Be Adaptive, Avoid Overcommitting

Zahra Jafargholi[1(✉)], Chethan Kamath[2], Karen Klein[2], Ilan Komargodski[3], Krzysztof Pietrzak[2], and Daniel Wichs[4]

[1] Aarhus University, Aarhus, Denmark
zahra@au.cs.dk
[2] IST Austria, Am Campus 1, 3400 Klosterneuburg, Austria
{ckamath,karen.klein,pietrzak}@ist.ac.at
[3] Weizmann Institute of Science, 76100 Rehovot, Israel
ilan.komargodski@weizmann.ac.il
[4] Northeastern University, Boston, USA
wichs@ccs.neu.edu

Abstract. For many cryptographic primitives, it is relatively easy to achieve *selective security* (where the adversary commits a-priori to some of the choices to be made later in the attack) but appears difficult to achieve the more natural notion of *adaptive security* (where the adversary can make all choices on the go as the attack progresses). A series of several recent works shows how to cleverly achieve adaptive security in several such scenarios including *generalized selective decryption* (Panjwani, TCC '07 and Fuchsbauer et al., CRYPTO '15), *constrained PRFs* (Fuchsbauer et al., ASIACRYPT '14), and *Yao garbled circuits* (Jafargholi and Wichs, TCC '16b). Although the above works expressed vague intuition that they share a common technique, the connection was never made precise. In this work we present a new framework that connects all of these works and allows us to present them in a unified and simplified fashion. Moreover, we use the framework to derive a new result for adaptively secure *secret sharing over access structures defined via monotone circuits*. We envision that further applications will follow in the future.

Underlying our framework is the following simple idea. It is well known that selective security, where the adversary commits to n-bits of information about his future choices, automatically implies adaptive security at the cost of amplifying the adversary's advantage by a factor of up to 2^n. However, in some cases the proof of selective security proceeds via a sequence of hybrids, where each pair of adjacent hybrids locally only

Z. Jafargholi—Supported by The Danish National Research Foundation and The National Science Foundation of China (under the grant 61061130540) for the Sino-Danish Center for the Theory of Interactive Computation, within which part of this work was performed; and by the Advanced ERC grant MPCPRO.

C. Kamath—Supported by the European Research Council, ERC consolidator grant (682815 - TOCNeT).

I. Komargodski—Supported by a Levzion fellowship and by a grant from the Israel Science Foundation.

D. Wichs—Research supported by NSF grants CNS-1314722, CNS-1413964.

J. Katz and H. Shacham (Eds.): CRYPTO 2017, Part I, LNCS 10401, pp. 133–163, 2017.
DOI: 10.1007/978-3-319-63688-7_5

requires some smaller partial information consisting of $m \ll n$ bits. The partial information needed might be completely different between different pairs of hybrids, and if we look across all the hybrids we might rely on the entire n-bit commitment. Nevertheless, the above is sufficient to prove adaptive security, at the cost of amplifying the adversary's advantage by a factor of only $2^m \ll 2^n$.

In all of our examples using the above framework, the different hybrids are captured by some sort of a *graph pebbling game* and the amount of information that the adversary needs to commit to in each pair of hybrids is bounded by the maximum number of pebbles in play at any point in time. Therefore, coming up with better strategies for proving adaptive security translates to various pebbling strategies for different types of graphs.

1 Introduction

Many security definitions come in two flavors: a stronger "adaptive" flavor, where the adversary can arbitrarily make various choices during the course of the attack, and a weaker "selective" flavor where the adversary must commit to some or all of his choices a-priori. For example, in the context of *identity-based encryption*, selective security requires the adversary to decide on the identity of the attacked party at the very beginning of the game whereas adaptive security allows the attacker to first see the master public key and some secret keys before making this choice. Often, it appears to be much easier to achieve selective security than it is to achieve adaptive security.

A series of recent works achieves adaptive security in several such scenarios where we previously only knew how to achieve selective security: *generalized selective decryption (GSD)* [8,23], *constrained PRFs* [9], and *garbled circuits* [16]. Although some of these works suggest a vague intuition that there is a general technique at play, there was no attempt to make this precise and to crystallize what the technique is or how these results are connected. In this work we present a new framework that connects all of these works and allows us to present them in a unified and simplified fashion. Moreover, we use the framework to derive a new result for adaptively secure secret sharing over access structures defined via monotone circuits.

At a high level, our framework carefully combines two basic tools commonly used throughout cryptography: *random guessing* (of the adaptive choices to be made by the adversary)[1] and *the hybrid argument*. Firstly, "random guessing" gives us a generic way to qualitatively upgrade selective security to adaptive security at a quantitative cost in the amount of security. In particular, assume

[1] In many previous works – including [8,9,16], and by the authors of this paper – this random guessing was referred to as "complexity leveraging", but this seems to be an abuse of the term. Instead, complexity leveraging [7] refers to the use of two different schemes, S_1, S_2, where the two schemes are chosen with different values of the security parameter, k_1 and k_2, where $k_1 < k_2$ and such that an adversary against S_2 (or perhaps even the honest user of S_2) can break the security of S_1.

we can prove the security of a selective game where the adversary commits to n-bits of information about his future choices. Then, we can also prove adaptive security by guessing this commitment and taking a factor of 2^n loss in the security advantage. However, this quantitative loss is often too high and hence we usually wish to avoid it or at least lower it. Secondly, the hybrid argument allows us to prove the indistinguishability of two games G_L and G_R by defining a sequence of hybrid games $\mathsf{G}_\mathsf{L} \equiv \mathsf{H}_0, \mathsf{H}_1, \ldots, \mathsf{H}_\ell \equiv \mathsf{G}_\mathsf{R}$ and showing that each pair of neighboring hybrids H_i and H_{i+1} are indistinguishable.

Our Framework. Our framework starts with two adaptive games G_L and G_R that we wish to show indistinguishable but we don't initially have any direct way of doing so. Let H_L and H_R be selective versions of the two games respectively, where the adversary initially has to commit to some information $w \in \{0,1\}^n$ about his future choices. Furthermore, assume there is some sequence of selective hybrids $\mathsf{H}_\mathsf{L} = \mathsf{H}_0, \mathsf{H}_1, \ldots, \mathsf{H}_\ell \equiv \mathsf{H}_\mathsf{R}$ such that we can show that H_i and H_{i+1} are indistinguishable. A naïve combination of the hybrid argument and random guessing shows that G_L and G_R are indistinguishable at a factor of $2^n \cdot \ell$ loss in security, but we want to do better.

Recall that the hybrids H_i are selective and require the adversary to commit to w. However, it might be the case that for each i we can prove that H_i and H_{i+1} would be indistinguishable even if the adversary didn't have to commit to all of w but only some partial-information $h_i(w) \in \{0,1\}^m$ for $m \ll n$ (formalizing this condition precisely requires great care and is the major source of subtlety in our framework). Notice that the partial information that we need to know about w may be completely different for different pairs of hybrids, and if we look across all hybrids then we may need to know all of w. Nevertheless, we prove that this suffices to show that the adaptive games G_L and G_R are indistinguishable with only a $2^m \cdot \ell \ll 2^n \cdot \ell$ loss of security.

Applications of Our Framework. We show how to understand all of the prior works mentioned above as applications of our framework. In many cases, this vastly simplifies prior works. We also use the framework to derive a new result, proving the adaptive security of Yao's secret sharing scheme for access structures defined via monotone circuits.

In all of the examples, we get a series of selective hybrids $\mathsf{H}_1, \ldots, \mathsf{H}_\ell$ that correspond to *pebbling configurations* in some graph pebbling game. The amount of information needed to show that neighboring hybrids H_i and H_{i+1} are indistinguishable only depends on the configuration of the pebbles in the i'th step of the game. Therefore, using our framework, we translate the problem of coming up with adaptive security proofs to the problem of coming up with pebbling strategies that only require a succinct representation of each pebbling configuration.

We now proceed to give a high level overview of each of our results applying our general framework to specific problems, and refer to the main body for technical details.

1.1 Adaptive Secret Sharing for Monotone Circuits

Secret sharing schemes, introduced by Blakley [4] and Shamir [27], are methods that enable a dealer, that has a secret piece of information, to distribute this secret among n parties such that a "qualified" subset of parties has enough information to reconstruct the secret while any "unqualified" subset of parties learns nothing about the secret. The monotone collection of "qualified" subsets is known as an *access structure*. Any access structure admits a secret sharing scheme but the share size could be exponential in n [14]. We are interested in efficient schemes in which the share size is polynomial (in n and possibly in a security parameter).

Many of the classical schemes for secret sharing are *perfectly* (information theoretically) secure. The largest class of access structures that admit such a (perfect and efficient) scheme was obtained by Karchmer and Wigderson [18] for the class of all functions that can be computed by monotone span programs. This result generalized a previous work of Benaloh and Leichter [3] (which, in turn, improved a result of Ito et al. [14]) that showed the same result but for a smaller class of access structures: those functions that can be computed by monotone Boolean formulas. Under cryptographic hardness assumptions, efficient schemes for more general access structures are known (but security is only for bounded adversaries). In particular, in an unpublished work (mentioned in [1], see also Vinod et al. [28]), Yao showed how to realize schemes for access structures that are described by monotone *circuits*. This construction could be used for access structures which are known to be computed by monotone circuits but are not known to be computed by monotone span programs, e.g., directed connectivity [17,24].[2] Komargodski et al. [21] showed how to realize the class of access structures described by monotone functions in NP[3] under the assumption that witness encryption for NP [10] and one-way functions exist.[4,5]

Selective vs. Adaptive Security. All of the schemes described above guarantee security against static adversaries, where the adversary chooses a subset of parties it controls before it sees any of the shares. A more natural security guarantee would be to require that even an adversary that chooses its set of parties in an *adaptive* manner (i.e., based on the shares it has seen so far) is unable to learn the secret (or any partial information about it).

It is known that the schemes that satisfy perfect security (including the works [3,14,18] mentioned above) actually satisfy this stronger notion of adaptive

[2] In the access structure for directed connectivity, the parties correspond to an edge in the complete *directed* graph and the "qualified" subsets are those edges that connect two distinguished nodes s and t.

[3] For access structures in NP, a qualified set of parties needs to know an NP witness that they are qualified.

[4] Witness encryption for a language $L \in \mathsf{NP}$ allows to encrypt a message relative to a statement $x \in L$ such that anyone holding a witness to the statement can decrypt the message, but if $x \notin L$, then the message is computationally hidden.

[5] One can relax the additional assumption of one-way functions to an average-case hardness assumption in NP [20].

security. However, the situation for the schemes that are based on cryptographic assumptions (including Yao's scheme and the scheme of [21]) is much less clear. Using random guessing (see Lemma 1) it can be shown that these schemes are adaptively secure, but this reduction loses an exponential (in the number of parties) factor in the security of the scheme. Additionally, as noted in [21], their scheme can be shown to be adaptively secure if the witness encryption scheme is *extractable*.[6] The latter is a somewhat controversial assumption that we prefer to avoid.

Our Results. We analyze the adaptive security of Yao's scheme under our framework and show that in some cases the security loss is much smaller than 2^n. Roughly, we show that if the access structure can be described by a monotone circuit of depth d and s gates (with unbounded fan-in and fan-out) the security loss is proportional to $s^{O(d)}$. Thus, for shallow circuits our analysis shows that an exponential loss is avoidable.

To exemplify the usefulness of the result, consider, for instance, the directed st-connectivity access structure mentioned in Footnote 6. It is known that it can be computed by a monotone circuit of size $O(n^3 \log n)$ and depth $O(\log^2 n)$, but its monotone formula and span-program complexity is $2^{\Omega(\log^2 n)}$ [17,24]. Thus, no efficient perfectly secure scheme is known, and our proof shows that Yao's scheme for this access structure is secure based on the assumption that quasi-polynomially-secure one-way functions exist.

Yao's Scheme. In this scheme, an access structure is described by a monotone circuit. The sharing procedure first labels the output wire of the circuit with the shared secret and then proceeds to assign labels to all wires of the circuit; in the end the label on each input wire is included in the share of the corresponding party. The procedure for assigning labels is recursive and in each step it labels the input wires of a gate g assuming its output wires are already labeled (recall that we assume unbounded fan-in and fan-out so there are many input and output wires). To do so, we first sample a fresh encryption key s for a symmetric-key encryption scheme. If the gate is an AND gate, then we label each input wire with a random string conditioned on their XOR being s, and if the gate is an OR gate, then we label each input wire with s. In either case, we encrypt the labels of the output wires under s and include these ciphertexts associated with the gate g as part of ever party's share. The reconstruction of the scheme works by reversing the above procedure from the leaves to the root. This scheme is indeed efficient for access structures that have polynomial-size monotone circuits.

