 1]| \leq \text{negl}(\kappa)$$

where $S \leftarrow \text{KeyGen}_{\mathcal{PRF}}(1^\kappa)$, $S\{y\} \leftarrow \text{Punc}(S, y)$ and U_κ denotes the uniform distribution over $\{0, 1\}^\kappa$.

Indistinguishability Obfuscator. We now define Indistinguishability obfuscator from [BGI+12,GGH+13b].

Definition 7. A PPT algorithm $i\mathcal{O}$ is an indistinguishability obfuscator for a family of circuits $\{C_\kappa\}_\kappa$ that satisfies the following properties:

- **Correctness:** For all κ and for all $C \in C_\kappa$ and for all x ,

$$\Pr[i\mathcal{O}(C)(x) = C(x)] = 1$$

where the probability is over the random choices of $i\mathcal{O}$.

- **Security:** For all $C_0, C_1 \in C_\kappa$ such that for all x , $C_0(x) = C_1(x)$ and for all poly sized adversaries \mathcal{A} ,

$$|\Pr[\mathcal{A}(i\mathcal{O}(C_0)) = 1] - \Pr[\mathcal{A}(i\mathcal{O}(C_1)) = 1]| \leq \text{negl}(\kappa)$$

Functional Encryption. We recall the notion of functional encryption with selective indistinguishability based security [BSW11,O’N10].

A functional encryption \mathcal{FE} is a tuple of PPT algorithms (FE.Setup, FE.Enc, FE.KeyGen, FE.Dec) with the message space $\{0, 1\}^*$ having the following syntax:

- FE.Setup(1^κ): Takes as input the unary encoding of the security parameter κ and outputs a public key PK and a master secret key MSK .
- FE.Enc $_{PK}(m)$: Takes as input a message $m \in \{0, 1\}^*$ and outputs an encryption C of m under the public key PK .
- FE.KeyGen(MSK, f): Takes as input the master secret key MSK and a function f (given as a circuit) as input and outputs the function key FSK_f .
- FE.Dec(FSK_f, C): Takes as input the function key FSK_f and the ciphertext C and outputs a string y .

Definition 8 (Correctness). The functional encryption scheme \mathcal{FE} is correct if for all κ and for all messages $m \in \{0, 1\}^*$,

$$\Pr \left[y = f(m) \left| \begin{array}{l} (PK, MSK) \leftarrow \text{FE.Setup}(1^\kappa) \\ C \leftarrow \text{FE.Enc}_{PK}(m) \\ FSK_f \leftarrow \text{FE.KeyGen}(MSK, f) \\ y \leftarrow \text{FE.Dec}(FSK_f, C) \end{array} \right. \right] = 1 \quad (1)$$

Definition 9 (Selective Security). For all κ and for all poly sized adversaries \mathcal{A} ,

$$|\Pr[\text{Expt}_{1^\kappa,0,\mathcal{A}} = 1] - \Pr[\text{Expt}_{1^\kappa,1,\mathcal{A}} = 1]| \leq \text{negl}(\kappa)$$

where $\text{Expt}_{1^\kappa,b,\mathcal{A}}$ is defined below:

- **Challenge Message Queries:** The adversary \mathcal{A} outputs two messages m_0, m_1 such that $|m_0| = |m_1|$ to the challenger.
- The challenger samples $(PK, MSK) \leftarrow \text{FE.Setup}(1^\kappa)$ and generates the challenge ciphertext $C \leftarrow \text{FE.Enc}_{PK}(m_b)$. It then sends (PK, C) to \mathcal{A} .

- **Function Queries:** \mathcal{A} submits function queries f to the challenger. The challenger responds with $\text{FSK}_f \leftarrow \text{FE.KeyGen}(\text{MSK}, f)$.
- If \mathcal{A} makes a query f to functional key generation oracle such that $f(m_0) \neq f(m_1)$, output of the experiment is \perp . Otherwise, the output is b' which is the output of \mathcal{A} .

Remark 1. We say that the functional encryption scheme \mathcal{FE} is **single-key, selectively secure** if the adversary \mathcal{A} in $\text{Expt}_{1^\kappa, b, \mathcal{A}}$ is allowed to query the functional key generation oracle $\text{FE.KeyGen}(\text{MSK}, \cdot)$ on a single function f .

Definition 10. (Compactness, [AJS15, BV15a, AJ15]). The functional encryption scheme \mathcal{FE} is said to be compact if for all $\kappa \in \mathbb{N}$ and for all $m \in \{0, 1\}^*$ the running time of the encryption algorithm FE.Enc is $\text{poly}(\kappa, |m|)$.

Prefix Puncturable Pseudo Random Functions. We now define the notion of prefix puncturable pseudo random function PPRF which is satisfied by the construction of the pseudo random function in [GGM86].

Definition 11. A prefix puncturable pseudo random function PPRF is a tuple of PPT algorithms $(\text{KeyGen}_{\text{PPRF}}, \text{PrefixPunc})$ satisfying the following properties:

- **Functionality is preserved under repeated puncturing:** For all κ , for all $y \in \bigcup_{k=0}^{\text{poly}(\kappa)} \{0, 1\}^k$ and for all $x \in \{0, 1\}^{\text{poly}(\kappa)}$ such that there exists a $z \in \{0, 1\}^*$ s.t. $x = y \parallel z$,

$$\Pr[\text{PrefixPunc}(\text{PrefixPunc}(S, y), z) = \text{PrefixPunc}(S, x)] = 1$$

where $S \leftarrow \text{KeyGen}_{\text{PPRF}}(1^\kappa)$.

