

Monotone Policy BARGs from BARGs and Additively Homomorphic Encryption

Shafik Nassar
UT Austin
shafik@cs.utexas.edu

Brent Waters
UT Austin and NTT Research
bwaters@cs.utexas.edu

David J. Wu
UT Austin
dwu4@cs.utexas.edu

Abstract

A monotone policy batch NP language $\mathcal{L}_{\mathcal{R},P}$ is parameterized by a monotone policy $P: \{0, 1\}^k \rightarrow \{0, 1\}$ and an NP relation \mathcal{R} . A statement (x_1, \dots, x_k) is a YES instance if there exists w_1, \dots, w_k where $P(\mathcal{R}(x_1, w_1), \dots, \mathcal{R}(x_k, w_k)) = 1$. For example, we might say that an instance (x_1, \dots, x_k) is a YES instance if a majority of the statements are true. A monotone policy batch argument (BARG) for NP allows a prover to prove that $(x_1, \dots, x_k) \in \mathcal{L}_{\mathcal{R},P}$ with a proof of size $\text{poly}(\lambda, |\mathcal{R}|, \log k)$, where λ is the security parameter, $|\mathcal{R}|$ is the size of the Boolean circuit that computes \mathcal{R} , and k is the number of instances. Recently, Brakerski, Brodsky, Kalai, Lombardi, and Paneth (CRYPTO 2023) gave the first monotone policy BARG for NP from the learning with errors (LWE) assumption.

In this work, we describe a generic approach for constructing monotone policy BARGs from any BARG for NP together with an additively homomorphic encryption scheme. This yields the first constructions of monotone policy BARGs from the k -Lin assumption in prime-order pairing groups as well as the (subexponential) DDH assumption in *pairing-free* groups. Central to our construction is a notion of a zero-fixing hash function, which is a relaxed version of a predicate-extractable hash function from the work of Brakerski et al. Our relaxation enables a direct realization of zero-fixing hash functions from BARGs for NP and additively homomorphic encryption, whereas the previous notion relied on leveled homomorphic encryption, and by extension, the LWE assumption.

As an application, we also show how to combine a monotone policy BARG with a puncturable signature scheme to obtain a monotone policy aggregate signature scheme. Our work yields the first (statically-secure) monotone policy aggregate signatures that supports general monotone Boolean circuits from standard pairing-based assumptions. Previously, this was only known from LWE.

1 Introduction

A non-interactive batch argument (BARG) for NP allows a prover to convince a verifier that a collection of k statements x_1, \dots, x_k is true with a proof whose size scales *sublinearly* with k . Beyond the immediate application to amortizing the communication cost of NP verification, batch arguments for NP also play a key role in constructing delegation for RAM programs (also known as a succinct non-interactive argument (SNARG) for P) [KVZ21, CJJ21b, KLVW23] and incrementally verifiable computation [DGKV22, PP22]. These objects have received extensive study recently, and to date, we have constructions from most standard algebraic assumptions in cryptography such as the learning with errors (LWE) assumption [CJJ21b], the k -Lin assumption on groups with bilinear maps [WW22], the (sub-exponential) decisional Diffie-Hellman (DDH) assumption in pairing-free groups [CGJ⁺23], or combinations of quadratic residuosity and (sub-exponential) DDH in pairing-free groups [CJJ21a, HJKS22].

Beyond batch NP and P. The recent successes in constructing succinct arguments for batch NP and for P from standard cryptographic assumptions has motivated the study of other (sub)-classes of NP for which we can build succinct non-interactive arguments from standard (falsifiable) assumptions. Very recently, Brakerski, Brodsky, Kalai, Lombardi, and Paneth [BBK⁺23] showed how to construct SNARGs for monotone policy batch NP. At a high level, the monotone policy batch NP language $\mathcal{L}_{\mathcal{R},P}$ is defined with respect to an NP relation \mathcal{R} together with a monotone policy $P: \{0, 1\}^k \rightarrow \{0, 1\}$ as follows:

$$\mathcal{L}_{\mathcal{R},P} = \{(x_1, \dots, x_k) \mid \exists (w_1, \dots, w_k) : P(\mathcal{R}(x_1, w_1), \dots, \mathcal{R}(x_k, w_k)) = 1\}.$$

In words, an instance (x_1, \dots, x_k) is true as long as an acceptable subset of the statements are true (as determined by the policy P). Such “monotone policy batch arguments” capture policies like majority, general thresholds, and more. The standard batch argument corresponds to the special case where the policy P is a simple conjunction.

Brakerski et al. [BBK⁺23] provided two constructions of monotone policy BARGs for NP. The first construction only relies on standard (somewhere extractable) BARGs and collision-resistant hash functions, but could only support monotone policies of logarithmic depth (i.e., monotone NC¹). To extend to monotone policies of arbitrary polynomial depth, they combine standard BARGs with a new notion of a predicate-extractable hash function, which they then build from the LWE assumption (specifically, they rely on leveled homomorphic encryption). This yields a monotone policy batch argument for arbitrary monotone policies from the LWE assumption. Due to the current reliance on leveled homomorphic encryption to construct the predicate-extractable hash function, instantiations of monotone policy BARGs for arbitrary-depth policies rely on the LWE assumption.

1.1 Our Results

Our main result in this work is showing how to construct BARGs for monotone policies by combining a (standard) BARG with an *additively* homomorphic encryption scheme (which can in turn be built from most number-theoretic assumptions [Gam84, Pai99, Reg05]). Combined with the recent progress on constructing BARGs from pairing-based groups [WW22] and pairing-free groups [CGJ⁺23], we obtain the first monotone policy BARGs for NP from the k -Lin assumption over pairing groups and from the (sub-exponentially) DDH assumption in pairing-free groups. We provide an overview of our techniques in Section 1.2 and summarize our main results in the following theorem:

Theorem 1.1 (Informal). *Assuming any of (1) the plain LWE assumption, (2) the k -Lin assumption over pairing groups for any constant $k \in \mathbb{N}$, or (3) the (sub-exponential) DDH assumption in pairing-free groups, there exists a monotone policy BARG for all polynomial-size monotone circuit policies. The monotone policy BARG satisfies non-adaptive soundness and the proof size is $\text{poly}(\lambda + |C| + \log |P|)$, where $|C|$ denotes the size of the Boolean circuit computing the NP relation, and $|P|$ is the size of the monotone policy.*

Monotone policy aggregate signatures. A key difference between Theorem 1.1 and the previous LWE-based construction [BBK⁺23] is that we obtain a *non-adaptively-sound* BARG for monotone circuit policies whereas the [BBK⁺23] construction satisfied a stronger “somewhere extractability” notion. That is, in [BBK⁺23], the common reference string (CRS) can be sampled in a trapdoor mode and the trapdoor can be used to recover a witness for some x_i given a valid proof on statements (x_1, \dots, x_k) . While extractability is often useful to have in a cryptographic primitive, it is not always essential.

As an illustrative example, we show how to use monotone policy BARGs in conjunction with (puncturable) signatures [GVW19] to construct a monotone policy aggregate multi-signature scheme. In an aggregate multi-signature scheme, there is a set of k signers, each with a signing/verification key-pair $(\text{sk}_i, \text{vk}_i)$. Given a policy P and a set of signatures σ_i for $i \in S$ (where σ_i verifies with respect to vk_i) on a common message m , if the set S satisfies the policy P , then it is possible to aggregate $\{\sigma_i\}_{i \in S}$ into a single short signature whose size is sublinear in $|S|$. For instance, P might encode a “threshold” policy that accepts all sets of size at least t . Crucially, static security of our monotone policy aggregate signature scheme only relies on *non-adaptive soundness* of the monotone policy BARG and security of the puncturable signature scheme. There is no need for an explicit extraction requirement. Very briefly, a puncturable signature scheme allows one to sample a “punctured” verification key vk (and associated signing key) for some message m^* . The punctured verification key is computationally indistinguishable from a normal verification key, but has the property that there does not *exist* any signatures on the punctured message m^* with respect to the punctured key. As shown in [GVW19], puncturable (or “all-but-one signatures”) can be constructed from many standard number-theoretic assumptions. We summarize this result in the following theorem:

Theorem 1.2 (Informal). *Assuming the existence of a non-adaptively sound monotone BARG and a puncturable signature scheme, there exists a monotone policy aggregate multi-signature scheme. The scheme satisfies static unforgeability and the size of the aggregate signature is $\text{poly}(\lambda + \log |P|)$, where $|P|$ denotes the size of the circuit computing the monotone policy.*

Theorem 1.2 shows that in combination with puncturable signatures, soundness alone is sufficient for building aggregate signatures for general monotone policies. Notably, **Theorem 1.2** also provides the *first* monotone policy aggregate signature from pairing-based assumptions (in the plain model). Previous work have shown how to build *vanilla* aggregate signatures using (vanilla) non-interactive batch arguments [WW22, DGKV22]. In an independent and concurrent work, [BCJP24] also show how to construct a monotone policy aggregate multi-signature. Their work provides two constructions of monotone policy aggregate (multi)-signatures. The first scheme supports monotone policies that can be implemented by a read-once, bounded-space Turing machine and is also *adaptively* secure. This scheme relies on somewhere extractable BARGs and a verifiable private information retrieval scheme [BKP22], and can be instantiated from standard pairing-based or lattice-based assumptions. The second scheme supports policies implemented by an arbitrary monotone Boolean circuit, but achieves a weaker security definition (closer to static security) and also relies on fully homomorphic encryption (which to date, is not known from pairing-based assumptions). **Theorem 1.2** gives a statically-secure monotone policy aggregate signature scheme that supports all monotone Boolean circuits, and does not rely on fully homomorphic encryption. This enables a new instantiation from pairings.

Soundness vs. extraction. While **Theorem 1.2** shows that extraction is *unnecessary* for all applications of monotone policy BARGs, our proof strategy for arguing soundness can nonetheless be extended to achieve a notion of extractability (see [Section 8](#)). The notion we achieve is similar to the somewhere extractability notion from [BBK⁺23], where for every monotone policy P , they define a notion of a “necessary set” associated with P (i.e., a set with the property that for every satisfying input (x_1, \dots, x_n) to P , there exists $i \in S$ where $x_i = 1$). The somewhere extractability notion from [BBK⁺23] programs S into the common reference string, and asserts that whenever the prover comes up with an accepting proof for statements (x_1, \dots, x_k) for an NP relation \mathcal{R} and policy P , then the extractor will output w_i for $i \in S$ where $\mathcal{R}(x_i, w_i) = 1$. Our construction satisfies a looser variant of this property where the success probability of the extractor is smaller by a factor of $1/k$. We refer to this notion as semi-somewhere extractability. While our construction does achieve this notion of extraction with essentially no modification (see [Section 8](#)), we choose to focus on the simpler notion of non-adaptive soundness in the core part of this paper. Our rationale is twofold:

- First, there is a lack of consensus on what the “right” notion of extraction is when it comes to the setting of monotone policy BARGs. Notably, the recent and concurrent work of [BCJP24] that builds monotone policy aggregate signatures highlighted the *inadequacy* of the somewhere extractability notion from [BBK⁺23] for their particular application to constructing monotone policy aggregate signatures. Indeed, the work of [BCJP24] propose two different and seemingly incomparable notions of extraction for their application. This illustrates that the most useful or desirable notion of extraction for monotone policy BARGs may be application-dependent.
- Second, while it is straightforward to show that our construction satisfies some notion of extractability, proving this property does not appear to confer additional capabilities. For the main application to statically-secure aggregate signatures, we showed above that non-adaptive soundness already suffices. There is no need for extraction if this is the end goal. The main advantage of having some kind of extractability definition is we can apply this construction to compile any digital signature scheme into a monotone policy aggregate signature scheme, as opposed to restricting ourselves to puncturable signatures (and we show this in [Section 8.1](#)). While there is a qualitative benefit to this, we do not view it as strong evidence that semi-somewhere extractability is a clearly more powerful or more useful notion than non-adaptive soundness.