Security Proof. Our goal is to show that as long as an adversary controls an unqualified set, he cannot learn anything about the secret. We start by outlining the selective security proof (following the argument of [28]), where the adversary first commits to the "corrupted" set. The proof is via a series of hybrids in

[6] This is a knowledge assumption that says that if an adversary can decrypt a witness encryption ciphertext, then it must *know* a witness which can be extracted from it.

which we slowly replace the ciphertexts associated with various gates g with bogus ciphertexts. Once we do this for the output gate, the shares become independent of the secret which proves security. The gates for which we can replace the ciphertexts with bogus ones are the gates for which the adversary cannot compute the corresponding encryption key. Since the adversary controls an unqualified set, a sequence which eventually results with replacing the encryption of the root gate must exist. Since in every hybrid we "handle" one gate and never consider it again, the number of hybrids is at most the number of gates in the circuit.

The problem with lifting this proof to the adaptive case is that it seems inherent to know the corrupted set of parties in order to know for which gates g to switch the ciphertexts from real to bogus (and in what order). However, in the adaptive game this set is not known during the sharing procedure. A naïve use of random guessing would result in an exponential security loss 2^n, where n is the number of parties.

To overcome this we associate each intermediate hybrid H_i with a *pebbling configuration* in which each gate in the circuit is either pebbled (ciphertexts are bogus) or unpebbled (ciphertexts are real). The pebbling rules are:

1. Can place or remove a pebble on any AND gate for which (at least) one input wire is either *not* corrupted or comes out of a gate with a pebble on it.
2. Can place or remove a pebble on any OR gate for which all of the incoming wires are either non-corrupted input wires or come out of gates all of which have pebbles on them.

The initial hybrid corresponds to the case in which all gates are unpebbled and the final hybrid corresponds to the case in which all gates are unpebbled except the root gate which has a pebble. Now, any pebbling strategy that takes us from the initial configuration to the final one, corresponds to a sequence of selective hybrids H_i. Furthermore, to prove indistinguishability of neighboring hybrids $\mathsf{H}_i, \mathsf{H}_{i+1}$ we don't need the adversary to commit to the entire set of corrupted parties ahead of time but it suffices if the adversary only commits to the pebble configuration in steps i and $i+1$. Therefore, if the pebbling strategy has the property that each configuration requires few bits to describe, then we would be able to use our framework. We show that for every corrupted set and any monotone circuit of depth d and s gates, there exists such a pebbling strategy, where the number of moves is roughly $2^{O(d)}$ and each configuration has a very succinct representation: roughly $d \cdot \log s$ bits. Plugging this into our framework, we get a proof of adaptive security with security loss proportional to $s^{O(d)}$. We refer to Sect. 4 for the precise details.

1.2 Generalized Selective Decryption

Generalized Selective Decryption (GSD), introduced by Panjwani [23], is a game that captures the difficulty of proving adaptive security of certain protocols, most notably the Logical Key Hierarchy (LKH) multicast encryption protocol. On a

high level, it deals with scenario where we have many secret keys k_i and various ciphertexts encrypting one key under another (but no cycles). We will discuss this problem in depth in the full version [15], here giving a high level overview on how our framework applies to this problem.

Let $(\mathsf{Enc}, \mathsf{Dec})$ be a CPA-secure symmetric encryption scheme with (probabilistic) $\mathsf{Enc}\colon \mathcal{K} \times \mathcal{M} \to \mathcal{C}$ and $\mathsf{Dec}\colon \mathcal{K} \times \mathcal{C} \to \mathcal{M}$. We assume $\mathcal{K} \subseteq \mathcal{M}$, i.e., we can encrypt keys. In the game, the challenger—either G_L or G_R—picks $n+1$ random keys $k_0, \ldots, k_n \in \mathcal{K}$, and the adversary A is then allowed to make three types of queries:[7]

- Encryption query: on input $(\texttt{encrypt}, i, j)$ receives $\mathsf{Enc}(k_i, k_j)$.
- Corruption queries: on input $(\texttt{corrupt}, i)$ receives k_i.
- Challenge query, only one is allowed: on input $(\texttt{challenge}, i)$ receives k_i in the real game G_L, and a random value in the random game G_R.

We think of this game as generating a directed graph, with vertex set $\mathcal{V} = \{0, \ldots, n\}$, where every $(\texttt{encrypt}, i, j)$ query adds a directed edge (i, j), and we say a vertex v_i is corrupted if a query $(\texttt{corrupt}, i)$ was made, or v_i can be reached from a corrupted vertex. The goal of the adversary is to distinguish the games G_L or G_R, with the restriction that the constructed graph has no cycles, and the challenge vertex is a sink. To prove security, i.e., reduce the indistinguishability of G_L or G_R to the security of Enc, we can consider a selectivized version of this game where A must commit to the graph as described above (which uses $<n^2$ bits). The security of this selectivized game can then be reduced to the security of Enc by a series of $<n^2$ hybrids, where a distinguisher for any two consecutive hybrids can be used to break the security of Enc with the same advantage. Using random guessing followed by a hybrid argument we conclude that if Enc is δ-secure, the GSD game is $\delta \cdot n^2 \cdot 2^{n^2}$-secure. Thus, we lose an exponential in n^2 factor in the reduction.

Fortunately, if we look at the actual protocols that GSD is supposed to capture, it turns out that the graphs that A can generate are not totally arbitrary. Two interesting cases are given by GSD restricted to graphs of bounded depth, and to trees. For these cases better reductions exist. Panjwani [23] shows that if the adversary is restricted to play the game such that the resulting graph is of depth at most d, a reduction losing a factor $(2n)^d$ exists. Moreover, Fuchsbauer et al. [8] give a reduction losing a factor $n^{3 \log n}$ when the underlying graph is a tree. In the full version we prove these results in our framework. Our proofs are much simpler than the original ones, especially than the proof of [23] which is very long and technical. This is thanks to our modular approach, where our general framework takes care of delicate probabilistic arguments, and basically just leaves us with the task of designing pebbling strategies, where each pebbling configuration has a succinct description, for various graphs, which is a clean combinatorial problem. The generic connection between adaptive security proofs of the GSD problem and graph pebbling is entirely new to this work.

[7] In the actual game the adversary can also make standard CPA encryption queries $\mathsf{Enc}(k_i, m)$ for chosen m, i. As this doesn't meaningfully change the security proof we ignore this here.

GSD on a Path. Let us sketch the proof idea for the [8] result, but for an even more restricted case where the graph is a path visiting every node exactly once. In other words there is a permutation σ over $\{0,\ldots,n\}$ and the adversary's queries are of the form $(\texttt{encrypt}, \sigma(i-1), \sigma(i))$ and $(\texttt{challenge}, \sigma(n))$. We first consider the selective game where A must commit to this permutation σ ahead of time. Let $\mathsf{H}_\mathsf{L}, \mathsf{H}_\mathsf{R}$ be the selectivized versions of G_L, G_R respectively.

To prove selective security, we can define a sequence of hybrid games $\mathsf{H}_\mathsf{L} = \mathsf{H}_0, \ldots, \mathsf{H}_\ell = \mathsf{H}_\mathsf{R}$. Each hybrid is defined by a path, $0 \to 1 \to \ldots \to n$, with a subset of the edges holding a black pebble. In the hybrid games, a pebble on $(i, i+1)$ means that instead of answering the query $(\texttt{encrypt}, \sigma(i), \sigma(i+1))$ with the "real" answer $\mathsf{Enc}(k_{\sigma(i)}, k_{\sigma(i+1)})$, we answer it with a "fake" answer $\mathsf{Enc}(k_{\sigma(i)}, r)$ for a random r. The goal is to move from a hybrid with no pebbles (this corresponds to H_L) to one with a single black pebble on the "sink" edge $(n-1, n)$ (this corresponds to H_R). We can prove that neighboring hybrids are indistinguishable via a reduction from CPA-security as long as the pebbling configurations are only modified via the following legal moves:

1. We can put/remove a pebble on the source edge $(0, 1)$ at any time.
2. We can put/remove a pebble on an edge $(i, i+1)$ if the preceding edge $(i-1, i)$ has a pebble.

This is because adding/removing a pebble $(i, i+1)$ means changing what we encrypt under key $k_{\sigma(i)}$ and therefore we need to make sure that either the edge is a source edge or there is already a pebble on the preceding edge to ensure that the key $k_{\sigma(i)}$ is never being encrypted under some other key.

The simplest "basic pebbling strategy" consists of $2n$ moves where we add pebbles on the path $0 \to 1 \to \ldots \to n$, one by one starting on the left and then remove one by one starting on the right, keeping only the pebble on the sink edge $(n-1, n)$. This is illustrated in Fig. 1(a) for $n = 8$. The strategy uses n pebbles. However, there are other pebbling strategies that allow us to trade off more moves for fewer pebbles. For example there is a "recursive strategy" (recursively pebble the middle vertex, then recursively pebble the right-most vertex, then recursively remove the pebble from the middle vertex) that uses at most $\log n + 1$ pebbles (instead of n), but requires $3^{\log n} + 1$ moves (instead of just $2n$). This is illustrated in Fig. 1(b).

As we described, each pebbling strategy with ℓ moves gives us a sequence of hybrids $\mathsf{H}_\mathsf{L} = \mathsf{H}_0, \ldots, \mathsf{H}_\ell = \mathsf{H}_\mathsf{R}$ that allows us to prove selective security. Furthermore, we can prove relatively easily that neighboring hybrids $\mathsf{H}_j, \mathsf{H}_{j+1}$ are indistinguishable even if the adversary doesn't commit to the entire permutation σ but only to the value $\sigma(i)$ of vertices i where either H_j or H_{j+1} has a pebble on the edge $(i-1, i)$. Using our framework, we therefore get a proof of adaptive security where the security loss is $\ell \cdot n^p$ where p is the maximum number of pebbles used and ℓ is the number of pebbling moves. In particular, if we use the recursive pebbling strategy described above we only suffer a quasipolynomial security loss $3^{\log n} \cdot n^{\log n + 1}$, as compared with $2n \cdot (n+1)!$ for naïve random guessing where the adversary commits to the entire permutation σ.

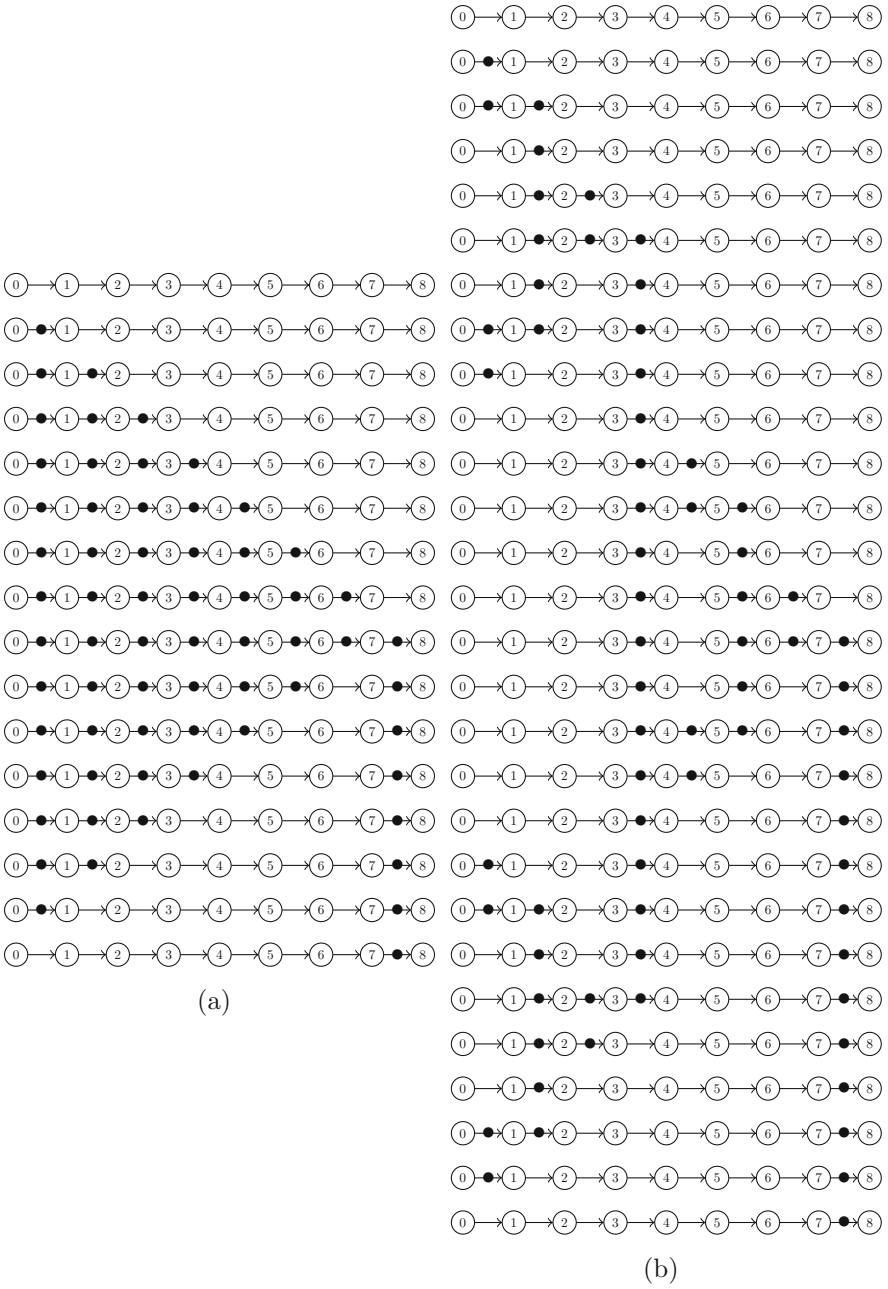

Fig. 1. "Classical" hybrid argument vs. improved hybrid argument. In both diagrams, the edges that carry a pebble are faked. (a) Illustration of the classical hybrids $\mathsf{H}_0, \ldots, \mathsf{H}_{15}$ for GSD on a path graph with $n = 8$ edges: the number of hybrids is $2n = 16$, and the number of fake edges is at most n. (b) A sequence of hybrids $\tilde{\mathsf{H}}_0, \ldots, \tilde{\mathsf{H}}_{27}$ that use fewer fake edges: even though the number of hybrids is $3^{\log n} + 1 = 28$, the number of fake edges is at most $\log n + 1 = 4$. The argument on the right is identical to the one using nested hybrids in [8].