- **Pseudorandomness at punctured prefix:** For all κ , for all $x \in \{0, 1\}^{\text{poly}(\kappa)}$, and for all poly sized adversaries \mathcal{A}

$$|\Pr[\mathcal{A}(\text{PrefixPunc}(S, x), \text{Keys}) = 1] - \Pr[\mathcal{A}(U_\kappa, \text{Keys}) = 1]| \leq \text{negl}(\kappa)$$

where $S \leftarrow \text{KeyGen}_{\text{PPRF}}(1^\kappa)$ and $\text{Keys} = \{\text{PrefixPunc}(S, x_{[i-1]}) \parallel (1 - x_i)\}_{i \in [\text{poly}(\kappa)]}$.

4 Hardness from Indistinguishability Obfuscation

In this section, we prove that SVL is hard on average assuming polynomial hardness of indistinguishability obfuscation, injective PRGs and puncturable pseudo random functions. Coupled with the fact that SVL reduces to EOL (Lemma 1) we have the following theorem.

Theorem 3. Assume the existence of one-way permutations and indistinguishability obfuscation against polynomial time adversaries then we have that EOL problem is hard for polynomial time algorithms.

4.1 Hard on Average SVL Instances

In this section, we describe an efficient sampler that provides hard on average instances $(x_s, \text{Succ}, \text{Ver}, 1^\kappa)$ of SVL. Here x_s is the source node and Succ is the successor circuit. We define a directed edge between u and v if and only if $\text{Succ}(u) = v$. Ver is the verification circuit and is used to test whether a given node is the k^{th} node from x_s . That is, $\text{Ver}(x, k) = 1$ iff $x = \text{Succ}^{k-1}(x_s)$. For the generated instances, we argue that it is hard to find the 1^κ node in the path.

The formal description of hard on average SVL instance sampler is provided in Fig. 3. Internally this sampler generates an obfuscation of the Next circuit provided in Fig. 2. Next we describe the SVL instances which we consider informally.

The instance we generate defines a line graph. The nodes in the graph are of the form: $(x, \sigma_1, \dots, \sigma_\kappa)$ where $x \in \{0, 1\}^\kappa$. The nodes satisfy the following relation: for all $i \in [\kappa]$, $\text{PRF}_{S_i}(x_{[i]}) = \sigma_i$ and in that case we say that $(x, \sigma_1, \dots, \sigma_\kappa)$ is *valid*. The node $(x, \sigma_1, \dots, \sigma_\kappa)$ is connected to $(x + 1, \sigma'_1, \dots, \sigma'_\kappa)$ through an outgoing edge and is connected to $(x - 1, \sigma''_1, \dots, \sigma''_\kappa)$ through an incoming edge where $\sigma'_1, \dots, \sigma'_\kappa$ and $\sigma''_1, \dots, \sigma''_\kappa$ satisfy the above described PRF relationship. The source node is given by $(0^\kappa, \text{PRF}_{S_1}(0), \dots, \text{PRF}_{S_\kappa}(0^\kappa))$.

At a very high level successor circuit of our SVL instances provides a method for moving forward from one node to the next. The successor circuit in our instances corresponds to an obfuscation of the Next circuit. This circuit on input a node of the form $(x, \sigma_1, \dots, \sigma_\kappa)$ checks for the validity of the input. If it is valid, it outputs the next node $(x + 1, \sigma'_1 \dots \sigma'_\kappa)$ where $\sigma'_i = \text{PRF}_{S_i}((x + 1)_{[i]})$ in the path. On an invalid input, it outputs \perp .

Input: $(x, \sigma_1, \dots, \sigma_\kappa)$
Hardcoded Parameters: S_1, \dots, S_κ

1. For any $i \in [\kappa]$, if $\sigma_i \neq \text{PRF}_{S_i}(x_{[i]})$ then output \perp .
2. If $x = 1^\kappa$, then output SOLVED.
3. Else output $(x + 1, \sigma'_1, \dots, \sigma'_\kappa)$, where for all $i \in [\kappa]$ compute $\sigma'_i = \text{PRF}_{S_j}((x + 1)_{[i]})$.

Padding: This circuit is padded so that total size of the circuit is $p(\kappa)$, for some polynomial $p(\cdot)$ specified later.

Fig. 2. $\text{Next}_{S_1, \dots, S_\kappa}$

For the hard SVL instances we additionally need to provide a verification circuit. The verification circuit just uses the successor circuit in a very natural manner. The verification circuit on input $(x, \sigma_1, \dots, \sigma_\kappa, j)$ outputs 1 if and only if $x = j - 1$ and $\text{Next}_{S_1, \dots, S_\kappa}(x, \sigma_1, \dots, \sigma_\kappa) \neq \perp$.

Due to space constraints we defer the proof of hardness to full version of this paper [GPS15].

- **Sampled Ingredients:** Sample $\{S_i\}_{i \in [\kappa]} \leftarrow \text{KeyGen}_{\mathcal{PRF}}(1^\kappa)$. For all $i \in [\kappa]$, S_i is a seed for a PRF mapping i bits to κ bits. That is, $\text{PRF}_{S_i} : \{0, 1\}^i \rightarrow \{0, 1\}^\kappa$.
- **Source Node:** The source node $x_s = (0^\kappa, \text{PRF}_{S_1}(0), \dots, \text{PRF}_{S_\kappa}(0^\kappa))$.
- **Successor Circuit:** The successor circuit is given by $i\mathcal{O}(\text{Next}_{S_1, \dots, S_\kappa})$ where the circuit $\text{Next}_{S_1, \dots, S_\kappa}$ is described in Figure 2.
- **Verification Circuit:** The verification circuit, given by Ver , on input $((x, \sigma_1 \dots \sigma_\kappa), j)$ checks if $x = j - 1$ and $i\mathcal{O}(\text{Next}_{S_1, \dots, S_\kappa})((x, \sigma_1 \dots \sigma_\kappa)) \neq \perp$.