A new application: general-policy BARGs for $\text{NP} \cap \text{coNP}$. We also highlight a simple application of BARGs for monotone policy batch NP to constructing a BARG that supports *arbitrary* policies over languages in $\text{NP} \cap \text{coNP}$. Our observation essentially follows the similar strategy of extending monotone closure of SZK to non-monotone closure [Vad06]. Specifically, for a language $\mathcal{X} \in \text{NP} \cap \text{coNP}$ and an arbitrary policy $P: \{0, 1\}^k \rightarrow \{0, 1\}$, we define the language

$$\mathcal{L}_{\mathcal{X}, P} = \{(x_1, \dots, x_k) \mid P(b_1, \dots, b_k) = 1 \text{ where } b_i = \mathbb{1}\{x_i \in \mathcal{X}\}\},$$

where $\mathbb{1}\{x_i \in \mathcal{X}\}$ is the indicator function that outputs 1 if $x_i \in \mathcal{X}$ and 0 otherwise. Importantly, in this context, we allow P to be any arbitrary (possibly non-monotone) Boolean circuit. It is not difficult to see that a BARG for

monotone policy batch NP immediately implies a BARG for $\mathcal{L}_{\mathcal{X},P}$. Namely, we first re-express the circuit P on k inputs b_1, \dots, b_k as a new *monotone* circuit P' on $2k$ inputs corresponding to the original input bits b_1, \dots, b_k as well as their negations $\bar{b}_1, \dots, \bar{b}_k$. We can then apply a BARG for monotone policy batch NP on the set of $2k$ inputs with the policy P' . For this transformation to work, it is important that for each statement x_i , the prover can either provide a proof of membership $x_i \in X$ (which sets $b_i = 1$) or a proof of non-membership $x_i \notin X$ (which sets $\bar{b}_i = 1$).

1.2 Technical Overview

The starting point of our BARG construction is the “canonical protocol” from [BBK⁺23, §2.1]. We recall this below. In our description, we will consider the NP relation of Boolean circuit satisfiability.

- Given a Boolean circuit $C: \{0, 1\}^n \times \{0, 1\}^h \rightarrow \{0, 1\}$, a monotone policy $P: \{0, 1\}^k \rightarrow \{0, 1\}$, statements $x_1, \dots, x_k \in \{0, 1\}^n$, witnesses $w_1, \dots, w_k \in \{0, 1\}^h$, the prover first computes $b_i \leftarrow C(x_i, w_i)$ for all $i \in [k]$.
- The prover then evaluates the circuit $P(b_1, \dots, b_k)$. The prover commits to all of the wire values in $P(b_1, \dots, b_k)$ using a succinct commitment com that supports local openings. We index the input wires with the integers $1, \dots, k$, the output wire by s (where s is the number of wires in P), and the intermediate wires with $k+1, \dots, s-1$.
- The prover uses a batch argument to prove the following statements with respect to the commitment com :
 - **Input wires:** For every input wire $j \in [k]$, it proves that there exists a local opening of com to a value $b_j \in \{0, 1\}$ at index j , and moreover, $b_j = C(x_j, w_j)$.
 - **Gate computation:** For every gate g in P with input wires j_L, j_R and output wire j , it proves that there exists a local opening of com to wire values $b_{j_L}, b_{j_R}, b_j \in \{0, 1\}$ at indices $j_L, j_R, j \in [s]$, respectively, and moreover, $b_j = g(b_{j_L}, b_{j_R})$.
 - **Output wire:** It proves that there exists a local opening to the value 1 at index s for com .

The proof consists of the commitment com together with the batch argument π .

When the policy circuit P has logarithmic depth, the authors of [BBK⁺23] describe a simple inductive argument to argue the security of this construction by relying on somewhere extractability of the underlying BARG. Somewhere extractability says that the common reference string of the BARG can be programmed at a small number of (hidden) indices i_1, \dots, i_ℓ . Given a valid proof π for (x_1, \dots, x_n) along with a trapdoor, one can extract witnesses for $x_{i_1}, \dots, x_{i_\ell}$. However, when P has super-logarithmic depth, the basic inductive argument no longer suffices (specifically, the security loss of the reduction decays *exponentially* in the depth of P).

Predicate-extractable hash functions for bit-fixing constraints. To construct monotone policy BARGs for policies P of arbitrary depth, the authors of [BBK⁺23] replace the Merkle hash of the wire values with a more sophisticated “predicate-extractable” hash function for bit-fixing constraints.¹

A predicate-extractable hash function for bit-fixing predicates is a hash function where the hash key can be programmed in one of two computationally indistinguishable modes: (1) a normal mode and (2) a bit-fixing mode. In bit-fixing mode, the setup algorithm takes as input a set of indices $S \subseteq [n]$ along with a collection of bits $\{(i, y_i)\}_{i \in S}$, where n is the input length. It outputs a hash key hk and an extraction trapdoor td . The correctness requirement says that if $\text{dig} = \text{Hash}(\text{hk}, \mathbf{x})$ for an input \mathbf{x} where $x_i = y_i$ for all $i \in S$, then $\text{Extract}(\text{td}, \text{dig}) = \text{Matching}$. Alternatively, if dig is a digest for an input \mathbf{x} where $x_i \neq y_i$ for some $i \in S$, then $\text{Extract}(\text{td}, \text{dig})$ should output $(\text{NotMatching}, i^*)$ where $i^* \in S$ is an index where $x_{i^*} \neq y_{i^*}$. Essentially, the extractor is deciding whether dig corresponds to the hash of an input that is consistent with $\{(i, y_i)\}_{i \in S}$. If the hash is declared inconsistent, the extractor outputs one of the inconsistent indices. Finally, the hash function supports succinct local openings to individual bits of an input. The two key security properties are as follows:

- For a hash digest dig where $\text{Extract}(\text{td}, \text{dig}) = \text{Matching}$, then it should be computationally difficult to construct an opening for dig to a value $x_i \neq y_i$ for any $i \in S$.

¹This is conceptually similar to the notion of function-binding hash functions introduced concurrently in [FWW23].

- For a hash digest dig where $\text{Extract}(\text{td}, \text{dig}) = (\text{NotMatching}, i^*)$, then it should be computationally difficult for the adversary to open index i^* to the value y_{i^*} .

In the monotone BARG construction, the prover takes the Boolean circuit C , the policy P , the statements (x_1, \dots, x_k) and the witnesses (w_1, \dots, w_k) , and computes $b_i \leftarrow C(x_i, w_i)$ and $P(b_1, \dots, b_k)$. Let (b_1, \dots, b_s) be the complete set of wire values in $P(b_1, \dots, b_k)$, arranged in topological order. The prover hashes the wire values (b_1, \dots, b_s) using the predicate-extractable hash function. In fact, the prover computes two independent hashes $\text{dig}_1, \text{dig}_2$ of the wire values, and the BARG will check validity of the openings against *both* hashes. To argue non-adaptive soundness, the authors of [BBK⁺23] first define the zero-set J associated with a circuit C , policy P , and statement (x_1, \dots, x_k) :

- For each $i \in [k]$, let $\beta_i^* = 1$ if there exists w_i such that $C(x_i, w_i) = 1$ and let $\beta_i^* = 0$ otherwise.
- Let $\beta_1^*, \dots, \beta_s^* = P(\beta_1^*, \dots, \beta_k^*)$ be the wire values in $P(\beta_1^*, \dots, \beta_k^*)$, where the wires are ordered topologically.
- Let $J = \{i \in [k] : \beta_i^* = 0\}$. For a layer index t , define $J_t \subseteq J$ to just contain the indices of wires in layer t of P .

The proof of non-adaptive soundness now proceeds as follows:²

- Take any circuit C , monotone policy P , and statements x_1, \dots, x_k . The invariant they use roughly says the following: if hk_0, hk_1 are programmed to bind to the all-zeroes string on the zero-sets J_i, J_{i-1} for layers i and $i-1$ of P , and the digest associated with the upper layer is NotMatching , then the digest associated with the lower layer is also NotMatching .
- To establish this invariant, the proof critically relies on BARG security and security of the predicate-extractable hash function. Namely, if the extractor declares an index $j \in J_i$ in the upper layer to be NotMatching and the BARG is set to be extracting on wire j , then that means the adversary must have opened one of the input wires j' (to the gate computing wire j) to a 1 where $j' \in J_{i-1}$ (since the policy P is monotone). Security of the hash function then says that the extractor must declare the digest associated with the lower layer to be NotMatching .
- To complete the proof, they argue that the output layer must be NotMatching (by programming the BARG to be extracting on the output wire). By propagating the invariant to the input wires, they conclude that the input layer must be NotMatching (when one of the hash keys is programmed to bind on the input layer). In this case, programming the BARG to be extracting on the wire identified by the NotMatching input (output by the extractor for the hash function) yields a contradiction (in this case, the BARG extractor would need to output a witness for a false NP statement).

The authors of [BBK⁺23] then show how to construct a predicate-extractable hash function for bit-fixing predicates using the learning with errors (LWE) assumption. Their construction specifically relies on leveled homomorphic encryption (similar to the construction of somewhere statistically binding hash functions [HW15]). In conjunction with BARGs for NP based on LWE [CJJ21b], this yields a monotone policy BARG for NP from LWE.

This work: zero-fixing hash functions. The starting point of our work is a relaxation of a predicate-extractable hash function for bit-fixing predicates we call a zero-fixing hash function. Like the predicate-extractable hash function, the zero-fixing hash function supports succinct local openings and moreover, the hash key for a zero-fixing hash function can be sampled in one of two computationally-indistinguishable modes: (1) a normal mode and (2) a zero-fixing mode. In zero-fixing mode, the setup algorithm takes as input a set $S \subseteq [n]$ of indices (that should be zero) and outputs a hash key hk along with a trapdoor td . There is also an extract algorithm Extract that takes as input the hash key hk and a digest dig , and outputs either Matching or NotMatching . The key distinction with predicate-extractable hash functions is that Extract only outputs the flag; it does *not* output an index when it declares a digest NotMatching . Correspondingly, the zero-fixing security requirement only imposes a requirement for matching digests:

²With a suitable strengthening of the notion of predicate-extractable hash functions, the authors of [BBK⁺23] also show how to obtain a somewhere *extractable* monotone policy BARG. In this work, we focus on achieving the core notion of non-adaptive soundness.

- **Zero-fixing:** Suppose (hk, td) are sampled in zero-fixing mode for a set S . Then, for any digest dig where $\text{Extract}(td, dig)$ outputs `Matching`, it should be hard to find an opening to an index $i \in S$ to the value 1.

While this distinction of having the extractor output a mismatching index j or not might seem like a small difference, it has two significant implications:

- **Simpler to construct:** By only requiring the zero-fixing hash function declare whether a digest is `Matching` or `NotMatching`, we significantly simplify the construction of the hash function. Whereas computing and propagating an index of a “mismatching bit” (as in [BBK⁺23]) relies heavily on (leveled) homomorphic encryption, checking whether there *exists* a mismatching index or not can be realized from simpler tools. As we show in this work (and describe later on), we can construct zero-fixing hash functions generically from BARGs for NP together with any *additively* homomorphic encryption scheme (Section 5). If we prefer to avoid non-black-box techniques altogether, we also describe a direct algebraic construction using composite-order pairing groups (Section 6). This is the critical distinction that allows us to obtain monotone policy BARGs from group-based assumptions (which give additively homomorphic encryption [Gam84] but not leveled homomorphic encryption).
- **Sufficient for monotone policy BARGs:** A second important fact is that our notion of zero-fixing hash function still suffices to build monotone policy BARGs. As noted in the preceding sketch, the soundness analysis from [BBK⁺23] critically relied on the hash function extractor outputting an index of a mismatching bit. This is so that when the BARG is programmed to bind on the wire associated with the mismatching index, the `NotMatching` invariant propagates from the output layer to the input layer. In our setting, the zero-fixing extractor only outputs `Matching` or `NotMatching`, and in the case where the extractor outputs `NotMatching`, we *cannot* definitively declare an index to be “mismatching.” This requires a new proof strategy as well as imposing additional security requirements on the zero-fixing hash function. We describe these properties as well as our new proof strategy in more detail below.

Monotone policy BARGs from zero-fixing hash functions. Our main construction is similar to the canonical protocol from [BBK⁺23] sketched above, except the prover commits to all of the wires of the policy circuit P using two zero-fixing hash functions (with hash keys hk_1 and hk_2). Our security analysis takes a different *bottom-up* approach rather than the previous top-down approach. The bottom-up approach is more natural when using our zero-fixing hash function. Here, we provide a sketch of our non-adaptive soundness analysis.

To argue non-adaptive soundness, fix a Boolean circuit C , a monotone policy P (assumed to be a layered Boolean circuit), and a false statement (x_1, \dots, x_k) . Similar to [BBK⁺23], we define the zero-set J associated with C , P , and (x_1, \dots, x_k) . The zero-set J contains the indices of the wires with value 0 in the computation $P(\beta_1^*, \dots, \beta_k^*)$ where $\beta_i^* = 1$ if there exists w_i where $C(x_i, w_i) = 1$ and 0 otherwise. Since P is monotone, for all w_1, \dots, w_k , the wire values of $P(C(x_1, w_1), \dots, C(x_k, w_k))$ on the set J will be zero. As before, let $J_i \subset J$ be the subset of wires in layer i of P .