GSD on Low Depth and Other Families of Graphs. The proof outline for GSD on paths is just a very special case of our general result for GSD for various classes of graphs, which we discuss in the full version. If we consider a class of graphs which can be pebbled using ℓ pebbling configurations, each containing at most q pebbles, we get a reduction showing that GSD for this class is $\delta \cdot \ell \cdot 2^q$ secure, assuming the underlying Enc scheme is δ-secure.

Unfortunately, this approach will not gain us much for graphs with high in-degree: we can only put a pebble on an edge (i, j) if all the edges $(*, i)$ going into node i are pebbled. So if we consider graphs which can have large in-degree d, any pebbling strategy must at some point have pebbled all the parents of i, and thus we'll lose at least a factor 2^d in the reduction. But remember that to apply our Theorem 2, we just need to be able to "compress" the information required to simulate the hybrids. So even if the hybrids correspond to configurations with many pebbles, that is fine as long as we can generate a short hint which will allow to emulate it (we use the same idea in the proof of adaptive security of the secret sharing scheme for monotone circuits with large fan-in).

Consider the selective GSD game, where the adversary commits to all of its queries, we can think of this as a DAG, where each edge comes with an index indicating in which query this node was added. Assume the adversary is restricted to choose DAGs of depth l (but no bound on the in-degree). One can show that there exists a pebbling sequence (of length $(2n)^l$), such that in any pebbling configuration, all pebbles lie on a path from a sink to a root (which is of length at most l), and on edges going into this path. Moreover, we can ensure that in any configuration the following holds: if for a node j on this path, there is a pebble on edge (i, j) with index t, then all edges of the form $(*, j)$ with index $<t$ must also have a pebble.

To describe such a configuration, we will output the $\leq l$ nodes on the path, specify for every edge on this path if it is pebbled, and for any node j on the path, the number of edges going into j that have a pebble (note that there are at most $2^l n^{2l}$ choices for this hint). The hint is sufficient to emulate a hybrid, as for any query $(\texttt{encrypt}, i, j)$ the adversary makes, we will know if the corresponding edge has a pebble or not. This is clear if the edge (i, j) is on the path, as we know this path in full. But also for the other edges that can hold a pebble, where j is on the path but i is not. The reason is that we just have to count which query of the form $(*, j)$ this is, as we got a number c telling us that the first c such edges will have a pebble.

Applying Theorem 2, we recover Panjwani's result [23] showing that if the GSD game restricted to graphs of depth l only loses a factor $n^{O(l)}$ in the reduction.

1.3 Yao's Garbled Circuits

Garbled circuits, introduced by Yao in (oral presentations of) [29,30], can be used to garble a circuit C and an input x in a way that reveals $C(x)$ but hides everything else. More precisely, a garbling scheme has three procedures; one to garble the circuit C and produce a garbled circuit $\widetilde{C}$, one to garble the input x and produce a garbled input $\widetilde{x}$, and one that evaluates the garbled circuit $\widetilde{C}$ on

the garbled input $\widetilde{x}$ to get $C(x)$. Furthermore, to prove security, there must be a simulator that only gets the output of the computation $C(x)$ and can simulate the garbled circuit $\widetilde{C}$ and input $\widetilde{x}$, such that no PPT adversary can distinguish them from the real garbling.

Adaptive vs. Selective Security. In the adaptive setting, the adversary A first chooses the circuit C and gets back the garbled circuit $\widetilde{C}$, then chooses the input x, and gets back garbled input $\widetilde{x}$. The adversary's goal is to decide whether he was interacting with the real garbling scheme or the simulator. In the selective setting, the adversary has to choose the circuit C as well as the input x at the very beginning and only then gets back $\widetilde{C}, \widetilde{x}$.

Prior Work. The work of Bellare et al. [2] raised the question of whether Yao's construction or indeed any construction of garbled circuits achieves adaptive security. The work of Hemenway et al. [12] gave the first construction of non-trivial adaptively secure garbled circuits based on one-way functions, by modifying Yao's construction with an added layer of encryption having some special properties. Most recently, the work of Jafargholi and Wichs [16] gives the first analysis of adaptive security for Yao's unmodified garbled circuit construction which significantly improves on the parameters of trivial random guessing. See [16] for a more comprehensive introduction and broader background on garbled circuits and adaptive security.

Here, we present the work of [16] as a special case of our general framework. Indeed, the work of [16] already implicitly follows our general framework fairly closely and therefore we only give a high level overview of how it fits into it.

Selective Hybrids. We start by outlining the selective security proof for Yao's garbled circuits, following the presentation of [12,16] which is in turn based on the proof of Lindell and Pinkas [22]. Essentially the proof proceeds via series of hybrids which modify one garbled gate at a time from the Real distribution to a $\mathsf{Simulated}$ one. However, this cannot be done directly in one step and instead requires going through an intermediate distribution called $\mathsf{InputDep}$ (we explain the name later). There are important restrictions on the order in which these steps can be taken:

1. We can switch a gate from Real to $\mathsf{InputDep}$ (and vice versa) if it is at the input level or if its predecessor gates are already $\mathsf{InputDep}$.
2. We can switch a gate from $\mathsf{InputDep}$ to $\mathsf{Simulated}$ (and vice versa) if it is at the output level or if its successor gates are already $\mathsf{Simulated}$.

The simplest strategy to switch all gates from Real to $\mathsf{Simulated}$ is to start with the input level and go up one level at a time switching all gates to $\mathsf{InputDep}$. Then start with the output level and go down one level at a time switching all gates to $\mathsf{Simulated}$. This corresponds to the basic proof of selective security of Yao garbled circuits.

However, the above is not the only possibility. In particular, any strategy for switching all gates from Real to $\mathsf{Simulated}$ following rules (1) and (2) corresponds

to a sequence of hybrid games for proving selective security. We can identify the above with a *pebbling game* where one can place pebbles on the gates of the circuit. The Real distribution corresponds to not having a pebble and there are two types of pebbles corresponding to the InputDep and Simulated distributions. The goal is to start with no pebbles and finish by placing a Simulated pebble on every gate in the circuit while only performing legal moves according to rules (1) and (2) above. Every pebbling strategy gives rise to a sequence of hybrid games $\mathsf{H}_0, \mathsf{H}_1, \ldots, \mathsf{H}_\ell$ for proving selective security, where the number of hybrids ℓ corresponds to the number of moves and each hybrid H_i is defined by the configuration of pebbles after i moves.

From Selective to Adaptive. The problem with translating selective security proofs into the adaptive setting lies with the InputDep distribution of a gate. This distribution depends on the input x (hence the name) and, in the adaptive setting, the input x that the adversary will choose is not yet known at the time when the garbled circuit is created. To be more precise, the InputDep distribution of a gate i only depends on the 1-bit value going over the output wire of that gate during the computation $C(x)$. Moreover, if we take any two fixed hybrid games $\mathsf{H}_i, \mathsf{H}_{i+1}$ corresponding to two neighboring pebble configurations (ones which differ by a single move) we can prove indistinguishability even if the adversary does not commit to the entire n-bit input x ahead of time but only commits to the bits going over the output wires of all gates i that are in InputDep mode in either configuration. This means that as long as the pebbling strategy only uses m pebbles of the InputDep type at any point in time, each pair of hybrids $\mathsf{H}_i, \mathsf{H}_{i+1}$ can proved indistinguishable in a partially selective setting where the adversary only commits to m bits of information about his input ahead of time, rather than committing to the entire n bit input x. Using our framework, this shows that whenever there is a pebbling strategy for the circuit C that requires ℓ moves and uses at most m pebbles of the InputDep type, we can translate the selective hybrids into a proof of adaptive security where the security loss is $\ell \cdot 2^m$.

It turns out that for any graph of depth d there is a pebbling strategy that uses $O(d)$ pebbles and $\ell = 2^{O(d)}$ moves, meaning that we can prove adaptive security with a $2^{O(d)}$ security loss. This leads to a proof of adaptive security for NC^1 circuits where the reduction has only polynomial security loss, but more generally we can often get a much smaller security loss than the trivial 2^n bound achieved by naïve random guessing.[8]

[8] The presentation in [16] follows the above outline fairly closely and the reader can easily match it with our general framework. The one conceptual difference is that we think of all the hybrids H_i as existing in the selective setting where the adversary commits to the entire input but then we analyze indistinguishability of neighboring hybrids in a partially selective setting. The work of [16] thought of the hybrids H_i as already being partially selective, which made it difficult to compare neighboring hybrids, since the adversary was expected to commit to different information in each one. We view our new framework as being conceptually simpler.

1.4 Constrained Pseudorandom Functions

Goldreich et al. [11] introduced the notion of a pseudorandom function (PRF). A PRF is an efficiently computable keyed function $\mathsf{F}\colon \mathcal{K} \times \mathcal{X} \to \mathcal{Y}$, where $\mathsf{F}(k, \cdot)$, instantiated with a random key $k \leftarrow \mathcal{K}$, cannot be distinguished from a function randomly chosen from the set of all functions $\mathcal{X} \to \mathcal{Y}$ with non-negligible probability. More recently, the notion of constrained pseudorandom functions (CPRF) was introduced as an extension of PRFs, by Boneh and Waters [5], Boyle et al. [6] and Kiayias et al. [19], independently. Informally, a constrained PRF allows the holder of a master key to derive keys which are constrained to a set, in the sense that such a key can be used to evaluate the PRF on that set, while the outputs on inputs outside of this set remain indistinguishable from random.

Goldreich et al., in addition to formally defining PRFs, gave a construction of a PRF from any length doubling pseudorandom generator (PRG). Their construction is depicted in Fig. 2. All three of the aforementioned results [5,6,19] show that this GGM construction already gives a so-called "prefix-constrained" PRF, which is a CPRF where for any $x \in \{0,1\}^*$, one can give out keys which allow to evaluate the PRF on all inputs whose prefix is x. This is a simple but already very interesting class of CPRFs as it can be used to construct a punctured PRF, which in turn is a major tool in constructing various sophisticated primitives based on indistinguishability obfuscation (see, for example, [5,13,26]).

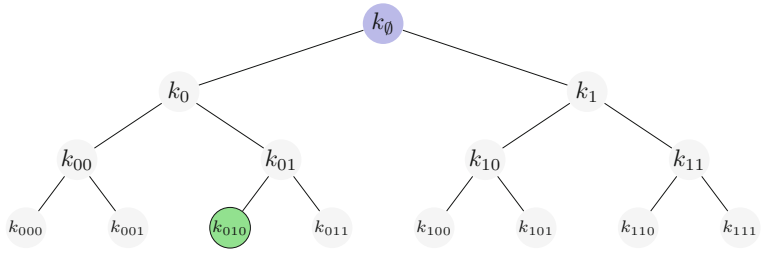

Fig. 2. Illustration of the GGM PRF. Every left child $k_{x\|0}$ of a node k_x is defined as the first half of $\mathsf{PRG}(k_x)$, the right child $k_{x\|1}$ as the second half. The circled node corresponds to $\mathsf{GGM}(k_\emptyset, 010)$.

Prior Work. To show that the GGM construction is a prefix-constrained PRF one must show how to transform an adversary that breaks GGM as a prefix-constrained PRF into a distinguisher for the underlying PRG. The proofs in [5,6,19] only show selective security, where the adversary must initially commit to the output he wants to be challenged on in the security game. There is a loss in tightness by a factor of $2n$. This can then be turned into a proof against adaptive adversaries via random guessing, losing an additional exponential factor 2^n in the input length n.

Fuchsbauer et al. [9] showed that it is possible to achieve adaptive security by losing only factor of $(3q)^{\log n}$, where q denotes the number of queries made by

the adversary—if q is polynomial, the loss is not exponential as before, but just quasi-polynomial. The bound relies on the so-called "nested hybrids" technique. Informally, the idea is to iterate random guessing and hybrid arguments several times. The random guessing is done in a way where one only has to guess some tiny amount of information, which although insufficient to get a full reduction using the hybrid argument, nevertheless reduces the complexity of the task significantly. Every such iteration "cuts" the domain in half, so after logarithmically many iterations the reduction is done. If the number of iterations is small, and the amount of information guessed in each iteration tiny, this can still lead to a reduction with much smaller loss than "single shot" random guessing.