Fig. 3. Sampler for hard on average instances of SVL based on hardness of $i\mathcal{O}$

5 Hardness Result Based on Functional Encryption

In this section we show that SVL is hard on average assuming polynomially hard functional encryption and one-way permutations. Coupled with the fact that SVL reduces to EOL (Lemma 1) we have the following theorem.

Theorem 4. *Assume the existence of one-way permutations and functional encryption against polynomial time adversaries then we have that EOL problem is hard for polynomial time algorithms.*

Recall that hard SVL instance based on $i\mathcal{O}$ (Sect. 4), required κ puncturable PRF keys. Basing hardness on polynomially hard functional encryption requires us to still maintain κ keys. However, now we need to use prefix-puncturing (see Definition 11) which is more delicate and needs to be handled carefully. Consequently the construction ends up being complicated. However, the special mechanism of prefix-puncturing that we use is crucial to understanding our construction. So towards simplifying exposition, we start by abstracting out the details of this puncturing and present a special tree structure and some properties about it next.

5.1 Special Tree Key Structure

Let $x_{[i]}$ denote the first i (higher order) bits of x i.e. $x_1 \dots x_i$. Now note that any $y \in \{0, 1\}^i$ can be identified with a node in a binary tree for which nodes at depth i correspond to strings $\{0, 1\}^i$. Note that the root of the tree corresponds to the empty string ϕ . As previously mentioned our construction needs κ PPRF keys, namely S_1, \dots, S_κ . The key S_i works on inputs of length i . We use $S_{i,x}$ to denote the key S_i prefix punctured at a string $x \in \{0, 1\}^{\leq i}$.

Looking ahead, in our hard-on-average instances of SVL each $x \in \{0, 1\}^\kappa$ will be attached with *associated signature* values $\sigma_1, \dots, \sigma_\kappa$ where for each $i \in [\kappa]$ we have that $\sigma_i = \text{PrefixPunc}(S_i, x_{[i]})$. Furthermore in our construction given x and the associated signature values, we will need to verify these values and provide the associated signature values for $x + 1$, but this has to be done in

a circuitous manner because of several security reasons. We do not delve into the security arguments right away, but focus on describing the prefix-puncturing that we need to perform.

We next describe the set V_x^i where $x \in \{0, 1\}^{\leq i}$, which contains suitable prefix-puncturings of the key S_i . Intuitively, we want this set to contain all keys that will allow us to perform the task of *checking the validity* of the i^{th} associated signature on any input of the form $x\|y$ where $y \in \{0, 1\}^{\kappa-|x|}$ as well as *computing* the i^{th} associated signature for $(x\|y) + 1$. Furthermore, it should suffice to generate $V_{x\|y}^i$ for all y . For any node $x \in \{0, 1\}^{\leq i}$, this very naturally translates to the keys $S_{i,x}$ and $S_{i,x+1}$. A careful reader might have noticed that instead of $S_{i,x+1}$, it in fact suffices to just have $S_{i,(x+1)\|0^{i-|x|}}$. As it turns out we must only include $S_{i,(x+1)\|0^{i-|x|}}$. Including $S_{i,x+1}$ prevents the Derivability Lemma (Lemma 4) from going through.

Recall that the key S_i corresponds to a PPRF key for inputs of length i . Therefore, for $x\|y$ such that $|x| = i$, the key S_i can be prefix-punctured only for the prefix $x = (x\|y)_{[i]}$. This raises the following question. Should we include $S_{i,x}$ and $S_{i,x+1}$ in all $V_{x\|y}^i$? As we will see later, in our construction, we carefully decouple the checking of associated signatures from the generation of new associated signatures. An important consequence, relevant here is that, even though the checks need to be performed for all $x\|y$, a new i^{th} associated signature needs to be generated for only one choice of y , namely $1^{\kappa-|x|}$ (the all 1 string of length $\kappa - |x|$). This design choice (which is crucial for polynomial security loss) also allows us to set $V_{x\|y}^i$ for all other choices of y to be \emptyset . In terms of the binary tree structure one can think of this as V_x^i getting passed only along the rightmost path in the subtree rooted at x . At a very high level, this allows us to argue that the key S_i (proved formally in Lemma 4) can be punctured at a special point by removing keys from V_x^i for only a polynomial number of choices of x and i . This is crucial for ensuring that our proof of security has only a polynomial number of hybrids.

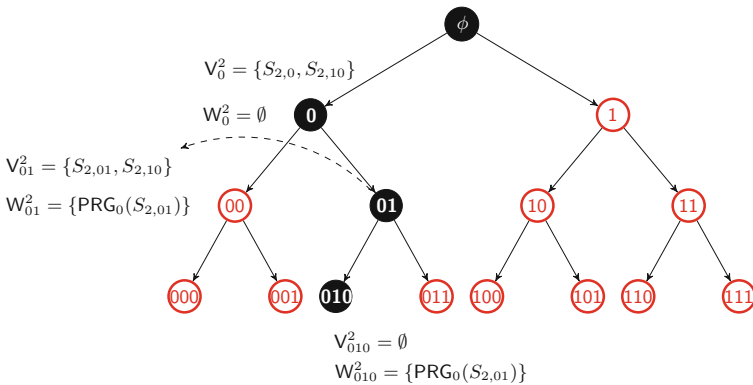


Fig. 4. Example of values contained in V_x^2 for $x \in \{0, 1\}^{\leq 3}$.

Next note that dropping keys from $V_{x\|y}^i$ (such that $|x| = i$) hinders the checking of associated signatures provided along with inputs $x\|y$ where $y \neq 1^{\kappa-i}$. We tackle this issue by introducing a vestigial set $W_{x\|y}^i$ corresponding to each $V_{x\|y}^i$. This vestigial set contains remnants of the keys that were dropped from V_x^i . We craft these remnants to be such that they suffice for performing the necessary checks. In particular, we set these remnants to be the left half of an left half injective PRG evaluation on the dropped key.