Our soundness argument proceeds layer-by-layer, starting from the input layer (i.e., layer 1) and progressing to the output layer (i.e., layer d , where d is the depth of P). Our goal establishes the following invariant: if the hash keys hk_1 and hk_2 are zero-fixing on J_i and J_{i+1} and the digest associated with the lower layer (i.e., layer i) is `Matching`, then the digest associated with the upper layer (i.e., layer $i + 1$) is also `Matching`. We provide a sketch of this step. For ease of exposition, suppose hk_1 is zero-fixing on J_i and the digest dig_1 is `Matching`. The goal is to show that hk_2 is zero-fixing on J_{i+1} , then the digest dig_2 is also `Matching`.³

- Initially, we set hk_2 to be binding on the empty set. We require in this case that dig_2 is always `Matching`. We refer to this property as an extractor validity property on the zero-fixing hash function.
- We now iteratively build up hk_2 . Let $J_{i+1}[1]$ be the first element of J_{i+1} . We set hk_2 to be binding on the set $\{J_{i+1}[1]\}$. Our goal is to argue that dig_2 is still `Matching`. While it might seem like this property should follow

³This step is straightforward if we had a predicate-extractable hash function where the extractor outputs a mismatching index. Namely, if the upper layer digest is `NotMatching`, then the extractor outputs an index $j \in J_{i+1}$ that is mismatching (i.e., cannot be opened to a 0). This means the efficient adversary can only open wire j to the value 1. Now, if the BARG is extracting on the statement associated with wire j , then we either (1) obtain the opening of some index $j' \in J_i$ to a 1, which breaks security of the hash function (since the lower layer digest is `Matching`); or (2) the value of wire j is inconsistent with the input wires associated with the gate computing wire j , which breaks security of the BARG.

assuming a basic index hiding property on the zero-fixing hash function (i.e., that the hash key hk hides which set it is binding on), this is *insufficient*. The reason is that when hk_2 is binding on \emptyset , the adversary might output a Matching digest dig_2 , but if hk_2 is binding on $\{J_{i+1}[1]\}$, the output digest dig_2 might be NotMatching. We cannot use such an adversary to construct an index hiding distinguisher, because in the index hiding security game, the distinguisher does *not* have the extraction trapdoor. As such, an attempted reduction algorithm cannot efficiently decide whether the adversary was successful or not. Indeed, this is a fundamental issue since knowledge of the extraction trapdoor would trivially break index hiding.

- To advance the proof, we introduce a *stronger* notion of index hiding security for zero-fixing hash functions, which essentially requires that no efficient adversary can output a digest dig that causes the output of Extract to differ depending on whether the hash key is binding on a set S or a set $S \setminus \{i\}$.⁴ Of course, this is only meaningful when the digest is computed over an input that is 0 on index i .⁵ Thus, we require this stronger index hiding with extracted guess property to hold only for digests dig where the adversary can provide an opening to index 0 for the target index i . We define this property formally in [Definition 3.1](#).
- To leverage the index hiding with extracted guess property, we need to enforce the fact that dig_2 opens to a 0 on index $J_{i+1}[1]$. We ensure this by appealing to the somewhere extractability of the BARG along with zero-fixing security of the hash function. Specifically, suppose that the BARG is binding on wire $J_{i+1}[1]$. The BARG extractor then produces openings to the wire $J_{i+1}[1]$ with respect to dig_2 as well as opening to the wires j_L, j_R with respect to dig_1 (corresponding to the input wires for the gate computing $J_{i+1}[1]$). Since dig_1 is zero-fixing on J_i and dig_1 is also Matching, if either $j_L, j_R \in J_i$, then the extracted openings must be openings to 0 (otherwise, we break zero-fixing of the hash function). But by monotonicity of P , this means the value of the output wire $J_{i+1}[1]$ must also be 0, and thus the BARG extractor produces an opening to 0 for wire $J_{i+1}[1]$. Now, by the index hiding with extracted guess property, we conclude that programming hk_2 to zero-fix on set $\{J_{i+1}[1]\}$ will still cause dig_2 to be Matching (except with a negligible loss in probability).
- We can now iteratively apply the argument and build up hk_2 until it is binding on all of J_{i+1} .

To complete the proof, we consider the input and output layers for P :

- **Handling the input layers:** The base case in our analysis is to show that if hk_1 is binding on J_1 (the input layer), then it is Matching. This follows using the same layer-wise strategy sketched above for proving our invariance, except for each index $J_1[i]$, we rely on the fact that the associated statement x_i is false (i.e., no witness exists) to argue that the only valid opening for dig_1 on index i is 0. Otherwise, we either break somewhere extractability of the BARG (i.e., extracting an invalid witness for index i) or the index hiding with extracted guess property.
- **Output layer:** Starting from the input layer, we now iteratively apply our basic invariant to argue that when the hash keys are binding to J_d (the output layer), the associated digests are also Matching. Now, if we have a valid proof, and the BARG is set to extract on the output layer, then the BARG extractor outputs an opening of the output wire to 1 with respect to the hash digests. However, since the output wire is contained in J_d (since the statement is false), and the digest is matching, this breaks zero-fixing security of the hash function.

Thus, the above analysis suffices to show non-adaptive soundness of our construction. The critical security requirement we require on our zero-fixing hash function is the strengthened index hiding with extracted guess property. This property allows us to complete the proof via an iterative approach without needing to rely on the extractor outputting a mismatching index as in previous work [[BBK⁺23](#)]. As we discuss below, this is an easier property to realize than full-fledged index extraction. We refer to [Section 3](#) for the formal definition of zero-fixing hash functions and [Section 4](#) for our construction of monotone policy BARGs.

⁴This type of property where the output of the extractor does not change for different choices of the CRS is often referred to as a “no-signaling” extraction property [[PR17](#), [KPY19](#), [GZ21](#), [KVZ21](#), [CJJ21b](#)].

⁵Otherwise, an honest digest on the input that is 1 at index i (and 0 everywhere else) would be declared Matching if the hash key was zero-fixing on a set S that contains i and NotMatching if the hash key was zero-fixing on the set $S \setminus \{i\}$.

Constructing zero-fixing hash functions. Our second contribution in this work is a generic construction of zero-fixing hash functions from vanilla BARGs together with an additively homomorphic encryption scheme. We start with a basic construction that captures the key ideas underlying our construction and refer to [Section 5](#) for the formal description and analysis:

- Let $n \in \mathbb{N}$ be the input length. For ease of exposition, we assume that $n = 2^k$ is a power-of-two. Suppose we want to zero-fix on a (possibly-empty) set $S \subseteq [n]$. The setup algorithm first samples a public/secret key-pair (pk, sk) for an additively homomorphic encryption scheme. For each $i \in [n]$, the setup algorithm constructs an encryption $ct_i \leftarrow \text{Enc}(pk, 1)$ of 1 if $i \in S$ and an encryption of $ct_i \leftarrow \text{Enc}(pk, 0)$ of 0 if $i \notin S$. It also constructs an encryption $ct_{\text{zero}} \leftarrow \text{Enc}(pk, 0)$ of 0. Finally, it constructs a commitment com_{hk} to the ciphertexts (ct_1, \dots, ct_n) . The hash key is then $\text{hk} = (pk, ct_{\text{zero}}, ct_1, \dots, ct_n, \text{com}_{\text{hk}})$, and the extraction trapdoor is the decryption key sk .
- To hash an input $x \in \{0, 1\}^n$, the user constructs a complete binary tree where each of the n leaves is associated with a ciphertext. If $x_i = 1$, then the user associates leaf i with ct_i , and if $x_i = 0$, then the user associates leaf i with ct_{zero} . The value of each internal node in the binary tree is defined to be the sum of the ciphertexts associated with its two children. By construction, the value of the root node is an encryption of the sum of the values associated with the n leaf nodes. We refer to the tree of ciphertexts as the “ciphertext-evaluation tree.” The digest dig then consists of the ciphertext ct_{root} associated with the root node along with a commitment com_{ct} to all of the ciphertexts in the ciphertext-evaluation tree.
- A local opening for index i^* and value $b_{i^*} \in \{0, 1\}$ for the digest $\text{dig} = (ct_{\text{root}}, \text{com}_{\text{ct}})$ is a BARG proof. The BARG statements correspond to the indices of the nodes in the ciphertext-evaluation tree. The associated relation is parameterized by the target index i^* , the root ciphertext ct_{root} , the encryption ct_{zero} of 0 from the hash key, and the commitment to the input ciphertexts com_{hk} . The BARG relation then checks the following:
 - **Leaf nodes:** For each leaf node i , com_{ct} opens to either ct_{zero} or ct_i at index i . For the particular index i^* , it checks that com_{ct} opens to ct_{zero} if $b_{i^*} = 0$ and com_{ct} opens to ct_{i^*} if $b_{i^*} = 1$. Since the BARG relation only has com_{hk} and not ct_i itself, the prover provides ct_i as part of its witness along with a proof of opening for ct_i with respect to com_{hk} . The proof of opening ensures that the correct ct_i is provided.
 - **Internal nodes:** For an internal node i (with children indexed j_L, j_R), the BARG checks that com_{ct} opens to ciphertexts ct_i, ct_{j_L}, ct_{j_R} where ct_i is the sum of ciphertexts ct_{j_L} and ct_{j_R} .
 - **Root node:** For the root node, the BARG checks that com_{ct} opens to ct_{root} .
- To test whether a digest $\text{dig} = (ct_{\text{root}}, \text{com}_{\text{ct}})$ is matching or not, the Extract algorithm outputs Matching if ct_{root} decrypts to 0 and NotMatching otherwise.

By definition, the ciphertext ct_{root} in any (honestly-generated) hash digest is the sum of the ciphertexts associated with the leaves of the ciphertext-evaluation tree. On an input x , if $x_i = 0$, then the associated ciphertext is an encryption of 0 and does not contribute to the sum. If $x_i = 1$, then the ciphertext associated with the leaf is an encryption of 1 if $i \in S$ and encryption of 0 otherwise. Thus, the sum is only incremented if $x_i = 1$ for some $i \in S$. This is precisely when Extract outputs NotMatching (i.e., the digest is for an input x where $x_i = 1$ for $i \in S$).

To argue that it is hard to open a Matching, but possibly-malformed digest to a 1 at an index $i \in S$, we appeal to soundness of the BARG. In this case, the root ciphertext ct_{root} in dig decrypts to a non-zero value, and yet the user constructed a valid BARG proof of opening for an index $i \in S$. The key observation is that the structure of the BARG used in the above construction is very similar to the structure of the canonical protocol from [\[BBK⁺23\]](#) described at the beginning of [Section 1.2](#) for demonstrating correct evaluation of a monotone circuit. Moreover, because the ciphertext-evaluation tree is *perfectly* balanced, it has depth $\log n$, where $n = \text{poly}(\lambda)$ is the input length. As such, we are able to adapt the proof strategy for arguing soundness of the monotone policy BARGs for *log-depth circuits* to directly argue zero-fixing security of our hash function. Specifically, we rely on BARG security to ensure that if the adversary uses an encryption of 1 as one of the leaves to the ciphertext (which it must if it opens an index $i \in S$ to a 1), then the root ciphertext necessarily is an encryption of a non-zero value. We provide the full details in [Section 5.1.3](#).

While the core construction described here satisfies zero-fixing security, we need to augment the construction to satisfy the additional security requirements we impose on a zero-fixing hash function. We summarize these here, and defer to the technical sections ([Sections 5](#), [5.1.4](#) and [5.1.5](#)) for the full details:

- **Extractor validity:** Recall that this property says that when the hash function is zero-fixing on the empty set, it should be hard for an adversary to come up with a “valid” digest that is NotMatching. To satisfy this property, we simply include a BARG proof of validity to the digest, where the BARG proof of validity simply checks that the ciphertext-evaluation tree was correctly constructed. When the hash key is binding to the empty set, all of the ciphertexts ct_i are an encryption of 0, so the root of a properly computed ciphertext-evaluation tree will also be an encryption of 0. We provide the details in [Section 5.1.4](#).
- **Index hiding with extracted guess:** Recall that this property says that the adversary cannot produce a digest dig where the extractor output disagrees depending on whether the hash key is zero-fixing on a set S or a set $S \setminus \{i\}$ (provided that the adversary provides an opening to 0 for index i). The only difference between the hash keys in these two cases is ct_i in the CRS changes from an encryption of 0 to an encryption of 1, which we could in principle show using semantic security. However, the reduction algorithm would have no way of checking whether a digest dig output by the adversary is Matching or NotMatching (since it does *not* and cannot know the decryption key). Thus, to argue this we adopt a Naor-Yung type of strategy [\[NY90\]](#) and *encrypt* twice. Namely, we introduce two *parallel* copies of the scheme (i.e., two independent public keys and two independent sets of ciphertexts). The digest now consists of two ciphertexts $ct_{root}^{(0)}, ct_{root}^{(1)}$ for the roots of the two ciphertext-evaluation trees. The same BARG would validate both roots. The key idea now is we can switch $ct_i^{(0)}$ from an encryption of 0 (i.e., zero-fixing at $S \setminus \{i\}$) to an encryption of 1 (i.e., zero-fixing at S) while being able to decrypt (i.e., extract) for the parallel encryption scheme. We can leverage soundness of the BARG to argue that for a valid digest/opening, both $ct_{root}^{(0)}$ and $ct_{root}^{(1)}$ encrypt *identical* values. This allows us to leverage semantic security to switch the ciphertexts for one scheme while being able to detect whether the output of Extract changed or not (using knowledge of the secret key for the parallel scheme). We provide the full details in [Section 5.1.5](#).