Our Results. We cast the result in [9] in our framework, giving an arguably simpler and more intuitive proof. To this aim, we first describe the GGM construction and sketch its security proof.

Given a $\mathsf{PRG}\colon \{0,1\}^m \to \{0,1\}^{2m}$, the PRF $\mathsf{GGM}\colon \{0,1\}^m \times \{0,1\}^n \to \{0,1\}^m$ is defined recursively as

$$\mathsf{GGM}(k,x) = k_x \text{ where } k_\emptyset = k \text{ and } k_{x\|0} \| k_{x\|1} = \mathsf{PRG}(k_x).$$

The construction is also a prefix-constrained PRF: given a key k_x for any $x \in \{0,1\}^*$, one can evaluate $\mathsf{GGM}(k,x')$ for all x' whose prefix is x.

The security of the GGM as a PRF is given in [11]. In particular, they show that if an adversary exists who distinguishes $\mathsf{GGM}(k,\cdot)$ (real experiment) from a uniformly random function (random experiment) with advantage ϵ making q (adaptive) queries, then an adversary of roughly the same complexity exists who distinguishes $\mathsf{PRG}(U_m)$ from U_{2m} with advantage ϵ/nq. Thus if we assume that PRG is δ-secure, then GGM is δnq-secure against any q-query adversary of the same complexity. This is one of the earliest applications of the hybrid argument.

The security definition for CPRFs is quite different from that of standard PRFs: the adversary will get to query the CPRF $\mathsf{F}(k,\cdot)$ in both, the real and random experiment (and can ask for constrained keys, not just regular outputs), and only at the very end the adversary will choose a challenge query x^*, which is then answered with either the correct CPRF output $\mathsf{F}(k,x^*)$ (in the real experiment) or a random value (in the random experiment). In the selective version of these security experiments, the adversary has to choose the challenge x^* before making any queries. In particular, for the case of prefix-constrained PRFs, the experiment is as follows. The challenger samples $k \in \{0,1\}^n$ uniformly at random. The adversary $\mathcal{A}$ first commits to some $x^* \in \{0,1\}^n$. Then it can make *constrain* queries $x \in \{0,1\}^*$ for any x which is not a prefix of x^*, and receives the constrained key k_x in return. Finally, $\mathcal{A}$ gets either $\mathsf{GGM}(k,x^*)$ (in the real game) or a random value, and must guess which is the case.

Selective Hybrids. A naïve sequence of selective hybrids, which is of length $2n$, relies just on the knowledge of x^*. For $n = 8$ the corresponding 16 hybrid games are illustrated in Fig. 1a. Each path $\mathrm{C}(n)$ corresponds to a hybrid, and it "encodes" how the value of the function F is computed on the challenge input

x^* (and this determines how the function is computed on the rest of the inputs too). An edge that does not carry a pebble is computed, normally, as defined in GGM—i.e., if the ith edge is not pebbled then $k_{x^*[1,i-1]\|0}\|k_{x^*[1,i-1]\|1}$ is set to $\mathsf{PRG}(k_{x[1,i-1]})$, where for $x \in \{0,1\}^n$, $x[1,i]$ denotes its i bit prefix. On the other hand, for an edge with a pebble, we replace the PRG output with a random value—i.e., $k_{x^*[1,i-1]\|0}\|k_{x^*[1,i-1]\|1}$ is set to a uniformly random string in $\{0,1\}^{2m}$. It's not hard to see that any distinguisher for two consecutive hybrids can be directly used to break the PRG with the same advantage by embedding the PRG-challenge – which is either U_{2m} or $\mathsf{PRG}(U_m)$ – at the right place. Using random guessing we can get adaptive security losing an additional factor 2^n in the distinguishing advantage by initially guessing $x^* \in \{0,1\}^n$.

From Selective to Adaptive. Before we explain the improved reduction, we take a step back and consider an even more selective game where A must commit, in addition to the challenge query $x_q = x^*$, also to the constrain queries $\{x_1, \ldots, x_{q-1}\}$. We can use the knowledge of $x_1, \ldots, x_{q-1}$ to get a better sequence of hybrids: this requires two tricks. First, as in GSD on a path, instead of using the pebbling strategy in Fig. 1a, we switch to the recursive pebbling sequence in Fig. 1b. Second, we need a more concise "indexing" for the pebbles: unlike in the proof for GSD, here we can't simply give the positions of the (up to $\log n + 1$) pebbles as hint to simulate the hybrids, as the graph has exponential size, thus even the position of a single pebble would require as many bits to encode as the challenge x^*. Instead, we assume there's an upper bound q on the number of queries made by the adversary. For a pebble on the ith edge, we just give the index of the first constrain query whose i bit prefix coincides with x^*, i.e., the minimum j such that $x_j[1,i] = x^*[1,i]$. This information is sufficient to tell when exactly during the experiment we have to compute a value that corresponds to a pebbled edge.

As there are $3^{\log n}$ hybrids, and each hint comes from a set of size $q^{\log n}$ (i.e., a value $\leq q$ for every pebble), our Theorem 2 implies that GGM is a $\delta(3q)^{\log n}$ secure prefix-constrained PRF if PRG is δ secure. Details are given in the full version [15].

2 Notation

Throughout, we use λ to denote the security parameter. We use capital letters like X to denote variables, small letters like x to denote concrete values, calligraphic letters like $\mathcal{X}$ to denote sets and sans-serif letters like X to denote algorithms. Our algorithms can all be modelled as (potentially interactive, probabilistic, polynomial time) Turing machines. With $\mathsf{X} \equiv \mathsf{Y}$ we denote that X has exactly the same input/output distribution as Y, and $X \sim Y$ denotes that X and Y have the same distributions. $U_{\mathcal{X}}$ denotes the uniform distribution over $\mathcal{X}$. In particular, U_n denotes the uniform distribution over $\{0,1\}^n$. For a set $\mathcal{X}$, $s_{\mathcal{X}}$ denotes the complexity of sampling uniformly at random from $\mathcal{X}$. For $a, b \in \mathbb{N}$, $a \geq b$, by $[a,b]$ we denote the set $\{a, a+1, \ldots, b\}$. For $x \in \{0,1\}^n$ we'll denote with $x[1,i]$ its i bit prefix.

3 The Framework

We consider a game described via a challenger G which interacts with an adversary A. At the end of the interaction, G outputs a decision bit b and we let $\langle \mathsf{A}, \mathsf{G} \rangle$ denote the random variable corresponding to that bit.

Definition 1. *We say that two games defined via challengers G_0 and G_1 are (s, ε)-indistinguishable if for any adversary A of size at most s:*

$$|\Pr[\langle \mathsf{A}, \mathsf{G}_0 \rangle = 1] - \Pr[\langle \mathsf{A}, \mathsf{G}_1 \rangle = 1]| \leq \varepsilon.$$

We say that two games are perfectly indistinguishable and write $\mathsf{G}_0 \equiv \mathsf{G}_1$ if they are $(\infty, 0)$-indistinguishable.

Selectivized Games. We define two operations that convert adaptive or partially selective games into further selective games.

Definition 2 (Selectivized Game). *Given an (adaptive) game G and some function $g\colon \{0,1\}^* \to \mathcal{W}$ we define the selectivized game $\mathsf{H} = \mathsf{SEL}_{\mathcal{W}}[\mathsf{G}, g]$ which works as follows. The adversary A first sends a commitment $w \in \mathcal{W}$ to H. Then H runs the challenger G against A, at the end of which G outputs a bit b'. Let* transcript *denote all communication exchanged between G and A. If $g(\mathsf{transcript}) = w$ then H outputs the bit b' and else it outputs 0. See Fig. 3(a).*

Note that the selectivized game gets a commitment w from the adversary but essentially ignores it during the rest of the game. Only, at the very end of the game, it checks that the commitment matches what actually happened during the game.

Definition 3 (Further Selectivized Game). *Assume $\hat{\mathsf{H}}$ is a (partially selective) game which expects to receive some commitment $u \in \mathcal{U}$ from the adversary in the first round. Given functions $g\colon \{0,1\}^* \to \mathcal{W}$ and $h\colon \mathcal{W} \to \mathcal{U}$ we define the further selectivized game $\mathsf{H} = \mathsf{SEL}_{\mathcal{U} \to \mathcal{W}}[\hat{\mathsf{H}}, g, h]$ as follows. The adversary A first sends a commitment $w \in \mathcal{W}$ to H and H begins running $\hat{\mathsf{H}}$ and passes it $u = h(w)$. It then continues running the game between $\hat{\mathsf{H}}$ and A at the end of which $\hat{\mathsf{H}}$ outputs a bit b'. Let* transcript *denote all communication exchanged between $\hat{\mathsf{H}}$ and A. If $g(\mathsf{transcript}) = w$ then H outputs the bit b' and else it outputs 0. See Fig. 3(b).*

Note that if $\hat{\mathsf{H}}$ is a (partially selective) game where the adversary sends some commitment u, then in the further selectivized game the adversary might have to commit to more information w. The further selectivized game essentially ignores w and only relies on the partial information $u = h(w)$ during the course of the game but only at the very end is still checks that the full commitment w matches what actually happened during the game.

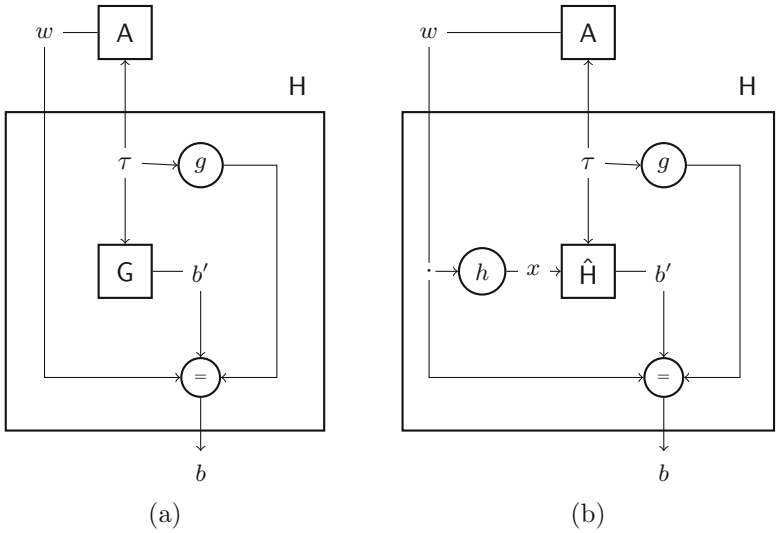

Fig. 3. *Selectivizing.* (a): $\mathsf{SEL}_{\mathcal{W}}[\mathsf{G}, g]$, and (b): $\mathsf{SEL}_{\mathcal{U}\to\mathcal{W}}[\hat{\mathsf{H}}, g, h]$. The symbol τ is short for transcript, the nodes with g and h compute the respective functions, whereas the node with $=$ outputs a bit b as prescribed in the consistency check.

Random Guessing. We first present the basic reduction using random guessing.

Lemma 1. *Assume we have two games defined via challengers* G_0 *and* G_1 *respectively. Let* $g\colon \{0,1\}^* \to \mathcal{W}$ *be an arbitrary function and define the selectivized games* $\mathsf{H}_b = \mathsf{SEL}_{\mathcal{W}}[\mathsf{G}_b, g]$ *for* $b \in \{0,1\}$*. If* H_0, H_1 *are* (s,ε)*-indistinguishable then* G_0, G_1 *are* $(s - s_{\mathcal{W}}, \varepsilon \cdot |\mathcal{W}|)$*-indistinguishable, where* $s_{\mathcal{W}}$ *denotes the complexity of sampling uniformly at random from* $\mathcal{W}$*.*

Proof. We prove the contrapositive. Assume that there is an adversary A of size $s' = s - s_{\mathcal{W}}$ such that

$$|\Pr[\langle \mathsf{A}, \mathsf{G}_0\rangle = 1] - \Pr[\langle \mathsf{A}, \mathsf{G}_1\rangle = 1]| > \varepsilon \cdot |\mathcal{W}|.$$

Let A^* be the adversary that first chooses a uniformly random $w \leftarrow \mathcal{W}$ and then runs A. Then for $b \in \{0,1\}$:

$$\Pr[\langle \mathsf{A}^*, \mathsf{H}_b\rangle = 1] = \Pr[\langle \mathsf{A}, \mathsf{G}_b\rangle = 1]/|\mathcal{W}|$$

and therefore

$$|\Pr[\langle \mathsf{A}^*, \mathsf{H}_0\rangle = 1] - \Pr[\langle \mathsf{A}^*, \mathsf{H}_1\rangle = 1]| > \varepsilon.$$

Moreover, since A^* is of size $s' + s_{\mathcal{W}} = s$ this shows that H_0 and H_1 are not (s,ε)-indistinguishable.