More formally, V_x^i and W_x are defined as follows. In the following, for any $i \in [\kappa]$ we treat $1^i + 1$ as 1^i , and $\phi + 1$ as ϕ . Here 1^i is a string of i 1s and ϕ is the empty string.

$$\begin{aligned}
 V_x &= \bigcup_{i \in [\kappa]} V_x^i & V_x^i &= \begin{cases} \{S_{i,x[i]}, S_{i,x[i]+1}\} & \text{if } |x| > i \text{ and } x = x_{[i]}\|1^{|x|-i} \\ \{S_{i,x}, S_{i,(x+1)\|0^{i-|x|}}\} & \text{if } |x| \leq i \\ \emptyset & \text{otherwise} \end{cases} \\
 W_x &= \bigcup_{i \in [\kappa]} W_x^i & W_x^i &= \begin{cases} \{\text{PRG}_0(S_{i,x[i]})\} & \text{if } |x| \geq i \\ \emptyset & \text{otherwise} \end{cases}
 \end{aligned}$$

For the empty string $x = \phi$, these sets can be initialized as follows.

$$\begin{aligned}
 V_\phi &= \bigcup_{i \in [\kappa]} V_\phi^i & V_\phi^i &= \{S_i\} \\
 W_\phi &= \bigcup_{i \in [\kappa]} W_\phi^i & W_\phi^i &= \emptyset
 \end{aligned}$$

Illustration with an Example. Finally we explain what sets V_x^2, W_x^2 contain when x is a prefix of 010 in Fig. 4. At the root node we have $V_\phi^2 = \{S_2\}$ and $W_\phi = \emptyset$. The set V_0^2 contains $S_{2,0}$ and $S_{2,10}$ and the set W_0^2 is still empty. Next note that V_{01}^2 contains $S_{2,01}, S_{2,10}$ and W_{01}^2 contains $\text{PRG}_0(S_{2,01})$. Finally set $V_{010}^2 = \emptyset$ and W_{010}^2 continues to contain $\text{PRG}_0(S_{2,01})$.

Properties of the Special Tree Key Structure. We now prove several properties about the special tree key structure. Intuitively speaking the crux of the lemmas is the claim V-set for can a node can be used to derive its children. Furthermore each element in V-set for any node can only be derived from the V-set of nodes in exactly two different paths.

Lemma 3 (Computability Lemma). *There exists an explicit efficient procedure that given V_x, W_x computes $V_{x\|0}, W_{x\|0}$ and $V_{x\|1}, W_{x\|1}$.*

Proof. We start by noting that it suffices to show that for each i , given V_x^i, W_x^i one can compute $V_{x\|0}^i, W_{x\|0}^i$ and $V_{x\|1}^i, W_{x\|1}^i$. We argue this next. Observe that two cases arise either $|x| < i$ or $|x| \geq i$. We deal with the two cases:

- $|x| < i$: In this case V_x^i is $\{S_{i,x}, S_{i,(x+1)\|0^{i-|x|}}\}$ and these values can be used to compute $S_{i,x\|0}, S_{i,x\|1}, S_{i,(x\|0)+1} = S_{i,x\|1}$ and $S_{i,((x\|1)+1)\|0^{i-|x|-1}} = S_{i,(x+1)\|0\|0^{i-|x|-1}} = S_{i,(x+1)\|0^{i-|x|}}$. Observe by case by case inspection that these values are sufficient for computing $V_{x\|0}^i, W_{x\|0}^i$ and $V_{x\|1}^i, W_{x\|1}^i$ in all cases.
- $|x| \geq i$: Note that according to the constraints placed on x by the definition, if $V_x^i = \emptyset$ then both $V_{x\|0}^i$ and $V_{x\|1}^i$ must be \emptyset as well. On the other hand if $V_x^i \neq \emptyset$ then $V_{x\|0}^i$ is still \emptyset while $V_{x\|1}^i = V_x^i$. Additionally, $W_{x\|0}^i = W_{x\|1}^i = W_x^i$.

This concludes the proof.

Lemma 4 (Derivability Lemma). *For every $i \in [\kappa], x \in \{0, 1\}^i$ and $x \neq 1^i$ we have that, $S_{i,x+1}$ can be derived from keys in V_y^i if and only if y is a prefix of $x\|1^{\kappa-i}$ or $(x+1)\|1^{\kappa-i}$. Additionally, $S_{i,0^i}$ can be derived from keys in V_y^i if and only if y is a prefix of $0^i\|1^{\kappa-i}$ (Fig. 5).*

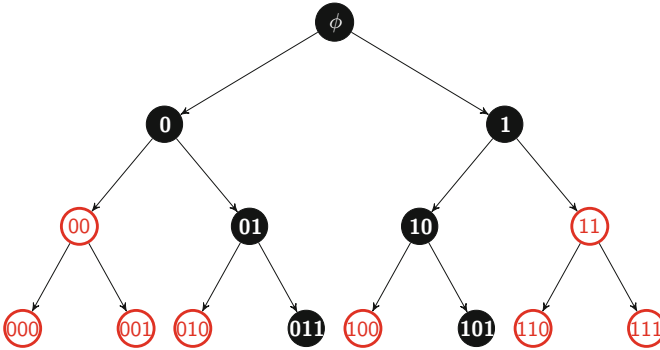


Fig. 5. Black nodes represent the choices of $x \in \{0, 1\}^{\leq 3}$ such that V_x^2 can be used to derive $S_{2,10}$.

Proof. We start by noting that for any $y \in \{0, 1\}^{>i} \cap \{0, 1\}^{\leq \kappa}$, by definition of V-sets we have that $V_y^i = V_{y[i]}^i$ or $V_y^i = \emptyset$. Hence it suffices to prove the above lemma for $y \in \{0, 1\}^{\leq i}$.