Taken together, we obtain a zero-fixing hash function from any standard BARG together with an additively-homomorphic encryption scheme. By instantiating with BARGs from the k -Lin assumption over groups with bilinear maps [\[WW22\]](#) or the (sub-exponential) DDH assumption over pairing-free groups [\[CGJ⁺23\]](#), we obtain zero-fixing hash functions from the same underlying assumptions. In conjunction with our generic construction from above, this yields [Theorem 1.1](#).

An algebraic construction of zero-fixing hash functions. As another contribution, we also describe an algebraic approach to construct zero-fixing hash functions directly from (composite-order) bilinear maps. This construction has the advantage that it only makes black-box use of cryptography. We give a brief sketch of the construction here, but defer the details to [Section 6](#). The basic version is an adaptation of the Catalano-Fiore vector commitment [\[CF13\]](#):

- Let $\mathcal{G} = (\mathbb{G}, \mathbb{G}_T, N, g, e)$ be a composite-order bilinear group of order N , generator g , and an efficiently-computable non-degenerate bilinear map $e: \mathbb{G} \times \mathbb{G} \rightarrow \mathbb{G}_T$. In the actual construction, we will require that N be a product of six primes. In the description here, we will just describe the basic scheme that operates primarily in just two subgroups. Let g_1 and g_2 be generators of two orthogonal subgroups of \mathbb{G} .
- To sample a hash key for a set $S \subseteq [n]$, the setup algorithm samples exponents $\alpha_i, \beta_i \xleftarrow{\mathbb{R}} \mathbb{Z}_N$. If $i \in S$, it sets $A_i \leftarrow (g_1 g_2)^{\alpha_i}$ and if $i \notin S$, it sets $A_i \leftarrow g_1^{\alpha_i}$. It sets $B_i \leftarrow g_1^{\beta_i}$ and for $i \neq j$, it computes the cross term $C_{i,j} \leftarrow g_1^{\alpha_i \beta_j}$. The hash key then contains A_i, B_i for $i \in [n]$ and $C_{i,j}$ for all $i \neq j$.
- The hash of an input $x \in \{0, 1\}^n$ is then $dig = \prod_{i \in [n]} A_i^{x_i}$. The opening to an index i is $V = \prod_{j \neq i} C_{j,i}^{x_j}$. To verify an opening to a bit b at index i , the verifier checks

$$e(dig, B_i) = e(A_i, B_i)^b \cdot e(g_1, V).$$

- To check whether a digest dig is Matching or not, the extraction algorithm output Matching if $e(dig, g_2) = 1$ and NotMatching otherwise.

The basic principle is to move the “encoding elements” A_i for $i \in S$ to have a component in the span of g_2 . The components A_i for $i \notin S$ are only in the span of g_1 . Then, any digest that includes an index $i \in S$ will contain a non-zero element in the span of g_2 , and thus, be declared `NotMatching`. Arguing the security of this scheme is more delicate and will require introducing a number of additional randomizing components (and subgroups). We refer to [Section 6](#) for the details.

Constructing monotone policy aggregate multi-signatures Our final contribution is a construction of monotone policy aggregate multi-signatures. While previous construction of aggregate signatures relied on extractable BARGs [[WW22](#), [DGKV22](#)], a similar implication is possible by combining a non-adaptively-sound BARG together with a “puncturable signature” scheme (also called an all-but-one signature scheme) [[GVW19](#)]. We sketch our construction below, and provide the full details in [Section 7](#).

In a puncturable signature scheme, it is possible to puncture a verification key on a message m^* . The property is that there does not *exist* signatures on m^* that verify with respect to the punctured verification key. Moreover, a punctured verification key is computationally indistinguishable from an honestly-generated verification key, even if the adversary is able to see signatures on arbitrary messages $m \neq m^*$. Goyal, Vusirikala, and Waters [[GVW19](#)] showed how to construct puncturable signatures from most standard number-theoretic assumptions (e.g., RSA, pairing-based assumptions, and LWE). We can use a non-adaptively-sound monotone policy BARG together with a puncturable signature scheme to construct a (statically-secure)⁶ aggregate multi-signature scheme for any policy computed by a monotone Boolean circuit. We provide a sketch below:

- **Setup:** Consider a scheme with k signers. Each signer $i \in [k]$ has a signing key sk_i and a verification key vk_i for the punctured signature scheme. The public parameters of the aggregation scheme contain the common reference string for a monotone policy BARG.
- **Signing:** To sign a message m , each user signs with their individual signing key.
- **Aggregation:** Given a set of signatures $\{\sigma_i\}_{i \in S}$ on the same message m and a (monotone) aggregation policy P , a user can aggregate the signatures by giving a monotone policy BARG proof for the policy P with respect to the natural relation $\mathcal{R}[m] = \{(vk, \sigma) : \text{Verify}(vk, m, \sigma)\}$. The aggregate signature is simply the BARG proof for the statements (vk_1, \dots, vk_k) with the witness $(\sigma_1, \dots, \sigma_k)$.
- **Verification:** To verify an aggregate multi-signature with respect to a policy P , the verifier just checks the BARG proof.

Note that one could also construct an aggregate multi-signature by sending the set S where $P(S) = 1$ and then use a vanilla BARG to prove knowledge of a signature σ_i for every $i \in S$. However, this approach would require communicating the set S as part of the aggregate signature. Using monotone policy BARGs, the aggregate signature only consists of the BARG proof, and thus has size, $\text{poly}(\lambda, \log |P|)$. It is straightforward to prove static security of the above multi-signature scheme just assuming *non-adaptive-soundness* on the underlying BARG. We sketch the reduction below:

- In the static security game, the adversary has to pre-commit to the message m^* it wants to forge on, the set of verification keys (vk_1^*, \dots, vk_k^*) it wants to use (which can be a mix of honest verification keys chosen by the challenger and verification keys chosen adversarially), and the aggregation policy P *before* seeing the aggregation parameters.
- Let $S \subseteq [k]$ be the set of indices i where the chosen key vk_i^* is uncorrupted (i.e., chosen by the challenger). The admissibility requirement is that $P(b_1, \dots, b_k) = 0$ where $b_i = 0$ if $i \in S$ and $b_i = 1$ otherwise; this is saying that the adversary cannot satisfy the policy P just by providing signatures under keys it controls.
- In the security reduction, we first puncture the honest users’ verification keys vk_i on the challenge message m^* . This means that there does not exist valid signatures on the challenge message m^* with respect to the honest users’ verification keys vk_i

⁶In the static security model, we require that the adversary declare the set of corrupted verification keys, its challenge message, and the aggregation policy at the beginning of the security game.

- Consider the relation $\mathcal{R}[m^*]$ used for verification. By definition of the set S and the fact that the honest verification keys are punctured at m^* , the statement (vk_1^*, \dots, vk_k^*) is *false* for the policy P with respect to the relation $\mathcal{R}[m^*]$. By non-adaptive soundness of the monotone policy BARG, the probability that the adversary can produce a valid aggregate signature (i.e., a valid proof on a false statement) is negligible.

Observe that in the above sketch, the verification time is linear in k . However, using a RAM delegation scheme, we can achieve fast verification. We refer to [Remark 7.8](#) for additional details.

2 Preliminaries

Throughout this work, we write λ to denote the security parameter. For $n \in \mathbb{N}$, we write $[n]$ to denote the set $\{1, \dots, n\}$. For $a, b \in \mathbb{N}$ we write $[a, b]$ to denote the set $\{a, a+1, \dots, b\}$. We write $\text{poly}(\lambda)$ to denote a function that is bounded by a fixed polynomial in λ , and $\text{negl}(\lambda)$ to denote a function that is $o(\lambda^{-c})$ for all $c \in \mathbb{N}$. We say an event happens with overwhelming probability if its complement occurs with negligible probability. For a finite set S , we write $x \stackrel{\mathbb{R}}{\leftarrow} S$ to denote that x is a uniformly random element of S . For a distribution \mathcal{D} we write $x \leftarrow \mathcal{D}$ to denote that x is a random draw from \mathcal{D} .

We say an algorithm is efficient if it runs in probabilistic polynomial time in the length of its input. A non-uniform algorithm \mathcal{A} consists of a pair of algorithms $(\mathcal{A}_1, \mathcal{A}_2)$ where \mathcal{A}_1 is a (possibly-unbounded) algorithm that takes as input 1^λ and outputs an advice string ρ_λ of $\text{poly}(\lambda)$ size. Algorithm \mathcal{A}_2 is an efficient algorithm. The output of \mathcal{A} on an input $x \in \{0, 1\}^\lambda$ is defined as first computing the advice string $\rho_\lambda \leftarrow \mathcal{A}_1(1^\lambda)$ and then outputting $\mathcal{A}_2(x, \rho_\lambda)$. We say two ensembles of distributions $\mathcal{D}_1 = \{\mathcal{D}_{1,\lambda}\}_{\lambda \in \mathbb{N}}$ and $\mathcal{D}_2 = \{\mathcal{D}_{2,\lambda}\}_{\lambda \in \mathbb{N}}$ are computationally indistinguishable if no efficient adversary can distinguish them with non-negligible probability. We say they are statistically indistinguishable if their statistical distance is bounded by $\text{negl}(\lambda)$.

2.1 Cryptographic Building Blocks

In this section, we recall the definition of a few standard cryptographic building blocks we use in this work.

Additively-homomorphic encryption. We start by reviewing the notion of an additively homomorphic encryption. For our applications, it suffices to consider constructions that only support decryption of values residing in a bounded message space. Such additively homomorphic encryption schemes can be built from most standard number-theoretic assumptions that imply public-key encryption such as the decisional Diffie-Hellman (DDH) assumption [[Gam84](#)], decisional composite residuosity (DCR) [[Pai99](#)], or the learning with errors (LWE) assumption [[Reg05](#)].

Definition 2.1 (Additively Homomorphic Encryption). An additively homomorphic encryption with bounded support is a tuple of polynomial time algorithms $\Pi_{\text{HE}} = (\text{Gen}, \text{Enc}, \text{Dec}, \text{Add})$ with the following syntax:

- $\text{Gen}(1^\lambda, 1^n) \rightarrow (\text{sk}, \text{pk})$: On input a security parameter $\lambda \in \mathbb{N}$ and a range parameter $n \in \mathbb{N}$, the key-generation algorithm outputs a secret key sk and a public key pk . We assume that the secret key and the public key includes an implicit description of the range parameter 1^n .
- $\text{Enc}(\text{pk}, \text{msg}) \rightarrow \text{ct}$: On input a public key pk and an integer $\text{msg} \in \{0, \dots, n\}$, the encryption algorithm outputs a ciphertext ct .
- $\text{Dec}(\text{sk}, \text{ct}) \rightarrow \text{msg}$: On input a secret key sk and a ciphertext ct , the decryption algorithm either outputs a plaintext $\text{msg} \in \{0, \dots, n\}$ or a special symbol $\text{msg} = \perp$. The decryption algorithm is deterministic.
- $\text{Add}(\text{pk}, \text{ct}_1, \text{ct}_2) \rightarrow \text{ct}'$: On input a public key pk and two ciphertexts ct_1, ct_2 , the homomorphic addition algorithm outputs a new ciphertext ct' . The addition algorithm is *deterministic*.