Partially Selective Hybrids. Consider the following setup. We have two adaptive games G_L and G_R. For some function $g\colon \{0,1\}^* \to \mathcal{W}$ we define the selectivized games $\mathsf{H}_\mathsf{L} = \mathsf{SEL}_\mathcal{W}[\mathsf{G}_\mathsf{L}, g]$, $\mathsf{H}_\mathsf{R} = \mathsf{SEL}_\mathcal{W}[\mathsf{G}_\mathsf{R}, g]$ where the adversary commits to some information $w \in \mathcal{W}$. Moreover, to show the indisitinguishability of $\mathsf{H}_\mathsf{L}, \mathsf{H}_\mathsf{R}$ we have a sequence of ℓ (selective) hybrid games $\mathsf{H}_\mathsf{L} = \mathsf{H}_0, \mathsf{H}_1, \ldots, \mathsf{H}_\ell = \mathsf{H}_\mathsf{R}$.

If we only assume that neighboring hybrids $\mathsf{H}_i, \mathsf{H}_{i+1}$ are indistinguishable then by combining the hybrid argument and random guessing we know that G_L and G_R are indistinguishable at a security loss of $\ell \cdot |\mathcal{W}|$.

Theorem 1. *Assume that for each $i \in \{0, \ldots, \ell - 1\}$, the games $\mathsf{H}_i, \mathsf{H}_{i+1}$ are (s, ε)-indistinguishable. Then G_L and G_R are $(s - s_\mathcal{W}, \varepsilon \cdot \ell \cdot |\mathcal{W}|)$-indistinguishable, where $s_\mathcal{W}$ denotes the complexity of sampling uniformly at random from $\mathcal{W}$.*

Proof. Follows from Lemma 1 and the hybrid argument.

Our goal is to avoid the loss of $|\mathcal{W}|$ in the above theorem. To achieve this, we will assume a stronger condition: not only are neighboring hybrids $\mathsf{H}_i, \mathsf{H}_{i+1}$ indistinguishable, but they are selectivized versions of less selective games $\hat{\mathsf{H}}_{i,0}, \hat{\mathsf{H}}_{i,1}$ which are already indistinguishable. In particular, we assume that for each pair of neighboring hybrids $\mathsf{H}_i, \mathsf{H}_{i+1}$ there exist some less selective hybrids $\hat{\mathsf{H}}_{i,0}, \hat{\mathsf{H}}_{i,1}$ where the adversary only commits to much less information $h_i(w) \in \mathcal{U}$ instead of $w \in \mathcal{W}$. In more detail, for each i there is some function $h_i\colon \mathcal{W} \to \mathcal{U}$ that lets us interpret H_{i+b} as a selectivized version of $\hat{\mathsf{H}}_{i,b}$ via $\mathsf{H}_{i+b} \equiv \mathsf{SEL}_{\mathcal{U}\to\mathcal{W}}[\hat{\mathsf{H}}_{i,b}, g, h_i]$. In that case, the next theorem shows that we only get a security loss proportional to $|\mathcal{U}|$ rather than $|\mathcal{W}|$. Note that different pairs of "less selective hybrids" $\hat{\mathsf{H}}_{i,0}, \hat{\mathsf{H}}_{i,1}$ rely on completely different partial information $h_i(w)$ about the adversary's choices. Moreover, the "less selective" hybrid that we associate with each H_i can be different when we compare $\mathsf{H}_{i-1}, \mathsf{H}_i$ (in which case it is $\hat{\mathsf{H}}_{i-1,1}$) and when we compare H_i and H_{i+1} (in which case it is $\hat{\mathsf{H}}_{i,0}$).

Theorem 2 (main). *Let G_L and G_R be two adaptive games. For some function $g\colon \{0,1\}^* \to \mathcal{W}$ we define the selectivized games $\mathsf{H}_\mathsf{L} = \mathsf{SEL}_\mathcal{W}[\mathsf{G}_\mathsf{L}, g]$, $\mathsf{H}_\mathsf{R} = \mathsf{SEL}_\mathcal{W}[\mathsf{G}_\mathsf{R}, g]$. Let $\mathsf{H}_\mathsf{L} = \mathsf{H}_0, \mathsf{H}_1, \ldots, \mathsf{H}_\ell = \mathsf{H}_\mathsf{R}$ be some sequence of hybrid games.*

Assume that for each $i \in \{0, \ldots, \ell - 1\}$, there exists a function $h_i\colon \mathcal{W} \to \mathcal{U}$ and games $\hat{\mathsf{H}}_{i,0}, \hat{\mathsf{H}}_{i,1}$ such that:

$$\mathsf{H}_i \equiv \mathsf{SEL}_{\mathcal{U}\to\mathcal{W}}[\hat{\mathsf{H}}_{i,0}, g, h_i], \quad \mathsf{H}_{i+1} \equiv \mathsf{SEL}_{\mathcal{U}\to\mathcal{W}}[\hat{\mathsf{H}}_{i,1}, g, h_i]. \tag{1}$$

Furthermore, assume that $\hat{\mathsf{H}}_{i,0}, \hat{\mathsf{H}}_{i,1}$ are (s, ε)-indistinguishable. Then G_L and G_R are $(s - s_\mathcal{U}, \varepsilon \cdot \ell \cdot |\mathcal{U}|)$-indistinguishable, where $s_\mathcal{U}$ denotes the complexity of sampling uniformly at random from $\mathcal{U}$.

Proof. Assume that A is an adaptive distinguisher for G_L and G_R of size s' such that

$$|\Pr[\langle \mathsf{A}, \mathsf{G}_\mathsf{L} \rangle = 1] - \Pr[\langle \mathsf{A}, \mathsf{G}_\mathsf{R} \rangle = 1]| > \varepsilon'.$$

Let A^* be a fully selective distinguisher that guesses $w \leftarrow \mathcal{W}$ uniformly at random in the first round and then runs A. By the same argument as in Lemma 1 and Theorem 1 we know that there exists some $i \in [0, \ell)$ such that:

$$|\Pr[\langle \mathsf{A}^*, \mathsf{H}_i \rangle = 1] - \Pr[\langle \mathsf{A}^*, \mathsf{H}_{i+1} \rangle = 1]| \geq \varepsilon'/(\ell \cdot |\mathcal{W}|) \tag{2}$$

Let A' be a partially selective distinguisher that guesses $u \leftarrow \mathcal{U}$ uniformly at random in the first round and then runs A. We want to relate the probabilities $\Pr[\langle \mathsf{A}^*, \mathsf{H}_{i+b} \rangle = 1]$ and $\Pr[\langle \mathsf{A}', \hat{\mathsf{H}}_{i,b} \rangle = 1]$.

Recall that the game $\langle \mathsf{A}^*, \mathsf{H}_{i+b} \rangle$ consists of A^* selecting a uniformly random value $w \leftarrow \mathcal{W}$ (which we denote by the random variable W) and then we run A against $\hat{\mathsf{H}}_{i,b}(u)$ (denoting the challenger $\hat{\mathsf{H}}_{i,b}$ that gets a commitment u in first round) which results in some $\mathsf{transcript}$ and an output bit b^*; if $g(\mathsf{transcript}) = w$ the final output is b^* else 0.

Similarly, the game $\langle \mathsf{A}', \hat{\mathsf{H}}_{i,b} \rangle$ consists of A' selecting a uniformly random value $u \leftarrow \mathcal{U}$ (which we denote by the random variable U) and then we run A against $\hat{\mathsf{H}}_{i,b}(u)$. Therefore:

$$\begin{aligned}
&\Pr[\langle \mathsf{A}^*, \mathsf{H}_{i+b} \rangle = 1] \\
&\quad = \sum_{u \in \mathcal{U}} \underbrace{\Pr[h_i(W) = u]}_{\text{I}} \cdot \underbrace{\Pr[\langle \mathsf{A}, \hat{\mathsf{H}}_{i,b}(u) \rangle = 1]}_{\text{II}} \cdot \Pr[W = g(\mathsf{transcript}) | \text{I}, \text{II}] \\
&\quad = \sum_{u \in \mathcal{U}} \frac{|h_i^{-1}(u)|}{|\mathcal{W}|} \cdot \Pr[\langle \mathsf{A}, \hat{\mathsf{H}}_{i,b}(u) \rangle = 1] \cdot \frac{1}{|h_i^{-1}(u)|} \\
&\quad = \frac{1}{|\mathcal{W}|} \cdot \sum_{u \in \mathcal{U}} \Pr[\langle \mathsf{A}, \hat{\mathsf{H}}_{i,b}(u) \rangle = 1] \\
&\quad = \frac{|\mathcal{U}|}{|\mathcal{W}|} \cdot \sum_{u \in \mathcal{U}} \Pr[\langle \mathsf{A}, \hat{\mathsf{H}}_{i,b}(u) \rangle = 1] \cdot \Pr[U = u] \\
&\quad = \frac{|\mathcal{U}|}{|\mathcal{W}|} \cdot \Pr[\langle \mathsf{A}', \hat{\mathsf{H}}_{i,b} \rangle = 1]
\end{aligned}$$

Combining the above with Eq. 2 we get:

$$|\Pr[\langle \mathsf{A}', \hat{\mathsf{H}}_{i,0} \rangle = 1] - \Pr[\langle \mathsf{A}', \hat{\mathsf{H}}_{i,1} \rangle = 1]| \geq \varepsilon'/(\ell \cdot |\mathcal{U}|)$$

Since by assumption $\hat{\mathsf{H}}_{i,0}, \hat{\mathsf{H}}_{i,1}$ are (s, ε)-indistinguishable and A' is of size $s' + s_{\mathcal{U}}$ this shows that when $s' = s - s_{\mathcal{U}}$ then $\varepsilon' \leq \varepsilon \cdot \ell \cdot |\mathcal{U}|$ which proves the theorem.

3.1 Example: GSD on a Path

As an example, we consider the problem of generalised selective decryption (GSD) on a path graph with n edges, where n is a power of two.

Let $(\mathsf{Enc}, \mathsf{Dec})$ be a symmetric encryption scheme with (probabilistic) $\mathsf{Enc}\colon \mathcal{K} \times \mathcal{M} \to \mathcal{C}$ and $\mathsf{Dec}\colon \mathcal{K} \times \mathcal{C} \to \mathcal{M}$. We assume $\mathcal{K} \subseteq \mathcal{M}$ so that we can

encrypt keys, and that the encryption scheme is (s, δ)-indistinguishable under chosen-plaintext attack.[9] In the game, the challenger—either $\mathsf{G_L}$ or $\mathsf{G_R}$—picks $n+1$ random keys $k_0, \ldots, k_n \in \mathcal{K}$, and the adversary A is then allowed to make two types of queries:

- Encryption queries, $(\texttt{encrypt}, v_i, v_j)$: it receives back $\mathsf{Enc}(k_i, k_j)$.
- Challenge query, $(\texttt{challenge}, v_{i^*})$: here the answer differs between $\mathsf{G_L}$ and $\mathsf{G_R}$, with $\mathsf{G_L}$ answering with k_{i^*} (real key) and $\mathsf{G_R}$ answering with $r \leftarrow \mathcal{K}$ (random, "fake" key).

A cannot ask arbitrary queries: it is restricted to encryption queries that form a path graph with the challenge query as the sink. That is, a valid attacker A is allowed exactly n encryption queries $(\texttt{encrypt}, v_{i_t}, v_{j_t})$, for $t = 1, \ldots, n$, and a single $(\texttt{challenge}, v_{i^*})$ query such that the directed graph $\mathrm{G}_\kappa = (\mathcal{V}, \mathcal{E})$ with $\mathcal{V} = \{v_0, \ldots, v_n\}$ and $\mathcal{E} = \{(v_{i_1}, v_{j_1}), \ldots, (v_{i_n}, v_{j_n})\}$ forms a path with sink v_{i^*}.