We first prove that if y is a prefix of x or $(x+1)$ then we can derive $S_{i,x+1}$ from V_y^i . Two cases arise:

- Observe that if y is a prefix of x then we must have that either y is a prefix of $x+1$ or $x+1 = (y+1)\|0^{i-|y|}$. Next note that by definition of V-sets we have that $V_y^i = \{S_{i,y}, S_{i,(y+1)\|0^{i-|y|}}\}$, and one of these values can be used to compute $S_{i,x+1}$.
- On the other hand if y is a prefix of $x+1$ then again by definition of V-sets we have that $V_y^i = \{S_{i,y}, S_{i,(y+1)\|0^{i-|y|}}\}$, and $S_{i,y}$ can be used to compute $S_{i,x+1}$.

Next we show that no other $y \in \{0, 1\}^{\leq i}$ allows for such a derivation. Note that by definition of V-sets we have that $V_y^i = \{S_{i,y}, S_{i,(y+1)\|0^{i-|y|}}\}$. We will argue that neither $S_{i,y}$ nor $S_{i,(y+1)\|0^{i-|y|}}$ can be used to derive $S_{i,x+1}$.

- We are given that y is not a prefix of $x + 1$. This implies that $S_{i,y}$ cannot be used to derive $S_{i,x+1}$.
- Now we need to argue that $S_{i,(y+1)\|0^{i-|y|}}$ cannot be used to compute $S_{i,x+1}$. For this, it suffices to argue that $x + 1 \neq (y + 1)\|0^{i-|y|}$. If $x + 1 = (y + 1)\|0^{i-|y|}$ then y must be prefix of x . However, we are given that this is not the case. This proves our claim.

The argument for the value $S_{i,0^i}$ follows analogously. This concludes the proof.

5.2 Hard on Average SVL Instances

In this section, we describe our construction for hard on average instance of SVL. In particular, we describe our sampler that samples hard on average instances $(x_s, \text{Succ}, \text{Ver}, 1^\kappa)$. Here x_s is the source node and Succ is the successor circuit. We define a directed edge between u and v if and only if $\text{Succ}(u) = v$. Ver is the verification circuit and is used to test whether a given node is the k^{th} node from x_s . That is, $\text{Ver}(x, k) = 1$ iff $x = \text{Succ}^{k-1}(x_s)$. For the generated instances, we argue that it is hard to find the 1^κ node in the path.

In our construction we use a selectively secure functional encryption scheme (FE.Setup, FE.KeyGen, FE.Enc, FE.Dec), a prefix-puncturable PRF (Definition 11), a semantically secure symmetric key encryption (SK.KeyGen, SK.Enc, SK.Dec) and injective PRGs having the left half injectivity property Definition 5. PRG_0 and PRG_1 denote the left and the right part of the output of this PRG.

The formal description of hard on average SVL instance sampler is provided in Fig. 6. Internally this sampler generates the successor circuit to include functional encryption secret keys for circuits provided in Fig. 7. Next we informally describe the SVL instances considered.

A sampled instance implicitly defines a line graph where each node in the graph is of the form $(x, \sigma_1, \dots, \sigma_\kappa)$ where $\sigma_i = \text{PrefixPunc}(S_i, x_{[i]})$ for all $i \in [\kappa]$. We say a node is *valid* if the above condition holds. The node $(x, \sigma_1, \dots, \sigma_\kappa)$ is connected to $(x + 1, \sigma'_1, \dots, \sigma'_\kappa)$ by an outgoing edge and to $(x - 1, \sigma''_1, \dots, \sigma''_\kappa)$ by an incoming edge. The successor circuit on input $(x, \sigma_1, \dots, \sigma_\kappa)$ checks for the validity of the node and if the node is valid it outputs $(x + 1, \sigma'_1, \dots, \sigma'_\kappa)$. The verification circuit on input $(x, \sigma_1, \dots, \sigma_\kappa, j)$ outputs if and only if $x = j - 1$ and $(x, \sigma_1, \dots, \sigma_\kappa)$ is valid.

We now explain how the successor circuit works. The successor circuit is described by a sequence of $\kappa + 1$ secret keys $\text{FSK}_1, \dots, \text{FSK}_{\kappa+1}$ for appropriate functions. These keys are generated corresponding to independent instances of functional encryption. Along with the keys the successor circuit also contains a ciphertext c_ϕ that encrypts the empty string, ϕ , under PK_1 along with the key values V_ϕ and W_ϕ . Intuitively, the function key FSK_i corresponds to a function