We require the following properties:

- **Correctness:** For all $\lambda, n \in \mathbb{N}$ and all messages $\text{msg} \in \{0, \dots, n\}$, it holds that:

$$\Pr \left[\text{Dec}(\text{sk}, \text{ct}) = \text{msg} : \begin{array}{l} (\text{sk}, \text{pk}) \leftarrow \text{Gen}(1^\lambda, 1^n) \\ \text{ct} \leftarrow \text{Enc}(\text{pk}, \text{msg}) \end{array} \right] = 1.$$

- **Evaluation correctness:** For all $\lambda, n \in \mathbb{N}$, all (sk, pk) in the support of $\text{Gen}(1^\lambda, 1^n)$ and all ciphertexts ct_1, ct_2 where $\text{Dec}(\text{sk}, \text{ct}_1) \neq \perp$, $\text{Dec}(\text{sk}, \text{ct}_2) \neq \perp$, and $\text{Dec}(\text{sk}, \text{ct}_1) + \text{Dec}(\text{sk}, \text{ct}_2) \in \{0, \dots, n\}$, it holds that

$$\text{Dec}(\text{sk}, \text{Add}(\text{pk}, \text{ct}_1, \text{ct}_2)) = \text{Dec}(\text{sk}, \text{ct}_1) + \text{Dec}(\text{sk}, \text{ct}_2).$$

- **Compactness:** There exists a polynomial $\text{poly}(\cdot)$ such that for all $\lambda, n \in \mathbb{N}$, all (sk, pk) in the support of $\text{Gen}(1^\lambda, 1^n)$, all ciphertexts ct in the support of $\text{Enc}(\text{pk}, \cdot)$ and $\text{Add}(\text{pk}, \cdot, \cdot)$, it holds that $|\text{pk}| \leq \text{poly}(\lambda + \log n)$ and $|\text{ct}| \leq \text{poly}(\lambda + \log n)$.
- **CPA-security:** For an adversary \mathcal{A} and a bit $b \in \{0, 1\}$, define the CPA-security experiment $\text{ExptSS}_{\mathcal{A}}(\lambda, b)$ as follows:

1. On input the security parameter 1^λ , the adversary \mathcal{A} starts by outputting a range parameter 1^n .
2. The challenger samples a key pair $(\text{sk}, \text{pk}) \leftarrow \text{Gen}(1^\lambda, 1^n)$ and sends pk to the adversary.
3. The adversary can now make (arbitrarily many) queries on pairs of messages $(\text{msg}_0, \text{msg}_1)$. On each query, the challenger replies with $\text{Enc}(\text{pk}, \text{msg}_b)$.
4. After the adversary \mathcal{A} is done making queries, it outputs a guess $b' \in \{0, 1\}$.

We say that Π_{HE} is semantically secure if for every efficient adversary \mathcal{A} , there exists a negligible function $\text{negl}(\cdot)$ such that $|\Pr[\text{ExptSS}_{\mathcal{A}}(\lambda, 1) = 1] - \Pr[\text{ExptSS}_{\mathcal{A}}(\lambda, 0) = 1]| = \text{negl}(\lambda)$.

Fact 2.2 (Additively Homomorphic Encryption [Gam84, Pai99, Reg05]). Assuming any of (1) the decisional Diffie-Hellman assumption (DDH), (2) the decisional composite residuosity assumption (DCR), or (3) the learning with errors (LWE) assumption, there exists an additively homomorphic encryption scheme with a bounded support.

Vector commitments. Next, we recall the notion of a vector commitment scheme with succinct local openings. Such commitments can be built from any collision-resistant hash function [Mer87].

Definition 2.3 (Vector Commitment). A vector commitment with local openings is a tuple of efficient algorithms $\Pi_{\text{Com}} = (\text{Setup}, \text{Commit}, \text{Verify})$ with the following properties:

- $\text{Setup}(1^\lambda, 1^n, \ell) \rightarrow \text{crs}$: On input the security parameter $\lambda \in \mathbb{N}$, the block length $n \in \mathbb{N}$, and the vector length $\ell \in \mathbb{N}$ (in *binary*), the setup algorithm outputs a common reference string crs . We assume the common reference string implicitly contains the parameters 1^n and ℓ .
- $\text{Commit}(\text{crs}, (x_1, \dots, x_t)) \rightarrow (\text{com}, \sigma_1, \dots, \sigma_t)$: On input the common reference string crs and a vector of $t \leq \ell$ messages $x_1, \dots, x_t \in \{0, 1\}^n$, the commit algorithm outputs a commitment com and openings $\sigma_1, \dots, \sigma_t$.
- $\text{Verify}(\text{crs}, \text{com}, i, y, \sigma) \rightarrow b'$: On input the common reference string crs , the commitment com , an index $i \in [\ell]$, a message $y \in \{0, 1\}^n$, and an opening σ , the verification algorithm outputs a bit $b' \in \{0, 1\}$.

Moreover, Π_{Com} should satisfy the following properties:

- **Correctness:** For all $\lambda, n, \ell \in \mathbb{N}$, and all positive $t \leq \ell$, all $x_1, \dots, x_t \in \{0, 1\}^n$, and indices $i \in [t]$,

$$\Pr \left[\text{Verify}(\text{crs}, \text{com}, i, x_i, \sigma_i) = 1 : \begin{array}{l} \text{crs} \leftarrow \text{Setup}(1^\lambda, 1^n, \ell), \\ (\text{com}, \sigma_1, \dots, \sigma_t) \leftarrow \text{Commit}(\text{crs}, (x_1, \dots, x_t)) \end{array} \right] = 1.$$

- **Computational binding:** For an adversary \mathcal{A} , define the computational binding experiment as follows:

1. On input the security parameter 1^λ , algorithm \mathcal{A} starts by outputting the block length 1^n and vector length ℓ .
2. The challenger responds with $\text{crs} \leftarrow \text{Setup}(1^\lambda, 1^n, \ell)$.
3. Algorithm \mathcal{A} outputs a commitment com , an index $i \in [\ell]$, and openings (y_0, σ_0) and (y_1, σ_1) .

4. The output of the experiment is $b = 1$ if $\text{Verify}(\text{crs}, \text{com}, i, y_0, \sigma_0) = 1 = \text{Verify}(\text{crs}, \text{com}, i, y_1, \sigma_1)$ and $y_0 \neq y_1$. Otherwise, the output is $b = 0$.

The commitment scheme is binding if for all efficient adversaries \mathcal{A} , there exists a negligible function $\text{negl}(\cdot)$ such that $\Pr[b = 1] = \text{negl}(\lambda)$ in the binding experiment.

- **Succinctness:** There exists a universal polynomial $\text{poly}(\cdot)$ such that for all $\lambda, n, \ell \in \mathbb{N}$, all crs in the support of $\text{Setup}(1^\lambda, 1^n, \ell)$, all $t \leq \ell$, and all $(\text{com}, \sigma_1, \dots, \sigma_t)$ in the support of $\text{Commit}(\text{crs}, \cdot)$, the following holds:
 - **Succinct CRS:** $|\text{crs}| = \text{poly}(\lambda + \log n + \log \ell)$.
 - **Succinct commitment:** $|\text{com}| = \text{poly}(\lambda + \log n + \log \ell)$.
 - **Succinct local opening:** For all $i \in [\ell]$, $|\sigma_i| = \text{poly}(\lambda + \log n + \log \ell)$.

Fact 2.4 (Vector Commitments from Collision-Resistant Hash Functions [Mer87]). Assuming the existence of collision-resistant hash functions, there exists a vector commitment scheme with local openings.

2.2 Batch Arguments for NP

In this section, we recall the notion of a non-interactive batch argument (BARG) for NP, the special case of a BARG for index languages (i.e., an “index BARG” [CJJ21b]) and the notion of a BARG for monotone policy batch NP [BBK⁺23].

Batch arguments for NP. We begin with the notion of a somewhere extractable batch argument for NP. Our presentation is adapted from [CJJ21b, WW22]. Here, we provide a more general syntax where the batch arguments supports extraction on up to ℓ indices.

Definition 2.5 (Boolean Circuit Satisfiability). We define the circuit satisfiability language $\mathcal{L}_{\text{CSAT}}$ as

$$\mathcal{L}_{\text{CSAT}} = \left\{ (C, x) \mid \begin{array}{l} C: \{0, 1\}^n \times \{0, 1\}^h \rightarrow \{0, 1\}, x \in \{0, 1\}^n \\ \exists w \in \{0, 1\}^* : C(x, w) = 1 \end{array} \right\}.$$

Definition 2.6 (Non-Interactive Batch Argument). A somewhere extractable non-interactive batch argument (BARG) for Boolean circuit satisfiability is a tuple of efficient algorithms $\Pi_{\text{BARG}} = (\text{Gen}, \text{Prove}, \text{Verify}, \text{TrapGen}, \text{Extract})$ with the following syntax:

- $\text{Gen}(1^\lambda, 1^k, 1^n, 1^s, 1^\ell) \rightarrow (\text{crs}, \text{vk})$: On input the security parameter $\lambda \in \mathbb{N}$, the number of instances $k \in \mathbb{N}$, the instance length $n \in \mathbb{N}$, a bound on the size of the Boolean circuit $s \in \mathbb{N}$, and a bound on the size of the extraction set $\ell \in \mathbb{N}$, the generator algorithm outputs a common reference string crs and a verification key vk .
- $\text{Prove}(\text{crs}, C, (x_1, \dots, x_k), (w_1, \dots, w_k)) \rightarrow \pi$: On input the common reference string crs , a Boolean circuit $C: \{0, 1\}^n \times \{0, 1\}^h \rightarrow \{0, 1\}$, statements $x_1, \dots, x_k \in \{0, 1\}^k$, and witnesses $w_1, \dots, w_k \in \{0, 1\}^h$, the prove algorithm outputs a proof π .
- $\text{Verify}(\text{vk}, C, (x_1, \dots, x_k), \pi) \rightarrow b$: On input the verification key vk , a Boolean circuit $C: \{0, 1\}^n \times \{0, 1\}^h \rightarrow \{0, 1\}$, statements $x_1, \dots, x_k \in \{0, 1\}^n$ and a proof π , the verification algorithm outputs a bit $b \in \{0, 1\}$.
- $\text{TrapGen}(1^\lambda, 1^k, 1^n, 1^s, 1^\ell, S) \rightarrow (\text{crs}, \text{vk}, \text{td})$: On input the security parameter $\lambda \in \mathbb{N}$, the number of instances $k \in \mathbb{N}$, the instance size $n \in \mathbb{N}$, a bound on the size of the Boolean circuit $s \in \mathbb{N}$, a bound on the size of the extraction set $\ell \in \mathbb{N}$, and a set $S \subseteq [k]$ of size at most ℓ , the trapdoor generator algorithm outputs a common reference string crs , a verification key vk and an extraction trapdoor td .
- $\text{Extract}(\text{td}, C, (x_1, \dots, x_k), \pi, i) \rightarrow w$. On input the trapdoor td , a Boolean circuit $C: \{0, 1\}^n \times \{0, 1\}^h \rightarrow \{0, 1\}$, a collection of statements $x_1, \dots, x_k \in \{0, 1\}^n$, a proof π and an index $i \in [k]$, the extraction algorithm outputs a witness w .

For notational convenience, when $\ell = 1$, we omit the final input 1^ℓ and instead, write $\text{Gen}(1^\lambda, 1^k, 1^n, 1^s)$ to denote $\text{Gen}(1^\lambda, 1^k, 1^n, 1^s, 1^1)$. Similarly, we write $\text{TrapGen}(1^\lambda, 1^k, 1^n, 1^s, i)$ to denote $\text{TrapGen}(1^\lambda, 1^k, 1^n, 1^s, 1^1, \{i\})$. Finally, we require that Π_{BARG} satisfy the following properties:

- **Completeness:** For all $\lambda, k, n, s, \ell \in \mathbb{N}$, all Boolean circuits $C: \{0, 1\}^n \times \{0, 1\}^h \rightarrow \{0, 1\}$ of size at most s , all statements $x = (x_1, \dots, x_k) \in \{0, 1\}^{kn}$ and witnesses $w = (w_1, \dots, w_k) \in \{0, 1\}^{kh}$ where $C(x_i, w_i) = 1$ for all $i \in [k]$,

$$\Pr \left[\text{Verify}(\text{vk}, C, (x_1, \dots, x_k), \pi) = 1 : \begin{array}{l} (\text{crs}, \text{vk}) \leftarrow \text{Gen}(1^\lambda, 1^k, 1^n, 1^s, 1^\ell) \\ \pi \leftarrow \text{Prove}(\text{crs}, C, x, w) \end{array} \right] = 1$$

- **Set hiding:** For an adversary \mathcal{A} and a bit $b \in \{0, 1\}$, define the set hiding experiment $\text{ExptSH}_{\mathcal{A}}(\lambda, b)$ as follows:

1. On input the security parameter 1^λ , algorithm \mathcal{A} starts by outputting the number of instances 1^k , the instance size 1^n , the bound on the circuit size 1^s , the bound on the size of the extraction set 1^ℓ , and a set $S \subseteq [k]$ of size at most ℓ .
2. If $b = 0$, the challenger gives $(\text{crs}, \text{vk}) \leftarrow \text{Gen}(1^\lambda, 1^k, 1^n, 1^s, 1^\ell)$ to \mathcal{A} . If $b = 1$, the challenger samples $(\text{crs}, \text{vk}, \text{td}) \leftarrow \text{TrapGen}(1^\lambda, 1^k, 1^n, 1^s, 1^\ell, S)$ and gives (crs, vk) to \mathcal{A} .
3. Algorithm \mathcal{A} outputs a bit $b' \in \{0, 1\}$, which is the output of the experiment.