Fully Selective Hybrids. Let's look at a naïve sequence of intermediate hybrids $\mathsf{H}_0, \ldots, \mathsf{H}_{2n-1}$. The fully selective challenger H_I receives as commitment the exact permutation σ that A will query—i.e., $v_{\sigma(i)}$ is the ith vertex on the path. Therefore, $\mathcal{W} = S_{n+1}$ (the symmetry group over $0, \ldots, n$) and g is the function that outputs the observed permutation from $\mathsf{transcript}$. Next, H_I samples $2(n+1)$ keys $k_0, \ldots, k_n, r_0, \ldots, r_n$, and when A makes a query $(\texttt{encrypt}, v_{\sigma(i)}, v_{\sigma(i+1)})$, it returns

for $0 \le I \le n$:

$$\begin{array}{ll} \mathsf{Enc}(k_{\sigma(i)}, r_{\sigma(i+1)}) \textbf{ if } (0 \le i \le I) & \text{(Fake edge)} \\ \mathsf{Enc}(k_{\sigma(i)}, k_{\sigma(i+1)}) \textbf{ otherwise}. & \text{(Real edge)} \end{array}$$

for $n < I \le 2n-1$:

$$\begin{array}{ll} \mathsf{Enc}(k_{\sigma(i)}, r_{\sigma(i+1)}) \textbf{ if } (0 \le i \le 2n-1-I) \vee (i = n-1) & \text{(Fake edge)} \\ \mathsf{Enc}(k_{\sigma(i)}, k_{\sigma(i+1)}) \textbf{ otherwise}. & \text{(Real edge)} \end{array} \tag{3}$$

Thus, in the sequence $\mathsf{H}_0, \ldots, \mathsf{H}_{2n-1}$, edges are "faked" sequentially down the path, and then "restored", except for the last edge, in the reverse order up the path—see Fig. 1a. By definition, $\mathsf{H}_0 = \mathsf{G_L}$ and $\mathsf{H}_{2n-1} = \mathsf{G_R}$. Moreover, H_I and H_{I+1} can be shown (s, δ)-indistinguishable: when A queries for $(\texttt{encrypt}, v_{\sigma(I)}, v_{\sigma(I+1)})$, the reduction R_I returns the challenge ciphertext

$$\begin{array}{ll} \mathsf{C}(\cdot, k_{\sigma(I+1)}, r_{\sigma(I+1)}) \textbf{ if } (I \le n) & \text{(Real to fake)} \\ \mathsf{C}(\cdot, r_{\sigma(I+1)}, k_{\sigma(I+1)}) \textbf{ otherwise}. & \text{(Fake to real)} \end{array} \tag{4}$$

For the rest of the queries, R_I works as prescribed in Eq. 3.[10] It is easy to see that R_I simulates H_I when the ciphertext corresponds to the first message, and

[9] To be precise, we only need the encryption scheme to be secure in a weaker model where encryptions of two random messages $m_0, m_1 \in \mathcal{K}$ under a random key $k \in \mathcal{K}$ are (s, δ)-indistinguishable, with the adversary having access to ciphertexts on random messages from $\mathcal{K}$.

[10] Even though R_I does not know the key $k_{\sigma(I)}$, the query $(\texttt{encrypt}, v_{\sigma(I-1)}, v_{\sigma(I)})$ does not cause a problem as its response is $\mathsf{Enc}(k_{\sigma(I)}, r_{\sigma(I-1)})$.

H_{I+1} otherwise. By Theorem 1, $(s - n \cdot s_{\mathsf{Enc}}, \delta(2n+1)(n+1)!)$-indistinguishability of G_{L} and G_{R} follows, where s_{Enc} is the complexity of the Enc algorithm and the $(n+1)!$ factor is the size of the set $\mathcal{W} = S_{n+1}$.

Partially Selective Hybrids. In order to simulate according to the strategy just described, it *suffices* for the hybrid (as well as the reduction) to guess the edges that are faked—however, this number can be at most n (e.g., in the middle hybrids) and, therefore, the simulator guesses the whole path anyway. Intuitively, this is where the overall looseness of the bound stems from. Now, consider the alternative sequence of hybrids $\tilde{\mathsf{H}}_0, \ldots, \tilde{\mathsf{H}}_{27}$ given in Fig. 1b: the edges in this sequence are faked and restored, *one* at a time, in a recursive manner to ensure that at most four edges end up fake per hybrid. In particular, the new hybrid $\tilde{\mathsf{H}}_I$, fakes all the edges that belong to a set $\mathcal{P}_I \subseteq \mathcal{E}$. That is, when A makes a query $(\texttt{encrypt}, v_i, v_j)$—instead of following Eq. 3,—$\tilde{\mathsf{H}}_I$ returns

$$\begin{array}{ll} \mathsf{Enc}(k_i, r_j) \textbf{ if } ((v_i, v_j) \in \mathcal{P}_I) & \text{(Fake edge)} \\ \mathsf{Enc}(k_i, k_j) \textbf{ otherwise}. & \text{(Real edge)} \end{array} \tag{5}$$

This strategy can be extended to arbitrary n, and there exists such a sequence of sets $\mathcal{P}_0, \ldots, \mathcal{P}_{3^{\log n}}$ where the sets are of size at most $\log n + 1$.[11]

$\tilde{\mathsf{H}}^{\mathsf{A}}_{I+b}$
1: Obtain $\sigma \in S_{n+1}$ from A
2: Compute $\mathcal{P} := \mathcal{P}_0, \ldots, \mathcal{P}_\ell$
3: Run $\hat{\mathsf{H}}_{I,b}((\mathcal{P}_I, \mathcal{P}_{I+1}))$
4: **if** $g(\mathsf{transcript}) = \sigma$ **then**
5: **return** $\hat{\mathsf{H}}_{I,b}$'s output
6: **else return** 0
7: **end if**

$\hat{\mathsf{H}}^{\mathsf{A}}_{I,b}((\mathcal{P}_I, \mathcal{P}_{I+1}))$
1: Choose $2n$ keys $r_1, \ldots, r_n, k_1, \ldots, k_n \leftarrow \mathcal{K}$
2: Whenever A queries $(\texttt{encrypt}, v_i, v_j)$:
3: **if** $(v_i, v_j) \in \mathcal{P}_{I+b}$ **then return** $\mathsf{Enc}(k_i, r_j)$
4: **else return** $\mathsf{Enc}(k_i, k_j)$
5: **end if**
6: **return** A's output

Algorithm 1: $\tilde{\mathsf{H}}_{I+b} = \mathsf{SEL}_{\mathcal{U} \to \mathcal{W}}[\hat{\mathsf{H}}_{I,b}, g, h_I]$

Next, we show that the above simulation strategy satisfies the requirements for applying Theorem 2. Firstly, as shown in Algorithm 1, the strategy is partially selective—i.e., $\tilde{\mathsf{H}}_{I+b} = \mathsf{SEL}_{\mathcal{U} \to \mathcal{W}}[\hat{\mathsf{H}}_{I,b}, g, h_I]$, where, for $I \in [0, \ell = 3^{\log n}]$, the function $h_I : S_{n+1} \to \mathcal{E}^{\log n + 1}$ computes $\mathcal{P}_I$. Secondly, as the simulation in $\hat{\mathsf{H}}_{I,0}$ and $\hat{\mathsf{H}}_{I,1}$ differ by exactly one edge—which is real in one and fake in the other—they can be shown to be (s, δ)-indistinguishable. To be precise, if $(v_{i^*}, v_{j^*}) := \mathcal{P}_I \triangle \mathcal{P}_{I+1}$, where $\triangle$ denotes the symmetric difference, when A queries for $(\texttt{encrypt}, v_{i^*}, v_{j^*})$, the reduction $\tilde{\mathsf{R}}_I$ returns

$$\begin{array}{ll} \mathsf{C}(\cdot, k_{j^*}, r_{j^*}) \textbf{ if } (\mathcal{P}_I \subset \mathcal{P}_{I+1}) & \text{(Real to fake)} \\ \mathsf{C}(\cdot, r_{j^*}, k_{j^*}) \textbf{ otherwise}. & \text{(Fake to real)} \end{array} \tag{6}$$

with the rest of the queries answered as in Eq. 5.

[11] In the full version, one can see that the sequence $\mathcal{P}_0, \ldots, \mathcal{P}_{3^{\log n}}$ corresponds to an "edge-pebbling" of the path graph.

Although, the number of hybrids is greater than in the previous sequence, the number of fake edges in any hybrid is at most $\log n + 1$. Thus, the reduction can work with less information than earlier. By Theorem 2, $(s - n \cdot s_{\mathsf{Enc}} - s_{\mathcal{P}}, \delta \cdot 3^{\log n} \cdot n^{2(\log n+1)})$-indistinguishability of G_{L} and G_{R} follows, where $s_{\mathcal{P}}$ is the size of the algorithm that generates the set $\mathcal{P} = \{\mathcal{P}_0, \ldots, \mathcal{P}_\ell\}$, and the $n^{2(\log n+1)}$ factor results from the fact that the compressed set $\mathcal{U} = \mathcal{E}^{\log n+1}$. Thus, the bound is improved considerably from exponential to quasi-polynomial. A more formal treatment is given in the full version [15].

4 Adaptive Secret Sharing for Monotone Circuits

Throughout history there have been many formulations of secret sharing schemes, each providing a different notion of correctness or security. We focus here on the computational setting and adapt the definitions of [21] for our purposes. Rogaway and Bellare [25] survey many different definitions, so we refer there for more information.

A computational secret sharing scheme involves a dealer who has a secret, a set of n parties, and a collection M of "qualified" subsets of parties called the access structure.

Definition 4 (Access structure). *An* access structure M *on parties* $[n]$ *is a monotone set of subsets of* $[n]$*. That is,* $M \subseteq 2^{[n]}$ *and for all* $X \in M$ *and* $X \subseteq X'$ *it holds that* $X' \in M$*.*

We sometimes think of M as a characteristic function $M\colon 2^{[n]} \to \{0,1\}$ that outputs 1 on input X if and only if X is in the access structure. Here, we mostly consider access structures that can be described by a *monotone Boolean circuit*. These are directed acyclic graphs (DAGs) in which leaves are labeled by input variables and every internal node is labeled by an OR or AND operation. We assume that the circuit has fan-in k_{in} and fan-out (at most) k_{out}. The computation is done in the natural way from the leaves to the root which corresponds to the output of the computation. A circuit in which every gate has fan-out $k_{\mathsf{out}} = 1$ is called a *formula*.

A secret sharing scheme for M is a method by which the dealer efficiently distributes shares to the parties such that (1) any subset in M can efficiently reconstruct the secret from its shares, and (2) any subset not in M cannot efficiently reveal any partial information on the secret. We denote by Π_i the share of party i and by Π_X the joint shares of parties $X \subseteq [n]$.

Definition 5 (Secret sharing). *Let* $M\colon 2^{[n]} \to \{0,1\}$ *be an access structure. A secret sharing scheme for* M *consists and secret space* $\mathcal{S}$ *of efficient sharing and reconstruction procedures* S *and* R*, respectively, that satisfy the following requirements:*

1. $\mathsf{S}(1^\lambda, n, S)$ *gets as input the unary representation of a security parameter, the number of parties and a secret* $S \in \mathcal{S}$*, and generates a share for each party.*

2. $\mathsf{R}(1^\lambda, \Pi_X)$ *gets as input the unary representation of a security parameter, the shares of a subset of parties* X*, and outputs a string* S'*.*
3. Completeness: *For a qualified set* $X \in M$ *the reconstruction procedure* R *outputs the shared secret:*

$$\Pr\left[\mathsf{R}(1^\lambda, \Pi_X) = S\right] = 1,$$

where the probability is over the randomness of the sharing procedure $\Pi_1, \ldots, \Pi_n \leftarrow \mathsf{S}(1^\lambda, n, S)$*.*
4. Adaptive security: *For every adversary* A *of size* s *it holds that*

$$|\Pr[\langle \mathsf{A}, \mathsf{G}_0 \rangle = 1] - \Pr[\langle \mathsf{A}, \mathsf{G}_1 \rangle = 1]| \leq \epsilon,$$

where the challenger G_b *is defined as follows:*
   (a) *The adversary* A *specifies a secret* $S \in \mathcal{S}$*.*
      i. *If* $b = 0$*: the challenger generate shares* $\Pi_1, \ldots, \Pi_n \leftarrow \mathsf{S}(1^\lambda, n, S)$*.*
      ii. *If* $b = 1$*: the challenger samples a random* $S' \in \mathcal{S}$ *and generate shares* $\Pi_1, \ldots, \Pi_n \leftarrow \mathsf{S}(1^\lambda, n, S')$*.*
   (b) *The adversary* adaptively *specifies an index* $i \in [n]$ *and if the set of parties he requested so far is unqualified, he gets back* Π_i*, the share of the* i*-th party.*
   (c) *Finally, the adversary outputs a bit* b'*, which is the output of the experiment.*

The selective security variant is obtained by changing item 4b in the definition of the challenger G_b so that the adversary first sends a commitment to the set of shares X he wants to see ahead of time before seeing any share. We denote this challenger by $\mathsf{H}_b = \mathsf{SEL}_{2^{[n]}}[\mathsf{G}_b, X]$.

4.1 The Scheme of Yao

Here we describe the scheme of Yao (mentioned in [1], see also Vinod et al. [28]). The access structure M is given by a monotone Boolean circuit that is composed of AND and OR gates with fan-in k_{in} and fan-out (at most) k_{out}. Each leaf in the circuit is associated with an input variable $x_1, \ldots, x_n$ (there may be multiple inputs corresponding to the same input variable). During the sharing process, each wire in the circuit is assigned a label and the shares of party $i \in [n]$ corresponds to the labels of the wires corresponding to the input variable x_i. The sharing is done from the output wire to the leaves. The reconstruction is done in reverse: using the shares of the parties (that correspond to labels of the input wires), we recover the label of the output wire which will correspond to the secret.

The scheme (S, R) uses a symmetric-key encryption scheme $\mathsf{SKE} = (\mathsf{Enc}, \mathsf{Dec})$ in which keys are uniformly random strings in $\{0,1\}^\lambda$ and is ϵ-secure: any polynomial-time adversary cannot distinguish the encryption of $m_1 \in \{0,1\}^\lambda$ from an encryption of $m_2 \in \{0,1\}^\lambda$ with probability larger than ϵ. The sharing procedure S is described in Fig. 4.