- **Sampled Ingredients:**
 1. Sample $\{S_i\}_{i \in [\kappa]}$ and K_ϕ from $\text{KeyGen}_{\mathcal{P}\mathcal{P}\mathcal{R}\mathcal{F}}(1^\kappa)$. Here S_i 's is a key that works for i bit inputs, namely $\text{PPRF}_{S_i} : \{0, 1\}^i \rightarrow \{0, 1\}^\kappa$ for all $i \in [\kappa]$. Similarly, K_ϕ works on inputs of length $\text{rand}(\kappa)$ where $\text{rand}(\cdot)$ would be specified later. Initialize $V_\phi^i = S_i$, $V_\phi = \bigcup_{i \in [\kappa]} V_\phi^i$ and $W_\phi = \emptyset$.
 2. Sample $(PK_i, MSK_i) \leftarrow \text{FE.Setup}(1^\kappa)$ for all $1 \leq i \leq \kappa + 1$.
 3. Sample $sk \leftarrow \text{SK.KeyGen}(1^\kappa)$ and let $\Pi \leftarrow \text{SK.Enc}_{sk}(\pi)$ and $\Lambda \leftarrow \text{SK.Enc}_{sk}(\lambda)$ where $\pi = 0^{\ell(\kappa)}$ and $\lambda = 0^{\ell'(\kappa)}$. Here $\ell(\cdot)$ and $\ell'(\cdot)$ are appropriate length functions specified later.
 4. Sample $v \leftarrow \{0, 1\}^{2\kappa}$.
- **Functional encryption ciphertext and keys to simulate obfuscation:**
 1. For each $i \in [\kappa]$ generate $\text{FSK}_i \leftarrow \text{FE.KeyGen}(MSK_i, F_i, PK_{i+1}, \Pi)$ and $\text{FSK}_{\kappa+1} \leftarrow \text{FE.KeyGen}(MSK_{\kappa+1}, G_{v,\Lambda})$, where F_i, PK_{i+1}, Π and $G_{v,\Lambda}$ are circuits described in Figure 7.
 2. Let $c_\phi = \text{FE.Enc}_{PK_1}(\phi, V_\phi, W_\phi, 0^\kappa, 0)$
- **Source node:** The source node x_s is given by $(0^\kappa, \sigma_1, \dots, \sigma_\kappa)$ where $\sigma_i = \text{PPRF}_{S_i}(0^i)$ for all $i \in [\kappa]$.
- **Successor Circuit:** The successor circuit Succ in our setting takes as input $x, \sigma_1, \dots, \sigma_\kappa$ and outputs $x + 1, \sigma'_1, \dots, \sigma'_\kappa$ if the associated signatures $\sigma_1, \dots, \sigma_\kappa$ are valid. It proceeds as follows:
 1. For $i \in [\kappa]$ compute $c_{x_{[i-1]}\|0}, c_{x_{[i-1]}\|1} := \text{FE.Dec}(\text{FSK}_i, c_{x_{[i-1]}})$.
 2. Obtain $d_x = ((\alpha_1, \dots, \alpha_\kappa), (\beta_j, \dots, \beta_\kappa))$ as output of $\text{FE.Dec}(\text{FSK}_{\kappa+1}, c_x)$. Here $j = f(x)$ where $f(x)$ is the smallest j such that $x = x_{[j]}\|1^{\kappa-j}$.
 3. Output \perp if $\text{PRG}_0(\sigma_i) \neq \alpha_i$ for any $i \in [\kappa]$ or if $d_x = \perp$.
 4. If $x = 1^\kappa$, output SOLVED.
 5. For each $i \in [j - 1]$ set $\sigma'_i = \sigma_i$.
 6. For each $i \in \{j, \dots, \kappa\}$ set $\gamma_i = \text{PRG}_1(\sigma_i)$ and σ'_i as $\text{SK.Dec}_{\gamma_j, \dots, \gamma_\kappa}(\beta_i)$, decrypting β_i encrypted under $\gamma_j, \dots, \gamma_\kappa$.
 7. Output $(x + 1, \sigma'_1, \dots, \sigma'_\kappa)$.
- **Verification Circuit:** The verification circuit Ver on input $x, \sigma_1, \dots, \sigma_\kappa, j$ outputs 1 if Succ on input $x, \sigma_1, \dots, \sigma_\kappa$ doesn't output \perp and $x = j - 1$ and 0 otherwise.

Fig. 6. Hard on average instance for SVL based on hardness of FE.

F_i that takes as input a binary string x of length i and outputs an encryption of $x\|0$ and $x\|1$ under PK_{i+1} . Additionally these ciphertexts, in addition to $x\|0$ and $x\|1$, also contain key values $V_{x\|0}, W_{x\|0}$ and $V_{x\|1}, W_{x\|1}$ respectively. Recall from Sect. 5.1 that the keys in these sets are used to test validity of signatures provides as input and to generate the new ones.

The successor circuit on an input of the form $(x, \sigma_1, \dots, \sigma_\kappa)$ does the following. It first obtains an encryption of x along with key values V_x and W_x under the public key $PK_{\kappa+1}$. This is done as follows. Start with c_ϕ and decrypt it using key FSK_1 to obtain encryptions of 0 and 1. Choose one of them based on which one is a prefix of x and continue the process. Repeating this process