Then, Π_{BARG} satisfies set hiding if for every efficient adversary \mathcal{A} , there exists a negligible function $\text{negl}(\cdot)$ such that

$$\left| \Pr[\text{ExptSH}_{\mathcal{A}}(\lambda, 0) = 1] - \Pr[\text{ExptSH}_{\mathcal{A}}(\lambda, 1) = 1] \right| = \text{negl}(\lambda).$$

When $\ell = 1$, we might refer to this property as *index hiding*.

- **Somewhere extractable in trapdoor mode:** For an adversary \mathcal{A} , define the somewhere extractable security game as follows:

1. On input the security parameter 1^λ , algorithm \mathcal{A} starts by outputting the number of instances 1^k , the instance size 1^n , the bound on the circuit size 1^s , a bound on the size of the extraction set 1^ℓ , and a nonempty set $S \subseteq [k]$ of size at most ℓ .
2. The challenger samples $(\text{crs}, \text{vk}, \text{td}) \leftarrow \text{TrapGen}(1^\lambda, 1^k, 1^n, 1^s, 1^\ell, S)$ and gives (crs, vk) to \mathcal{A} .
3. Algorithm \mathcal{A} outputs a Boolean circuit $C: \{0, 1\}^n \times \{0, 1\}^h \rightarrow \{0, 1\}$ of size at most s , statements $x_1, \dots, x_m \in \{0, 1\}^n$, and a proof π .
4. The output of the game is $b = 1$ if $\text{Verify}(\text{vk}, C, (x_1, \dots, x_m), \pi) = 1$ and there exists an index $i \in S$ for which $C(x_i, w_i) \neq 1$ where $w_i \leftarrow \text{Extract}(\text{td}, C, (x_1, \dots, x_k), \pi, i)$. Otherwise, the output is $b = 0$.

Then Π_{BARG} is somewhere extractable in trapdoor mode if for every adversary \mathcal{A} , there exists a negligible function $\text{negl}(\cdot)$ such that $\Pr[b = 1] = \text{negl}(\lambda)$ in the somewhere extractable game.

- **Succinctness:** There exists a fixed polynomial $\text{poly}(\cdot)$ such that for all $\lambda, k, n, s, \ell \in \mathbb{N}$, all crs in the support of $\text{Gen}(1^\lambda, 1^k, 1^n, 1^s, 1^\ell)$, and all Boolean circuits $C: \{0, 1\}^n \times \{0, 1\}^h \rightarrow \{0, 1\}$ of size at most s , the following properties hold:

- **Succinct proofs:** The proof π output by $\text{Prove}(\text{crs}, C, \cdot, \cdot)$ satisfies $|\pi| \leq \text{poly}(\lambda + \log k + s + \ell)$.
- **Succinct CRS:** $|\text{crs}| \leq \text{poly}(\lambda + k + n + \ell) + \text{poly}(\lambda + \log k + s + \ell)$.
- **Succinct verification key:** $|\text{vk}| \leq \text{poly}(\lambda + \log k + s + \ell)$.

Fact 2.7 (Batch Arguments for NP [CJJ21b, WW22, KLVW23, CGJ⁺23]). Assuming any of (1) the plain LWE assumption, (2) the k -Lin assumption over pairing groups for any constant $k \in \mathbb{N}$, or (3) the (sub-exponential) DDH assumption in pairing-free groups, there exists a non-interactive batch argument for NP.

Set hiding with extraction. For our main construction (Section 5), we require a slight strengthening of the somewhere extractability property from Definition 2.6. Our stronger set-hiding property essentially says that if the extraction key is programmed to extract either on $S_0 \subseteq [k]$ or $S_1 \subseteq [k]$, then the extracted witness on “common indices” $i^* \in S_0 \cap S_1$ is computationally indistinguishable in the two cases. This type of property is often referred to as a “no-signaling” extraction property [PR17, KPY19, GZ21, KVZ21, CJJ21b]. We define this formally below and show that it follows generically from the standard vanilla extractability in Appendix A.

Definition 2.8 (Set Hiding with Extraction). Let $\Pi_{\text{BARG}} = (\text{Gen}, \text{Prove}, \text{Verify}, \text{TrapGen}, \text{Extract})$ be a somewhere extractable batch argument for Boolean circuit satisfiability (Definition 2.6). For an adversary \mathcal{A} and a bit $b \in \{0, 1\}$, define the set hiding with extraction experiment $\text{ExptSHwE}_{\mathcal{A}}(\lambda, b)$ as follows:

1. On input the security parameter λ , algorithm \mathcal{A} starts by outputting the number of instances 1^k , the instance length 1^n , the bound on the circuit size 1^s , the bound on the extraction set 1^ℓ , a set $S \subseteq [k]$ of size at most ℓ , and an index $i^* \in S$.
2. If $b = 0$, the challenger samples $(\text{crs}, \text{vk}, \text{td}) \leftarrow \text{TrapGen}(1^\lambda, 1^k, 1^n, 1^s, 1^\ell, S)$. If $b = 1$, the challenger samples $(\text{crs}, \text{vk}, \text{td}) \leftarrow \text{TrapGen}(1^\lambda, 1^k, 1^n, 1^s, 1^\ell, \{i^*\})$. The challenger replies to \mathcal{A} with (crs, vk) .
3. Algorithm \mathcal{A} outputs a Boolean circuit $C: \{0, 1\}^n \times \{0, 1\}^h \rightarrow \{0, 1\}$, statements $x_1, \dots, x_k \in \{0, 1\}^n$, and a proof π .
4. If $\text{Verify}(\text{vk}, C, (x_1, \dots, x_k), \pi) \neq 1$, then the experiment halts with output 0. Otherwise, the challenger replies with $w^* \leftarrow \text{Extract}(\text{td}, C, (x_1, \dots, x_k), \pi, i^*)$.
5. Algorithm \mathcal{A} outputs a bit $b' \in \{0, 1\}$, which is the output of the experiment.

Then, Π_{BARG} satisfies set hiding with extraction if for every efficient adversary \mathcal{A} , there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$,

$$|\Pr[\text{ExptSHwE}_{\mathcal{A}}(\lambda, 0) = 1] - \Pr[\text{ExptSHwE}_{\mathcal{A}}(\lambda, 1) = 1]| = \text{negl}(\lambda).$$

Index BARGs. An index BARG [CJJ21b] is a batch argument for the batch index language where the instance is *always* the tuple $(1, \dots, k)$. Since the statements are the integers, they have a succinct description, so we can impose a stronger requirement on the running time of the Verify algorithm. We define this below:

Definition 2.9 (Index BARG [CJJ21b]). An index BARG is a special case of a BARG where the instances (x_1, \dots, x_k) are restricted to the integers $(1, \dots, k)$. In this setting, the Gen algorithm to the index BARG does *not* separately take in the instance length n as a separate input. Moreover, instead of providing x_1, \dots, x_k as input to the Prove, Verify, and Extract algorithms, we just give the single index k (in binary). Moreover, we require the *additional* succinctness property on the running time of Verify:

- **Succinct verification time:** There exists a fixed polynomial $\text{poly}(\cdot)$ such that for all $\lambda, k, n, s, \ell \in \mathbb{N}$, all (crs, vk) in the support of $\text{Gen}(1^\lambda, 1^k, 1^n, 1^s, 1^\ell)$ and all Boolean circuits $C: \{0, 1\}^n \times \{0, 1\}^h \rightarrow \{0, 1\}$ of size at most s , the running time of $\text{Verify}(\text{vk}, C, k, \cdot)$ is bounded by $\text{poly}(\lambda + \log k + s + \ell)$.

Monotone policy BARG. Next, we recall the notion of a SNARG for monotone policy BatchNP [BBK⁺23], which we refer to more succinctly as a “monotone policy BARG.” In this work, we just focus on the simplest notion of non-adaptive soundness.

Definition 2.10 (Monotone Policy BatchNP). A Boolean circuit $P: \{0, 1\}^k \rightarrow \{0, 1\}$ is a monotone Boolean policy if P is a Boolean circuit comprised entirely of AND and OR gates. Let $C: \{0, 1\}^n \times \{0, 1\}^h \rightarrow \{0, 1\}$ be a Boolean circuit and $P: \{0, 1\}^k \rightarrow \{0, 1\}$ be a monotone Boolean policy. We define the monotone policy BatchNP language $\mathcal{L}_{\text{MP-CSAT}}$ to be

$$\mathcal{L}_{\text{MP-CSAT}} = \left\{ (C, P, x_1, \dots, x_k) \mid \begin{array}{l} C: \{0, 1\}^n \times \{0, 1\}^h \rightarrow \{0, 1\}, P: \{0, 1\}^k \rightarrow \{0, 1\}, x_1, \dots, x_k \in \{0, 1\}^n \\ \exists w_1, \dots, w_k \in \{0, 1\}^h : P(C(x_1, w_1), \dots, C(x_k, w_k)) = 1 \end{array} \right\}.$$

Definition 2.11 (Monotone Policy BARG [BBK⁺23, adapted]). A monotone policy BARG is a tuple $\Pi_{\text{MP-BARG}} = (\text{Gen}, \text{Prove}, \text{Verify})$ of efficient algorithms with the following syntax:

- $\text{Gen}(1^\lambda, 1^n, 1^{s_c}, 1^{s_p}) \rightarrow \text{crs}$: On input the security parameter $\lambda \in \mathbb{N}$, the instance size $n \in \mathbb{N}$, a bound on the size of the Boolean circuit $s_c \in \mathbb{N}$, and a bound on the size of the policy $s_p \in \mathbb{N}$, the generator algorithm outputs a common reference string crs .
- $\text{Prove}(\text{crs}, C, P, (x_1, \dots, x_k), (w_1, \dots, w_k)) \rightarrow \pi$: On input the common reference string crs , a Boolean circuit $C: \{0, 1\}^n \times \{0, 1\}^h \rightarrow \{0, 1\}$, a monotone Boolean policy $P: \{0, 1\}^k \rightarrow \{0, 1\}$, statements $x_1, \dots, x_k \in \{0, 1\}^n$, and witnesses $w_1, \dots, w_k \in \{0, 1\}^h$, the prove algorithm outputs a proof π .
- $\text{Verify}(\text{crs}, C, P, (x_1, \dots, x_k), \pi) \rightarrow b$: On input the common reference string crs , a Boolean circuit $C: \{0, 1\}^n \times \{0, 1\}^h \rightarrow \{0, 1\}$, a monotone Boolean policy $P: \{0, 1\}^k \rightarrow \{0, 1\}$, statements $x_1, \dots, x_k \in \{0, 1\}^n$, and a proof π , the verification algorithm outputs a bit $b \in \{0, 1\}$.

Moreover, $\Pi_{\text{MP-BARG}}$ should satisfy the following properties:

- **Completeness:** For all $\lambda, n, s_c, s_p \in \mathbb{N}$, Boolean circuits $C: \{0, 1\}^n \times \{0, 1\}^h \rightarrow \{0, 1\}$ of size at most s_c , monotone Boolean policies $P: \{0, 1\}^k \rightarrow \{0, 1\}$ of size at most s_p , statements $x = (x_1, \dots, x_k) \in \{0, 1\}^{kn}$ and witnesses $w = (w_1, \dots, w_k) \in \{0, 1\}^{kh}$ where $P(C(x_1, w_1), \dots, C(x_k, w_k)) = 1$, it holds that

$$\Pr \left[\text{Verify}(\text{crs}, C, P, (x_1, \dots, x_k), \pi) = 1 : \begin{array}{l} \text{crs} \leftarrow \text{Gen}(1^\lambda, 1^n, 1^{s_c}, 1^{s_p}) \\ \pi \leftarrow \text{Prove}(\text{crs}, C, P, x, w) \end{array} \right] = 1.$$

- **Non-adaptive soundness:** For any adversary \mathcal{A} , define the non-adaptive soundness game as follows:
 1. On input the security parameter 1^λ , algorithm \mathcal{A} starts by outputting the instance size 1^n , the bound on the size of the NP relation 1^{s_c} , the bound on the size of the policy 1^{s_p} , a Boolean circuit $C: \{0, 1\}^n \times \{0, 1\}^h \rightarrow \{0, 1\}$ of size at most s_c , a monotone Boolean circuit $P: \{0, 1\}^k \rightarrow \{0, 1\}$ of size at most s_p , and statements $x_1, \dots, x_k \in \{0, 1\}^n$.
 2. The challenger samples $\text{crs} \leftarrow \text{Gen}(1^\lambda, 1^n, 1^{s_c}, 1^{s_p})$ and gives it to \mathcal{A} .
 3. Algorithm \mathcal{A} outputs a proof π .
 4. The output of the game is $b = 1$ if $\text{Verify}(\text{crs}, C, P, (x_1, \dots, x_k), \pi) = 1$ and $(C, P, (x_1, \dots, x_k)) \notin \mathcal{L}_{\text{MP-CSAT}}$.