The sharing procedure S:

Input: A secret $S \in \{0,1\}^\lambda$.

1. Initialize $\Pi(S,i) = \emptyset$ for every $i \in [n]$.
2. Label the output wire ow with the secret $\ell_{ow} = S$.
3. Repeat the following until all input wires of the circuit are labeled.
   (a) Let g be a gate with k_{in} input wires and (at most) k_{out} output wires. Let $w'_1, \ldots, w'_{k_{\mathsf{out}}}$ be the output wires of g having labels $\ell_{w'_1}, \ldots, \ell_{w'_{k_{\mathsf{out}}}}$ and $w_1, \ldots, w_{k_{\mathsf{in}}}$ be the input wires. Associate with g a fresh encryption key $s_g \leftarrow \{0,1\}^\lambda$.
   (b) If $g = \text{AND}$, assign the label of $w_1, \ldots, w_{k_{\mathsf{in}}}$ to be random conditioned on $\ell_{w_1} \oplus \cdots \oplus \ell_{w_{k_{\mathsf{in}}}} = s_g$.
   (c) If $g = \text{OR}$, assign the label of $w_1, \ldots, w_{k_{\mathsf{in}}}$ to be s_g.
   (d) For every $i \in [n]$, add to the share of the ith party an encryption of the labels of the w'_i's under s_g. That is,

$$\Pi(S,i) = \Pi(S,i) \cup \{(g, \mathsf{Enc}_{s_g}(\ell_{w'_1}), \ldots, \mathsf{Enc}_{s_g}(\ell_{w'_{k_{\mathsf{out}}}}))\}.$$

4. For every input wire w associated with the input variable x_i, add to the share of the ith party the tuple that consists of the name of the wire and its label:

$$\Pi(S,i) = \Pi(S,i) \cup \{(w, \ell_w)\}.$$

5. Output $\Pi(S,1), \ldots, \Pi(S,n)$.

Fig. 4. Yao's secret sharing scheme (S, R) for an access structure M described by a monotone Boolean circuit.

The reconstruction procedure R of the scheme is essentially applying the reverse operations from the leaves of the circuit to the root. Given the labels of the input wires of an AND gate g, we recover the key associated with g by applying a XOR operation on the labels of the input wires, and then recover the labels of the output wires by decrypting the corresponding ciphertexts. Given the labels of the input wires of an OR gate g, we recover the key associated with g by setting it to be the label of any input wire, and then recover the labels of the output wires by decrypting the corresponding ciphertexts. The label of the output wire of the root gate is the recovered secret.

The scheme is efficient in the sense that the share size of each party is bounded by $k_{\mathsf{out}} \cdot \lambda \cdot s$, where s is the number of gates in the circuit. So, if the circuit is of polynomial-size (in n), then the share size is also polynomial (in n and in the security parameter).

Correctness of the scheme follows by an induction on the depth of the circuit and we omit further details here. Vinod et al. [28] proved that this scheme[12] is selectively secure by a sequence of roughly s hybrid arguments, where s is the

[12] To be more precise, the scheme that Vinod et al. presented and analyzed is slightly different. Specifically, they considered AND and OR gates with fan-out 1 and showed how to separately handle FAN-OUT gates (gates that have fan-in 1 and fan-out 2). Their analysis can be modified to handle our scheme.

number of gates in the circuit representation of M. By the basic random guessing lemma (Lemma 1), this scheme is also adaptively secure but the security loss is exponential in the number of players the adversary requests to see. The later can be exponential in $O(n)$ so for our scheme to be adaptively secure, we need the encryption scheme to be exponentially secure.

Theorem 3 [28]**.** *Assume that* SKE *is a* ϵ*-secure symmetric-key encryption scheme. Then, for any polynomial-time adversary* A *and any access structure on* n *parties described by a monotone circuit with* s *gates, it holds that*

$$|\Pr[\langle \mathsf{A}, \mathsf{H}_0 \rangle = 1] - \Pr[\langle \mathsf{A}, \mathsf{H}_1 \rangle = 1]| \leq k_{\mathsf{out}} \cdot s \cdot \epsilon,$$

and (using Lemma 1),

$$|\Pr[\langle \mathsf{A}, \mathsf{G}_0 \rangle = 1] - \Pr[\langle \mathsf{A}, \mathsf{G}_1 \rangle = 1]| \leq 2^n \cdot k_{\mathsf{out}} \cdot s \cdot \epsilon,$$

In the following subsection we prove that the scheme is adaptively secure and the security loss is roughly $2^{d \cdot \log s}$, where d and s are the depth and number of gates, respectively, in the circuit representing the access structure.

Theorem 4. *Assume that* SKE *is* ϵ*-secure. Then, for any polynomial-time adversary* A *and any access structure on* n *parties described by a monotone circuit of depth* d *and* s *gates with fan-in* k_{in} *and fan-out* k_{out}*, it holds that*

$$|\Pr[\langle \mathsf{A}, \mathsf{G}_0 \rangle = 1] - \Pr[\langle \mathsf{A}, \mathsf{G}_1 \rangle = 1]| \leq 2^{d \cdot (\log s + \log k_{\mathsf{in}} + 2)} \cdot (2k_{\mathsf{in}})^{2d} \cdot k_{\mathsf{out}} \cdot \epsilon.$$

4.2 Hybrids and Pebbling Configurations

To prove Theorem 4 we rely on the framework introduced in Theorem 2 that we briefly recall here. Our goal is to prove that an adversary cannot distinguish the challengers $\mathsf{G}_{\mathsf{L}} = \mathsf{G}_0$ and $\mathsf{G}_{\mathsf{R}} = \mathsf{G}_1$, corresponding to the adaptive game. We define the selective version of the games $\mathsf{H}_{\mathsf{L}} = \mathsf{SEL}_{2^{[n]}}[\mathsf{G}_{\mathsf{L}}, X]$ and $\mathsf{H}_{\mathsf{R}} = \mathsf{SEL}_{2^{[n]}}[\mathsf{G}_{\mathsf{R}}, X]$, where the adversary has to commit to the whole set of shares it wished to see ahead of time. We construct a sequence of ℓ selective hybrid games $\mathsf{H}_{\mathsf{L}} = \mathsf{H}_0, \mathsf{H}_1, \ldots, \mathsf{H}_{\ell-1}, \mathsf{H}_\ell = \mathsf{H}_{\mathsf{R}}$. For each H_i we define two selective games $\hat{\mathsf{H}}_{i,0}$ and $\hat{\mathsf{H}}_{i,1}$ and show that for every $i \in \{0, \ldots, \ell-1\}$, there exists a mapping h_i such that the games H_{i+b} and $\hat{\mathsf{H}}_{i,b}$ (for $b \in \{0,1\}$) are equivalent up to the encoding of the inputs to the games (given by h_i). Then, we can apply Theorem 2 and obtain our result.

The Fully-Selective Hybrids. The sequence of fully selective hybrids $\mathsf{H}_{\mathsf{L}} = \mathsf{H}_0$, $\mathsf{H}_1, \ldots, \mathsf{H}_{\ell-1}, \mathsf{H}_\ell = \mathsf{H}_{\mathsf{R}}$ is defined such that each experiment is associated with a pebbling configuration. In a pebbling configuration, each gate is either pebbled or unpebbled. A configuration is specified by a compressed string that fully specifies the names of the gates which have a pebble on them (and the rest of the gates implicitly do not). We will define the possible pebbling configurations later but for now let us denote by Q the number of possible pebbling configurations.

We define for every $j \in [Q]$, a hybrid experiment H_j in which the adversary first commits to the set X of parties it wishes to see their shares, and then the challenger executes a *new sharing procedure* S^j that depends on the j-th pebbling configuration. Roughly, this sharing procedure acts exactly as the original sharing procedure S, but whenever it encounters a gate with a pebble it generates bogus ciphertexts rather than the real ones. This sharing procedure is described in Fig. 5.

The sharing procedure S^j:

Input: A secret $S \in \{0,1\}^\lambda$.

1. Initialize $\Pi_i = \emptyset$ for every $i \in [n]$.
2. Label the output wire ow with the secret $\ell_{ow} = S$.
3. Repeat the following until all input wires of the circuit are labeled.
   (a) Let g be a gate with k_{in} input wires and (at most) k_{out} output wires. Let $w'_1, \ldots, w'_{k_{\mathsf{out}}}$ be the output wires of g having labels $\ell_{w'_1}, \ldots, \ell_{w'_{k_{\mathsf{out}}}}$ and $w_1, \ldots, w_{k_{\mathsf{in}}}$ be the input wires. Associate with g a fresh encryption key $s_g \leftarrow \{0,1\}^\lambda$.
   (b) If $g = \text{AND}$, assign the label of $w_1, \ldots, w_{k_{\mathsf{in}}}$ to be random conditioned on $\ell_{w_1} \oplus \cdots \oplus \ell_{w_{k_{\mathsf{in}}}} = s_g$.
   (c) If $g = \text{OR}$, assign the label of $w_1, \ldots, w_{k_{\mathsf{in}}}$ to be s_g.
   (d) If g has no pebble on it: For every $i \in [n]$, add to the share of the ith party an encryption of the labels of the w'_i's under s_g. That is,
   $$\Pi_i = \Pi_i \cup \left(g, \mathsf{Enc}_{s_g}(\ell_{w'_1}), \ldots, \mathsf{Enc}_{s_g}(\ell_{w'_{k_{\mathsf{out}}}})\right).$$
   (e) If g has a pebble on it: Sample fresh random strings $r_1, \ldots, r_{k_{\mathsf{out}}}$ and for every $i \in [n]$, add to the share of the ith party an encryption of r_i and under s_g. That is,
   $$\Pi_i = \Pi_i \cup \{(g, \mathsf{Enc}_{s_g}(r_1), \ldots, \mathsf{Enc}_{s_g}(r_{k_{\mathsf{out}}}))\}.$$
4. For every input wire w associated with the input variable x_i, add to the share of the ith party the tuple that consists of the name of the wire and its label:
   $$\Pi_i = \Pi_i \cup \{(w, \ell_w)\}.$$
5. Output $\Pi_1, \ldots, \Pi_n$.

Fig. 5. The sharing procedure S^j for an access structure M, described by a monotone Boolean circuit, and the j-th pebbling configuration which encodes the color of the pebble on each gates.

Observe that the hybrid that corresponds to the configuration in which all gates are unpebbled is identical to the experiment H_L and the configuration in which there is a pebble only on the root gate corresponds to the experiment H_R.

Pebbling Rules and Strategies. The rules of the pebbling game depend on the subset of parties whose shares the adversary sees. The rules are:

1. Can place or remove a pebble on any AND gate for which (at least) one input wire is either *not* in X or comes out of a gate with a pebble on it.
2. Can place or remove a pebble on any OR gate for which all of the incoming wires are either input wires *not* in X or come out of gates all of which have pebbles on them.

Our goal is to find a sequence of pebbling rules so that starting with the initial configuration (in which there are no pebbles at all) will end up with a pebbling configuration in which only the root has a pebble. Jumping ahead, we would like for the sequence of pebbling rules to have the property that each configuration is as short to describe as possible (i.e., minimize Q). One way to achieve this is to have at any configuration along the way, as few pebbles as possible. An even more succinct representation can be obtained if we allow many pebbles but have a way to succinctly represent their location. This is what we achieve in the following lemma.

Lemma 2. *For every subset of parties X and any monotone circuit of depth d, fan-in k_{in}, and s gates, there exists a sequence of $(2k_{\mathsf{in}})^{2d}$ pebbling rules such that every pebbling configuration can be uniquely described by at most $d \cdot (\log s + \log k_{\mathsf{in}} + 1)$ bits.*

Proof. A pebbling configuration is described by a list of pairs (gate name, counter), where the counter is a number between 1 and k_{in}, and another bit b to specify whether the root gate has a pebble or not. The counter will represent the number of predecessors, ordered from left to right, that have a pebble on them. Any encoding uniquely defines a pebbling configuration (but notice that the converse is not true).

Denote by $T_X(d)$ the number of pebbling rules needed (i.e., the length of the sequence) and by $P_X(d)$ the maximum size of the description of the pebbling configuration during the sequence. The sequence of pebbling rules is defined via a recursive procedure in the depth d. We first pebble each of the k_{in} predecessors of the root *from left to right* and add a pair (root gate, counter) to the configuration. After we finish pebbling each predecessor we increase the counter by 1 to keep track of how many predecessors have been pebbled. To pebble all predecessors we used $k_{\mathsf{in}} \cdot T_X(d-1)$ pebbling rules and the maximal size of a configuration is at most $P_X(d-1) + (\log s + \log k_{\mathsf{in}} + 1)$. The $\log s$ term comes from specifying the name of the root gate, the $\log k_{\mathsf{in}}$ term come from the number of predecessors of the root gate that have a pebble on them, and the single bit is to signal whether the root gate is pebbled or not.