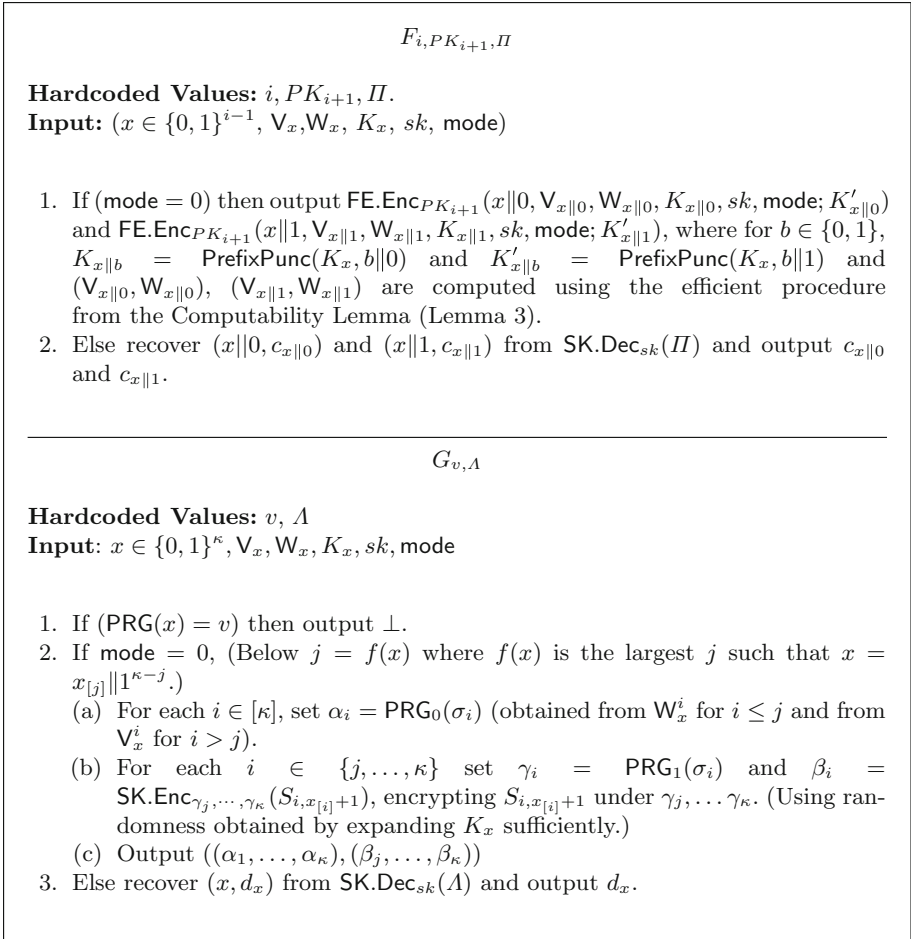


Fig. 7. Circuits for which functional encryption secret keys are given out.

κ times results in the desired ciphertext. Next decrypt the obtained ciphertext using $\text{FSK}_{\kappa+1}$ and it provides some information essential for checking validity of provided input signatures and additional information to generate the signatures for the next node. More details are provided in Figs. 6 and 7.

Setting $\text{rand}(\cdot)$ We set $\text{rand}(\kappa) = 2\kappa + r(\kappa)$ where $r(\kappa)$ is the maximum number of random bits used for generating encryptions of $S_{i, x_{[i]}+1}$ under $\gamma_j, \dots, \gamma_\kappa$ for every $i \in [j, \kappa]$.

Due to space constraints, we defer the proof of hardness of the sampled SVL instance to the full version of the paper [GPS15].

Acknowledgements. The first author would like to thank Sidharth Telang for useful discussions on related topics. Research supported in part from DARPA Safeware Award W911NF15C0210, AFOSR Award FA9550-15-1-0274, and NSF CRII Award 1464397. The views expressed are those of the author and do not reflect the official policy or position of the Department of Defense, the National Science Foundation, or the U.S. Government.

References

- [AB15] Applebaum, B., Brakerski, Z.: Obfuscating circuits via composite-order graded encoding. In: Dodis, Y., Nielsen, J.B. (eds.) TCC 2015, Part II. LNCS, vol. 9015, pp. 528–556. Springer, Heidelberg (2015)
- [ABSV15] Ananth, P., Brakerski, Z., Segev, G., Vaikuntanathan, V.: From selective to adaptive security in functional encryption. In: Gennaro, R., Robshaw, M. (eds.) CRYPTO 2015. LNCS, vol. 9216, pp. 657–677. Springer, Heidelberg (2015)
- [AJ15] Ananth, P., Jain, A.: Indistinguishability obfuscation from compact functional encryption. In: Gennaro, R., Robshaw, M.J.B. (eds.) CRYPTO 2015. LNCS, vol. 9215, pp. 308–326. Springer, Heidelberg (2015)
- [AJS15] Ananth, P., Jain, A., Sahai, A.: Achieving compactness generically: Indistinguishability obfuscation from non-compact functional encryption. IACR Cryptology ePrint Archive, 2015:730 (2015)
- [AKV04] Abbot, T., Kane, D., Valiant, P.: On Algorithms for Nash Equilibria (2004). <http://web.mit.edu/tabbott/Public/final.pdf>
- [BCC+14] Bitansky, N., Canetti, R., Cohn, H., Goldwasser, S., Kalai, Y.T., Paneth, O., Rosen, A.: The impossibility of obfuscation with auxiliary input or a universal simulator. In: Garay, J.A., Gennaro, R. (eds.) CRYPTO 2014, Part II. LNCS, vol. 8617, pp. 71–89. Springer, Heidelberg (2014)
- [BGI+12] Barak, B., Goldreich, O., Impagliazzo, R., Rudich, S., Sahai, A., Vadhan, S.P., Yang, K.: On the (im)possibility of obfuscating programs. J. ACM **59**(2), 6 (2012)
- [BGI14] Boyle, E., Goldwasser, S., Ivan, I.: Functional signatures and pseudorandom functions. In: Krawczyk, H. (ed.) PKC 2014. LNCS, vol. 8383, pp. 501–519. Springer, Heidelberg (2014)
- [BGK+14] Barak, B., Garg, S., Kalai, Y.T., Paneth, O., Sahai, A.: Protecting obfuscation against algebraic attacks. In: Nguyen, P.Q., Oswald, E. (eds.) EUROCRYPT 2014. LNCS, vol. 8441, pp. 221–238. Springer, Heidelberg (2014)
- [BPR15] Bitansky, N., Paneth, O., Rosen, A.: On the cryptographic hardness of finding a nash equilibrium. In: FOCS (2015)
- [BR14] Brakerski, Z., Rothblum, G.N.: Virtual black-box obfuscation for all circuits via generic graded encoding. In: Lindell, Y. (ed.) TCC 2014. LNCS, vol. 8349, pp. 1–25. Springer, Heidelberg (2014)
- [BSW11] Boneh, D., Sahai, A., Waters, B.: Functional encryption: definitions and challenges. In: Ishai, Y. (ed.) TCC 2011. LNCS, vol. 6597, pp. 253–273. Springer, Heidelberg (2011)
- [BV15a] Bitansky, N., Vaikuntanathan, V.: Indistinguishability obfuscation from functional encryption. In: 56th FOCS, pp. 171–190. IEEE Computer Society Press (2015)