We say that $\Pi_{\text{MP-BARG}}$ is non-adaptively sound if for every efficient adversary \mathcal{A} , there exists a negligible function $\text{negl}(\cdot)$ such that $\Pr[b = 1] = \text{negl}(\lambda)$ in the non-adaptive soundness game.

- **Succinctness:** There exists a fixed polynomial $\text{poly}(\cdot)$ such that for all $\lambda, n, s_c, s_p \in \mathbb{N}$, all crs in the support of $\text{Gen}(1^\lambda, 1^n, 1^{s_c}, 1^{s_p})$, all Boolean circuits $C: \{0, 1\}^n \times \{0, 1\}^h \rightarrow \{0, 1\}$ of size at most s_c , and all monotone Boolean policies $P: \{0, 1\}^k \rightarrow \{0, 1\}$ of size $|P| \leq s_p$, the following properties hold:
 - **Slightly succinct proofs:** The proof π output by $\text{Prove}(\text{crs}, C, P, \cdot, \cdot)$ satisfies $|\pi| \leq \text{poly}(\lambda + s_c + \log s_p)$.
 - **Succinct proofs:** The proof π output by $\text{Prove}(\text{crs}, C, P, \cdot, \cdot)$ satisfies $|\pi| \leq \text{poly}(\lambda + s_c + \log |P|)$.

Remark 2.12 (Slightly Succinct Proofs to Succinct Proofs). In a “slightly succinct” proof system, the size of the proof scales logarithmically with the bound s_p on the size of the policy circuit, rather than the size of the policy circuit itself. It is straightforward to transform a scheme with slightly succinct proofs into one that satisfies the standard notion of succinctness. We use a “powers-of-two” construction. Namely, we generate $\ell = \lceil \log s_p \rceil$ different common reference strings, where the i^{th} CRS supports policies of size at most 2^i . The prover and verifier will use the CRS for scheme i when proving or verifying statements with respect to policies of size between 2^{i-1} and 2^i . In this case, the size of the proof scales polylogarithmically with the size of the policy P rather than the bound s_p . This approach only incurs logarithmic overhead in the CRS size. In the rest of this work, we will focus on constructions satisfying the simpler requirement of having slightly succinct proofs.

Remark 2.13 (Short Verification Key via RAM Delegation). In [Definition 2.11](#), the setup algorithm outputs a single CRS that is used both for generating proofs and for verifying proofs. The size of the CRS is allowed to grow with the size of both the circuit C and the size of the monotone policy P . It is possible to obtain a construction with a short verification key (that grows polylogarithmically with $|C|$ and $|P|$) by “delegating” the verification process using a RAM delegation scheme [[CJJ21b](#), [WW22](#), [KLVW23](#), [CGJ⁺23](#)]. In this case, the verification key would be a succinct hash of the actual CRS. Since this provides a generic approach for realizing a short verification key, in our constructions, we will not explicitly decompose the CRS into a proving key and a separate short verification key. A similar approach was also used in [[BBK⁺23](#)] in their construction of predicate-extractable hash functions.

3 Zero-Fixing Hash Functions

In this section, we formally introduce the notion of a zero-fixing hash function. As we show in [Section 4](#), we can combine a zero-fixing hash function with a vanilla BARG to obtain a monotone policy BARG. Recall from [Section 1.2](#) that a zero-fixing hash function is a keyed hash function that supports succinct local openings. Moreover, the hash key is associated with a set of indices $S \subseteq [n]$, where n is the input length. Moreover, there is a trapdoor td associated with the hash key hk that can be used to decide whether a hash digest dig is Matching or NotMatching on the set S . The zero-fixing security requirement then says that if the extractor outputs Matching for a digest dig , it must be computationally hard to open dig to a 1 on any index $i \in S$.

As discussed in [Section 1.2](#), our zero-fixing hash function is similar to the predicate-extractable hash function for bit-fixing predicates from [[BBK⁺23](#)]. A key distinction is that when the extraction algorithm outputs NotMatching, the predicate-extractable hash function also outputs an index $i \in [n]$ where it is computationally infeasible to open the digest to a 1. In contrast, with our zero-fixing hash function, the extraction algorithm only outputs a single Matching or NotMatching flag. At the same time, we require our zero-fixing hash functions to satisfy additional security requirements that were not required in [[BBK⁺23](#)]. These additional security properties are necessary for our construction of monotone policy BARGs ([Section 4](#)). We now give the formal definition:

Definition 3.1 (Zero-Fixing Hash Function). A zero-fixing hash function is a tuple of polynomial-time algorithms $\Pi_H = (\text{Setup}, \text{Hash}, \text{ProveOpen}, \text{VerOpen}, \text{Extract}, \text{ValidateDigest})$ with the following syntax:

- $\text{Setup}(1^\lambda, 1^n, S) \rightarrow (hk, vk, td)$: On input a security parameter λ , an input length n , and a set $S \subseteq [n]$, the setup algorithm outputs a hash key hk , a verification key vk and a trapdoor td . We implicitly assume that hk includes λ and n .
- $\text{Hash}(hk, x) \rightarrow dig$: On input a hash key hk and a string $x \in \{0, 1\}^n$, the hash algorithm outputs a digest dig . This algorithm is deterministic.
- $\text{ValidateDigest}(vk, dig) \rightarrow b$: On input a hash key vk and a digest dig , the digest validation algorithm outputs a bit $b \in \{0, 1\}$. This algorithm is deterministic.
- $\text{ProveOpen}(hk, x, i) \rightarrow \sigma$: On input a hash key hk , a string $x \in \{0, 1\}^n$ and an index $i \in [n]$, the prove algorithm outputs an opening σ .
- $\text{VerOpen}(vk, dig, i, b, \sigma) \rightarrow b'$: On input a hash key vk , a digest dig , an index $i \in [n]$, a bit $b \in \{0, 1\}$ and an opening σ , the verification algorithm outputs a bit $b' \in \{0, 1\}$. The verification algorithm is deterministic.
- $\text{Extract}(td, dig) \rightarrow m$: On input a trapdoor td and a digest dig , the extraction algorithm outputs a value $m \in \{\text{Matching}, \text{NotMatching}\}$. This algorithm is deterministic.

We require Π_H satisfy the following efficiency and correctness properties:

- **Succinctness:** There exists a universal polynomial $\text{poly}(\cdot)$ such that for all parameters $\lambda, n \in \mathbb{N}$, all (hk, vk, td) in the support of $\text{Setup}(1^\lambda, 1^n, \cdot)$, all inputs $x \in \{0, 1\}^n$ and all indices $i \in [n]$, the following properties hold:
 - **Succinct verification key:** $|vk| \leq \text{poly}(\lambda + \log n)$.

- **Succinct digest:** The digest dig output by $\text{Hash}(\text{hk}, x)$ satisfies $|\text{dig}| \leq \text{poly}(\lambda + \log n)$.
- **Succinct openings:** The opening σ output by $\text{ProveOpen}(\text{hk}, x, i)$ satisfies $|\sigma| \leq \text{poly}(\lambda + \log n)$.
- **Succinct verification:** The running time of $\text{VerOpen}(\text{vk}, \cdot, \cdot, \cdot, \cdot)$ is $\text{poly}(\lambda + \log n)$.
- **Correctness:** For all $\lambda, n \in \mathbb{N}$, every $x \in \{0, 1\}^n$, and every $i \in [n]$, the following properties hold:
 - **Opening correctness:**

$$\Pr \left[\text{VerOpen}(\text{vk}, \text{dig}, i, x_i, \sigma) = 1 \quad : \quad \begin{array}{l} (\text{hk}, \text{vk}, \text{td}) \leftarrow \text{Setup}(1^\lambda, 1^n, \emptyset) \\ \text{dig} \leftarrow \text{Hash}(\text{hk}, x), \sigma \leftarrow \text{ProveOpen}(\text{hk}, x, i) \end{array} \right] = 1.$$

- **Digest correctness:**

$$\Pr [\text{ValidateDigest}(\text{vk}, \text{dig}) = 1 : (\text{hk}, \text{vk}, \text{td}) \leftarrow \text{Setup}(1^\lambda, 1^n, \emptyset), \text{dig} \leftarrow \text{Hash}(\text{hk}, x)] = 1.$$

We additionally require the following security properties:

- **Set hiding:** For a bit $b \in \{0, 1\}$ and an adversary \mathcal{A} , we define the set hiding game $\text{ExptSH}_{\mathcal{A}}(\lambda, b)$ as follows:
 1. On input 1^λ , the adversary \mathcal{A} outputs 1^n and a set $S \subseteq [n]$.
 2. If $b = 0$, the challenger samples $(\text{hk}, \text{vk}, \text{td}) \leftarrow \text{Setup}(1^\lambda, 1^n, \emptyset)$ and if $b = 1$, the challenger samples $(\text{hk}, \text{vk}, \text{td}) \leftarrow \text{Setup}(1^\lambda, 1^n, S)$. It gives (hk, vk) to \mathcal{A} .
 3. Algorithm \mathcal{A} outputs a bit b' which is the output of the experiment.

The hash function satisfies set binding if for all efficient adversaries \mathcal{A} , there exists a negligible function $\text{negl}(\cdot)$ such that

$$|\Pr[\text{ExptSH}_{\mathcal{A}}(\lambda, 0) = 1] - \Pr[\text{ExptSH}_{\mathcal{A}}(\lambda, 1) = 1]| = \text{negl}(\lambda).$$

- **Index hiding with extracted guess:** For an adversary \mathcal{A} and a bit $b \in \{0, 1\}$, we define the index hiding with extracted guess game $\text{ExptIHE}_{\mathcal{A}}(\lambda, b)$ as follows:
 1. On input 1^λ , algorithm \mathcal{A} outputs 1^n , a set $S \subseteq [n]$, and an index $i^* \in S$.
 2. If $b = 0$, the challenger samples $(\text{hk}, \text{vk}, \text{td}) \leftarrow \text{Setup}(1^\lambda, 1^n, S \setminus \{i^*\})$. Otherwise, it samples $(\text{hk}, \text{vk}, \text{td}) \leftarrow \text{Setup}(1^\lambda, 1^n, S)$. The challenger sends (hk, vk) to \mathcal{A} .
 3. Algorithm \mathcal{A} outputs a digest dig and an opening σ .
 4. The output of the experiment is 1 if $\text{VerOpen}(\text{hk}, \text{dig}, i^*, 0, \sigma) = 1$ and $\text{Extract}(\text{td}, \text{dig})$ outputs Matching. Otherwise, the output is 0.

The hash function satisfies index hiding with extracted guess if for all efficient adversaries \mathcal{A} , there exists a negligible function $\text{negl}(\cdot)$ such that

$$|\Pr[\text{ExptIHE}_{\mathcal{A}}(\lambda, 0) = 1] - \Pr[\text{ExptIHE}_{\mathcal{A}}(\lambda, 1) = 1]| = \text{negl}(\lambda).$$

- **Zero fixing:** For an adversary \mathcal{A} , we define the adaptive zero-fixing game $\text{ExptZF}_{\mathcal{A}}(\lambda)$ as follows:
 1. On input 1^λ , algorithm \mathcal{A} outputs 1^n and a set $S \subseteq [n]$.
 2. The challenger samples $(\text{hk}, \text{vk}, \text{td}) \leftarrow \text{Setup}(1^\lambda, 1^n, S)$ and gives (hk, vk) to \mathcal{A} .
 3. Algorithm \mathcal{A} outputs a digest dig , an index $i \in S$ and an opening σ .
 4. The output of the experiment is 1 if $\text{VerOpen}(\text{hk}, \text{dig}, i, 1, \sigma) = 1$ and $\text{Extract}(\text{td}, \text{dig})$ outputs Matching. Otherwise, the output is 0.