After this recursive pebbling each of the predecessors have a pebble only at their root gate and the root (of the depth d circuit) has no pebble. Now, we need to remove the pebble from the root of every predecessor of the root gate and put a pebble on the root gate. For the latter we apply one pebbling rule and put a pebble on the root gate. To remove the pebbles from the predecessors of the root

gate we reverse the recursive pebbling procedure (by "unpebbling" from right to left and updating the counter appropriately), resulting in the application of additional $k_{\mathsf{in}} \cdot T_X(d-1)$ pebbling rules. When we finish unpebbling, since the root has no predecessors with pebbles, we remove from the description of the configuration the pair corresponding to the root gate. Thus, we get that the maximum size of a pebbling configuration at any point in time is

$$P_X(d) \leq P_X(d-1) + (\log s + \log k_{\mathsf{in}} + 1) \Rightarrow P_X(d) \leq d \cdot (\log s + \log k_{\mathsf{in}} + 1).$$

The total number of pebbling rules we apply is

$$T_X(d) \leq 2k_{\mathsf{in}} \cdot T_X(d-1) + 1 \Rightarrow T_X(d) \leq (2k_{\mathsf{in}})^{2d}.$$

This completes the proof of the lemma.

Recall that we denote by Q the number of possible pebbling configurations. Using the pebbling strategy from Lemma 2, we get that

$$Q \leq 2^{d \cdot (\log s + \log k_{\mathsf{in}} + 1)}.$$

The Partially-Selective Hybrids. We define the partially selective hybrids $\hat{\mathsf{H}}_{j,0}$ and $\hat{\mathsf{H}}_{j,1}$ for every H_j and $j \in [Q]$. In both hybrid games $\hat{\mathsf{H}}_{j,0}$ and $\hat{\mathsf{H}}_{j,1}$, the adversary first commits to the *j-th pebbling configuration* and *the next pebbling rule to apply*. Denote by $j' \in [Q]$ the index of the pebbling configuration resulting from applying the next configuration rule to the j-th pebbling configuration. In $\hat{\mathsf{H}}_{j,0}$ the challenger samples the shares from S^j and in $\hat{\mathsf{H}}_{j,1}$ the challenger samples the shares from $\mathsf{S}^{j'}$ (but other than this the games do not change).

Denote by $\mathcal{U}$ the space of messages that the adversary has to commit in the partially selective hybrids $\hat{\mathsf{H}}_{j,b}$. This space includes all tuples of pebbling configurations and an additional valid pebbling rule. First, recall that there are Q possible pebbling configurations. Seocnd, observe that a pebbling rule can be described by a gate name: a pebbling rule is just flipping the color of the pebble on that gate. For a circuit with s gates this requires additional $\log s$ bits. Thus, $\mathcal{U} = \{(i, g) \mid i \in [Q], g \in [s]\}$ and this means that the size of $\mathcal{U}$ is bounded by

$$|\mathcal{U}| \leq Q \cdot s \leq 2^{d \cdot (\log s + \log k_{\mathsf{in}} + 1)} \cdot s.$$

By semantic security of the symmetric-key encryption scheme and the fact that we replace k_{out} ciphertexts with bogus ones, we have that the games $\hat{\mathsf{H}}_{j,0}$ and $\hat{\mathsf{H}}_{j,1}$ are indistinguishable. The proof is by planting the challenge ciphertext as the ciphertext in the gate where the "next pebbling rule" is applied. In $\hat{\mathsf{H}}_{j,0}$ it is a "real" ciphertext while in $\hat{\mathsf{H}}_{j,1}$ it is a bogus one.

Lemma 3. *Assume that* SKE *is ϵ-secure. Then, for any polynomial-time adversary* A *and any access structure on n parties described by a monotone circuit it holds that*

$$|\Pr[\langle \mathsf{A}, \hat{\mathsf{H}}_{j,0} \rangle = 1] - \Pr[\langle \mathsf{A}, \hat{\mathsf{H}}_{j,1} \rangle = 1]| \leq k_{\mathsf{out}} \cdot \epsilon.$$

Applying Theorem 2 with the fact that $\ell \leq (2k_{\mathsf{in}})^{2d}$ and $|\mathcal{U}| \leq 2^{d\cdot(\log s+\log k_{\mathsf{in}}+1)} \cdot s$, we get that if SKE is ϵ-secure, then for any polynomial-time adversary A and any access structure on n parties described by a monotone circuit of depth d and s gates of fan-in k_{in} and fan-out k_{out}, it holds that

$$\begin{aligned} |\Pr[\langle \mathsf{A}, \mathsf{G}_0\rangle = 1] - \Pr[\langle \mathsf{A}, \mathsf{G}_1\rangle = 1]| &\leq 2^{d\cdot(\log s+\log k_{\mathsf{in}}+1)} \cdot s \cdot (2k_{\mathsf{in}})^{2d} \cdot k_{\mathsf{out}} \cdot \epsilon \\ &\leq 2^{d\cdot(\log s+\log k_{\mathsf{in}}+2)} \cdot (2k_{\mathsf{in}})^{2d} \cdot k_{\mathsf{out}} \cdot \epsilon. \end{aligned}$$

5 Open Problems

In this work we presented a framework for proving adaptive security of various schemes including secret sharing over access structures defined via monotone circuits, generalized selective decryption, constrained PRFs, and Yao's garbled circuits. The most natural future direction is to find more applications where our framework can be used to prove adaptive security with better security loss than using the standard random guessing. Also, improving our results in terms of security loss is an open problem.

In all of our applications of the framework, the security loss of a scheme is captured by the existence of some pebbling strategy. Does there exist a connection in the opposite direction between the security loss of a scheme and possible pebbling strategies? That is, is it possible to use lower bounds for pebbling strategies to show that various security losses are necessary?

Acknowledgments. The fourth author thanks his advisor Moni Naor for asking whether Yao's secret sharing scheme is adaptively secure and for his support.

References

1. Beimel, A.: Secret-sharing schemes: a survey. In: Chee, Y.M., Guo, Z., Ling, S., Shao, F., Tang, Y., Wang, H., Xing, C. (eds.) IWCC 2011. LNCS, vol. 6639, pp. 11–46. Springer, Heidelberg (2011). doi:10.1007/978-3-642-20901-7_2
2. Bellare, M., Hoang, V.T., Rogaway, P.: Adaptively secure garbling with applications to one-time programs and secure outsourcing. In: Wang, X., Sako, K. (eds.) ASIACRYPT 2012. LNCS, vol. 7658, pp. 134–153. Springer, Heidelberg (2012). doi:10.1007/978-3-642-34961-4_10
3. Benaloh, J.C., Leichter, J.: Generalized secret sharing and monotone functions. In: Goldwasser, S. (ed.) CRYPTO 1988. LNCS, vol. 403, pp. 27–35. Springer, New York (1990). doi:10.1007/0-387-34799-2_3
4. Blakley, G.R.: Safeguarding cryptographic keys. In: Proceedings of AFIPS 1979 National Computer Conference, vol. 48, pp. 313–317 (1979)
5. Boneh, D., Waters, B.: Constrained pseudorandom functions and their applications. In: Sako, K., Sarkar, P. (eds.) ASIACRYPT 2013. LNCS, vol. 8270, pp. 280–300. Springer, Heidelberg (2013). doi:10.1007/978-3-642-42045-0_15
6. 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). doi:10.1007/978-3-642-54631-0_29

7. Canetti, R., Goldreich, O., Goldwasser, S., Micali, S.: Resettable zero-knowledge (extended abstract). In: 32nd ACM STOC, pp. 235–244. ACM Press, May 2000
8. Fuchsbauer, G., Jafargholi, Z., Pietrzak, K.: A quasipolynomial reduction for generalized selective decryption on trees. In: Gennaro, R., Robshaw, M. (eds.) CRYPTO 2015. LNCS, vol. 9215, pp. 601–620. Springer, Heidelberg (2015). doi:10.1007/978-3-662-47989-6_29
9. Fuchsbauer, G., Konstantinov, M., Pietrzak, K., Rao, V.: Adaptive security of constrained PRFs. In: Sarkar, P., Iwata, T. (eds.) ASIACRYPT 2014. LNCS, vol. 8874, pp. 82–101. Springer, Heidelberg (2014). doi:10.1007/978-3-662-45608-8_5
10. Garg, S., Gentry, C., Sahai, A., Waters, B.: Witness encryption and its applications. In: Boneh, D., Roughgarden, T., Feigenbaum, J. (eds.) 45th ACM STOC, pp. 467–476. ACM Press, June 2013
11. Goldreich, O., Goldwasser, S., Micali, S.: On the cryptographic applications of random functions (extended abstract). In: Blakley, G.R., Chaum, D. (eds.) CRYPTO 1984. LNCS, vol. 196, pp. 276–288. Springer, Heidelberg (1985). doi:10.1007/3-540-39568-7_22
12. Hemenway, B., Jafargholi, Z., Ostrovsky, R., Scafuro, A., Wichs, D.: Adaptively secure garbled circuits from one-way functions. In: Robshaw, M., Katz, J. (eds.) CRYPTO 2016, Part III. LNCS, vol. 9816, pp. 149–178. Springer, Heidelberg (2016). doi:10.1007/978-3-662-53015-3_6
13. Hohenberger, S., Sahai, A., Waters, B.: Replacing a random oracle: full domain hash from indistinguishability obfuscation. In: Nguyen, P.Q., Oswald, E. (eds.) EUROCRYPT 2014. LNCS, vol. 8441, pp. 201–220. Springer, Heidelberg (2014). doi:10.1007/978-3-642-55220-5_12
14. Ito, M., Saito, A., Nishizeki, T.: Secret sharing schemes realizing general access structure. In: Proceedings of IEEE Global Telecommunication Conference (Globecom 1987), pp. 99–102 (1987)
15. Jafargholi, Z., Kamath, C., Klein, K., Komargodski, I., Pietrzak, K., Wichs, D.: Be adaptive, avoid overcommitting. Cryptology ePrint Archive, Report 2017/515 (2017). http://eprint.iacr.org/2017/515
16. Jafargholi, Z., Wichs, D.: Adaptive security of Yao's garbled circuits. In: Hirt, M., Smith, A. (eds.) TCC 2016-B. LNCS, vol. 9985, pp. 433–458. Springer, Heidelberg (2016). doi:10.1007/978-3-662-53641-4_17
17. Karchmer, M., Wigderson, A.: Monotone circuits for connectivity require superlogarithmic depth. In: 20th ACM STOC, pp. 539–550. ACM Press, May 1988
18. Karchmer, M., Wigderson, A.: On span programs. In: Proceedings of Structures in Complexity Theory, pp. 102–111 (1993)
19. Kiayias, A., Papadopoulos, S., Triandopoulos, N., Zacharias, T.: Delegatable pseudorandom functions and applications. In: Sadeghi, A.R., Gligor, V.D., Yung, M. (eds.) ACM CCS 2013, pp. 669–684. ACM Press, November 2013
20. Komargodski, I., Moran, T., Naor, M., Pass, R., Rosen, A., Yogev, E.: One-way functions and (im)perfect obfuscation. In: 55th FOCS, pp. 374–383. IEEE Computer Society Press, October 2014
21. Komargodski, I., Naor, M., Yogev, E.: Secret-sharing for NP. J. Cryptol. **30**(2), 444–469 (2017). http://dx.doi.org/10.1007/s00145-015-9226-0
22. Lindell, Y., Pinkas, B.: A proof of security of Yao's protocol for two-party computation. J. Cryptol. **22**(2), 161–188 (2009)
23. Panjwani, S.: Tackling adaptive corruptions in multicast encryption protocols. In: Vadhan, S.P. (ed.) TCC 2007. LNCS, vol. 4392, pp. 21–40. Springer, Heidelberg (2007). doi:10.1007/978-3-540-70936-7_2

24. Robere, R., Pitassi, T., Rossman, B., Cook, S.A.: Exponential lower bounds for monotone span programs. In: 57th FOCS, pp. 406–415. IEEE Computer Society Press (2016)
25. Rogaway, P., Bellare, M.: Robust computational secret sharing and a unified account of classical secret-sharing goals. In: Ning, P., di Vimercati, S.D.C., Syverson, P.F. (eds.) ACM CCS 2007, pp. 172–184. ACM Press, October 2007
26. Sahai, A., Waters, B.: How to use indistinguishability obfuscation: deniable encryption, and more. In: Shmoys, D.B. (ed.) 46th ACM STOC, pp. 475–484. ACM Press, May/Jun 2014
27. Shamir, A.: How to share a secret. Commun. ACM **22**(11), 612–613 (1979)
28. Vinod, V., Narayanan, A., Srinathan, K., Rangan, C.P., Kim, K.: On the power of computational secret sharing. In: Johansson, T., Maitra, S. (eds.) INDOCRYPT 2003. LNCS, vol. 2904, pp. 162–176. Springer, Heidelberg (2003). doi:10.1007/978-3-540-24582-7_12
29. Yao, A.C.C.: Protocols for secure computations (extended abstract). In: 23rd FOCS, pp. 160–164. IEEE Computer Society Press, November 1982
30. Yao, A.C.C.: How to generate and exchange secrets (extended abstract). In: 27th FOCS, pp. 162–167. IEEE Computer Society Press, October 1986