- [BV15b] Bitansky, N., Vaikuntanathan, V.: Indistinguishability obfuscation from functional encryption. *IACR Cryptology ePrint Archive*, 2015:163 (2015)
- [BW13] Boneh, D., Waters, B.: Constrained pseudorandom functions and their applications. In: Sako, K., Sarkar, P. (eds.) *ASIACRYPT 2013, Part II*. LNCS, vol. 8270, pp. 280–300. Springer, Heidelberg (2013)
- [CDT09] Chen, X., Deng, X., Teng, S.-H.: Settling the complexity of computing two-player nash equilibria. *J. ACM* **56**(3), 1–57 (2009)
- [DGP09] Daskalakis, C., Goldberg, P.W., Papadimitriou, C.H.: The complexity of computing a nash equilibrium. *Commun. ACM* **52**(2), 89–97 (2009)
- [GGH13a] Garg, S., Gentry, C., Halevi, S.: Candidate multilinear maps from ideal lattices. In: Johansson, T., Nguyen, P.Q. (eds.) *EUROCRYPT 2013*. LNCS, vol. 7881, pp. 1–17. Springer, Heidelberg (2013)
- [GGH+13b] Garg, S., Gentry, C., Halevi, S., Raykova, M., Sahai, A., Waters, B.: Candidate indistinguishability obfuscation and functional encryption for all circuits. In: *54th FOCS*, pp. 40–49, Berkeley, CA, USA. IEEE Computer Society Press, 26–29 October 2013
- [GGHZ16] Garg, S., Gentry, C., Halevi, S., Zhandry, M.: Fully secure functional encryption from multilinear maps. In: *TCC* (2016)
- [GGM86] Goldreich, O., Goldwasser, S., Micali, S.: How to construct random functions. *J. ACM* **33**(4), 792–807 (1986)
- [GK05] Goldwasser, S., Kalai, Y.T.: On the impossibility of obfuscation with auxiliary input. In: *FOCS*, pp. 553–562 (2005)
- [GL89] Goldreich, O., Levin, L.A.: A hard-core predicate for all one-way functions. In: *Proceedings of the 21st Annual ACM Symposium on Theory of Computing*, Seattle, Washington, USA, pp. 25–32, 14–17 May 1989
- [GLSW15] Gentry, C., Lewko, A.B., Sahai, A., Waters, B.: Indistinguishability obfuscation from the multilinear subgroup elimination assumption. In: *56th FOCS*, pp. 151–170. IEEE Computer Society Press (2015)
- [GMS16] Garg, S., Mukherjee, P., Srinivasan, A.: Obfuscation without the vulnerabilities of multilinear maps. *IACR Cryptology ePrint Archive*, 2016:390 (2016)
- [GPS15] Garg, S., Pandey, O., Srinivasan, A.: On the exact cryptographic hardness of finding a nash equilibrium. *Cryptology ePrint Archive*, Report 2015/1078 (2015). <http://eprint.iacr.org/2015/1078>
- [GPSZ16] Garg, S., Pandey, O., Srinivasan, A., Zhandry, M.: Breaking the sub-exponential barrier in obfuscation. *Cryptology ePrint Archive*, Report 2016/102 (2016). <http://eprint.iacr.org/2016/102>
- [GS16] Garg, S., Srinivasan, A.: Unifying security notions of functional encryption. *Cryptology ePrint Archive*, Report 2016/524 (2016). <http://eprint.iacr.org/>
- [HY16] Hubáček, P., Yogev, E.: Hardness of continuous local search: query complexity and cryptographic lower bounds. *Electron. Colloquium Comput. Complex. (ECCC)* **23**, 63 (2016)
- [Jer12] Emil Jerábek. Integer factoring and modular square roots. *CoRR* abs/1207.5220 (2012)
- [KPTZ13] Kiayias, A., Papadopoulos, S., Triandopoulos, N., Zacharias, T.: Delegatable pseudorandom functions and applications. In: *2013 ACM SIGSAC Conference on Computer and Communications Security, CCS 2013*, Berlin, Germany, pp. 669–684, 4–8 November 2013

- [MP91] Megiddo, N., Papadimitriou, C.H.: On total functions, existence theorems and computational complexity. *Theor. Comput. Sci.* **81**(2), 317–324 (1991)
- [Nao03] Naor, M.: On cryptographic assumptions and challenges. In: Boneh, D. (ed.) *CRYPTO 2003*. LNCS, vol. 2729, pp. 96–109. Springer, Heidelberg (2003)
- [Nas51] Nash, J.: Non-cooperative games. *Ann. Math.* **54**(2), 286–295 (1951)
- [O’N10] O’Neill, A.: Definitional issues in functional encryption. *IACR Cryptology ePrint Archive*, 2010:556 (2010)
- [Pap94] Papadimitriou, C.H.: On the complexity of the parity argument and other inefficient proofs of existence. *J. Comput. Syst. Sci.* **48**(3), 498–532 (1994)
- [PST14] Pass, R., Seth, K., Telang, S.: Indistinguishability obfuscation from semantically-secure multilinear encodings. In: Garay, J.A., Gennaro, R. (eds.) *CRYPTO 2014, Part I*. LNCS, vol. 8616, pp. 500–517. Springer, Heidelberg (2014)
- [RSS16] Rosen, A., Segev, G., Shahaf, I.: Can PPAD hardness be based on standard cryptographic assumptions? *Electron. Colloquium Comput. Complex. (ECCC)* **23**, 59 (2016)
- [SW14] Sahai, A., Waters, B.: How to use indistinguishability obfuscation: deniable encryption, and more. In: *Symposium on Theory of Computing, STOC 2014, New York, NY, USA, 31 May–03 June 2014*, pp. 475–484 (2014)
- [Zim15] Zimmerman, J.: How to obfuscate programs directly. In: Oswald, E., Fischlin, M. (eds.) *EUROCRYPT 2015*. LNCS, vol. 9057, pp. 439–467. Springer, Heidelberg (2015)