The hash function satisfies zero-fixing if for all efficient adversaries \mathcal{A} , there exists a negligible function $\text{negl}(\cdot)$ such that $\Pr[\text{ExptZF}_{\mathcal{A}}(\lambda) = 1] = \text{negl}(\lambda)$.

- **Extractor validity:** For an adversary \mathcal{A} , we define the extractor validity game $\text{ExptEV}_{\mathcal{A}}(\lambda)$ as follows:
 1. On input 1^λ , the adversary \mathcal{A} outputs 1^n .
 2. The challenger samples $(\text{hk}, \text{vk}, \text{td}) \leftarrow \text{Setup}(1^\lambda, 1^n, \emptyset)$ and sends (hk, vk) to the adversary.
 3. Algorithm \mathcal{A} outputs a digest dig .
 4. The output of the experiment is 1 if $\text{ValidateDigest}(\text{hk}, \text{dig}) = 1$ and $\text{Extract}(\text{td}, \text{dig})$ outputs NotMatching . Otherwise, the output is 0.

The hash function satisfies the extractor validity property if for every efficient adversary \mathcal{A} , there exists a negligible function $\text{negl}(\cdot)$ such that $\Pr[\text{ExptEV}_{\mathcal{A}}(\lambda) = 1] = \text{negl}(\lambda)$.

Remark 3.2 (Selective Zero-Fixing Security). We can define a weaker selective notion of zero-fixing security where the adversary outputs the index $i \in S$ at the beginning of the security game (i.e., before seeing hk and vk). Note that the selective zero-fixing security definition is equivalent to the zero-fixing definition in [Definition 3.1](#). To see that selective zero fixing implies standard zero-fixing, consider a reduction algorithm that guesses the index $i \xleftarrow{\mathcal{R}} S$ at the beginning of the security reduction and aborts whenever the guess is incorrect. This reduction succeeds with probability $1/|S|$; since $|S| = \text{poly}(\lambda)$, this incurs only a polynomial loss in advantage. In our construction ([Construction 4.4](#)), we will work with the adaptive notion of security, but in our constructions ([Constructions 5.2](#) and [6.3](#)), we will work with the simpler selective definition.

One-sided index hiding. For our application, it suffices to consider a weaker notion of “one-sided” index hiding where we only require that the adversary’s advantage cannot increase (but could decrease). Proving one-sided security is often easier than proving two-sided security, so we define the simpler notion here:

Definition 3.3 (One-Sided Index-Hiding with Extracted Guess). We say a zero-fixing hash function Π_H satisfies *one-sided* index-hiding with extracted guess security if for all efficient adversaries \mathcal{A} , there exists a negligible function $\text{negl}(\cdot)$ such that

$$\Pr[\text{ExptIHE}_{\mathcal{A}}(\lambda, 1) = 1] \geq \Pr[\text{ExptIHE}_{\mathcal{A}}(\lambda, 0) = 1] - \text{negl}(\lambda).$$

4 Constructing Monotone Policy BARGs

In this section, we describe how to construct monotone policy BARGs from a standard batch argument for NP together with a zero-fixing hash function. We start by defining the conventions we use for describing Boolean circuits.

Definition 4.1 (Monotone Circuit Wire Indexing). Let $P: \{0, 1\}^k \rightarrow \{0, 1\}$ be a monotone Boolean circuit consisting exclusively of AND and OR gates with fan-in two. Let s be the size of P (i.e., the number of wires in P). A topological indexing of the wires of C is an assignment of an index $i \in [s]$ to each wire in P with the following properties:

- **Input wire:** For $i \in [k]$, the i^{th} input to P is associated with the index i .
- **Output wire:** The output wire is associated with the index s .
- **Intermediate wires:** The intermediate wires are associated with an index $i \in \{k + 1, \dots, s - 1\}$ with the property that the value of index i is completely determined by the values of the wires with indices $j_{i,1}, j_{i,2} \in \{1, \dots, i - 1\}$.

Every monotone circuit P has a canonical topological indexing that can be computed efficiently (e.g., by applying a deterministic topological sort to the wires of P).

Definition 4.2 (Layered Monotone Circuit). Let $P: \{0, 1\}^k \rightarrow \{0, 1\}$ be a (monotone) Boolean circuit of size s . We denote by $L_P(i)$ the *layer* of the wire i and define it as follows:

- If $i \in [k]$ (i.e., an input wire), then $L_P(i) = 1$.

- If $i > k$ then $L_P(i) = 1 + \max\{L_P(j_{i,1}), L_P(j_{i,2})\}$, where $j_{i,1}, j_{i,2} \in \{1, \dots, i-1\}$ are the indices of the input wires to the gate that computes the value of wire i .

The depth of the circuit is defined to be the layer associated with the output wire: $d = L_P(s)$. A circuit is *layered* if for every $i \in \{k+1, \dots, s\}$, it holds that $L_P(j_{i,1}) = L_P(j_{i,2})$. For a layer index $\ell \in [d]$, we define $\text{layer}_\ell(P) = \{i \in [s] : L_P(i) = \ell\}$ to be the set of wire indices in layer ℓ of the circuit.

Remark 4.3 (Layered Monotone Circuit). Every monotone circuit $P: \{0, 1\}^k \rightarrow \{0, 1\}$ of size s can be converted into a layered monotone circuit of size $\text{poly}(s)$. Thus, without loss of generality, we exclusively consider layered monotone circuits in the remainder of this work.

4.1 Monotone Policy BARG Construction

We now describe our construction of a monotone policy BARG for NP.

Construction 4.4 (Monotone Policy BARG). Let $\Pi'_{\text{BARG}} = (\text{Gen}', \text{Prove}', \text{Verify}', \text{TrapGen}', \text{Extract}')$ be a somewhere extractable BARG for Boolean circuit satisfiability. Let $\Pi_{\text{H}} = (\text{H.Setup}, \text{H.Hash}, \text{H.ProveOpen}, \text{H.VerOpen}, \text{H.Extract}, \text{H.ValidateDigest})$ be a zero-fixing hash function. We construct a monotone policy BARG $\Pi_{\text{MP-BARG}} = (\text{Gen}, \text{Prove}, \text{Verify})$ as follows:

- $\text{Gen}(1^\lambda, 1^n, 1^{s_c}, 1^{s_p})$: On input the security parameter λ , the input length n , the bound on the size of the Boolean circuit s_c , and the bound on the size of the monotone policy s_p , the setup algorithm proceeds as follows:

- Sample two hash keys

$$(\text{hk}_0, \text{vk}_0, \text{td}_0) \leftarrow \text{H.Setup}(1^\lambda, 1^{s_p}, \emptyset)$$

$$(\text{hk}_1, \text{vk}_1, \text{td}_1) \leftarrow \text{H.Setup}(1^\lambda, 1^{s_p}, \emptyset).$$

- Let s' be a bound on the size of the circuit that computes the relation $\mathcal{R}[C, k, s_p, \text{vk}_0, \text{vk}_1, \text{dig}_0, \text{dig}_1]$ from Fig. 1 when instantiated with an arbitrary Boolean circuit C of size at most s_c , an input length $k \leq s_p$ and digests $\text{dig}_0, \text{dig}_1$ associated with the hash and verification keys $(\text{hk}_0, \text{vk}_0)$ and $(\text{hk}_1, \text{vk}_1)$. Let $n' = 3 \cdot \lceil \log s_p \rceil + 1$ be the bound on the statement length. Sample $(\text{crs}_{\text{BARG}}, \text{vk}_{\text{BARG}}) \leftarrow \text{Gen}'(1^\lambda, 1^{s_p}, 1^{n'}, 1^{s'})$.⁷

It outputs the common reference string $\text{crs} = (\text{crs}_{\text{BARG}}, \text{vk}_{\text{BARG}}, \text{hk}_0, \text{hk}_1, \text{vk}_0, \text{vk}_1)$.

- $\text{Prove}(\text{crs}, C, P, (x_1, \dots, x_k), (w_1, \dots, w_k))$: On input a CRS $\text{crs} = (\text{crs}_{\text{BARG}}, \text{vk}_{\text{BARG}}, \text{hk}_0, \text{hk}_1, \text{vk}_0, \text{vk}_1)$, a circuit $C: \{0, 1\}^n \times \{0, 1\}^h \rightarrow \{0, 1\}$, a monotone layered Boolean policy circuit $P: \{0, 1\}^k \rightarrow \{0, 1\}$, statements $x_1, \dots, x_k \in \{0, 1\}^n$, and witnesses $w_1, \dots, w_k \in \{0, 1\}^h$, the prove algorithm does the following:

- Let s be the size of P . Index the wires of P under a canonical topological ordering (Definition 4.1). For each wire $i \in \{k+1, \dots, s\}$, let $g_i \in \{\text{AND}, \text{OR}\}$ be its type. Let $j_{i,1}, j_{i,2} \in \{1, \dots, i-1\}$ be the indices of the input wires to the gate i .
- For each $i \in [s]$, let $\beta_i \in \{0, 1\}$ be the value of the i^{th} wire in the evaluation of P on the input bits $(C(x_1, w_1), \dots, C(x_k, w_k))$. For $i \in \{s+1, \dots, s_p\}$, let $\beta_i = 0$. (This corresponds to “padding” the $s_p - s$ unused slots).
- Compute the digest $\text{dig}_0 \leftarrow \text{H.Hash}(\text{hk}_0, (\beta_1, \dots, \beta_{s_p}))$ and $\text{dig}_1 \leftarrow \text{H.Hash}(\text{hk}_1, (\beta_1, \dots, \beta_{s_p}))$.
- For each $i \in [s]$ and each $b \in \{0, 1\}$, compute $\sigma_i^{(b)} \leftarrow \text{H.ProveOpen}(\text{hk}_b, (\beta_1, \dots, \beta_{s_p}), i)$.
- Let C_{aug} be the circuit that computes the relation $\mathcal{R}[C, k, s, \text{vk}_0, \text{vk}_1, \text{dig}_0, \text{dig}_1]$ shown in Fig. 1.
- For each $i \in [s_p]$, construct the statement \hat{x}_i and witness \hat{w}_i as follows:
 - * If $i \in [k]$, let $\hat{x}_i = (i, x_i)$ and $\hat{w}_i = (\beta_i, \sigma_i^{(0)}, \sigma_i^{(1)}, w_i)$.

⁷Recall that when the bound on the extraction set parameter ℓ is not given, it defaults to the value 1.

Statement: index i and auxiliary statement x

Witness: value b , openings $(\sigma^{(0)}, \sigma^{(1)})$ and auxiliary witness w

Hard-Coded: circuit $C: \{0, 1\}^n \times \{0, 1\}^h \rightarrow \{0, 1\}$, the number of inputs k , the policy size s , the hash keys hk_0, hk_1 , and the digests dig_0, dig_1

On input a statement (i, x) and a witness $(b, \sigma^{(0)}, \sigma^{(1)}, w)$:

- If $i \leq k$, output 1 if all of the following conditions are met, otherwise output 0:
 - * **Opening validity:** For all $\alpha \in \{0, 1\}$, $H.VerOpen(vk_\alpha, dig_\alpha, i, b, \sigma^{(\alpha)}) = 1$.
 - * **Wire consistency:** $C(x, w) = b$.
- If $i \in \{k + 1, \dots, s\}$, parse $x = (g, j_1, j_2)$ where $g \in \{AND, OR\}$ and $j_1, j_2 \in \{1, \dots, i - 1\}$. Parse $w = (b_1, \sigma_1^{(0)}, \sigma_1^{(1)}, b_2, \sigma_2^{(0)}, \sigma_2^{(1)})$. Check each of the following conditions for $\alpha \in \{0, 1\}$: Output 1 if all of the following conditions are met, otherwise output 0:
 - * **Opening validity:** For all $\alpha \in \{0, 1\}$, all of the following holds:
 - $H.VerOpen(vk_\alpha, dig_\alpha, j_1, b_1, \sigma_1^{(\alpha)}) = 1$;
 - $H.VerOpen(vk_\alpha, dig_\alpha, j_2, b_2, \sigma_2^{(\alpha)}) = 1$;
 - $H.VerOpen(vk_\alpha, dig_\alpha, i, b, \sigma^{(\alpha)}) = 1$.
 - * **Wire consistency:** $b = g(b_1, b_2)$.
 - * **Output gate:** If $i = s$, check that $b = 1$.
- If $i > s$, then output 1.

Figure 1: The relation $\mathcal{R}[C, k, s, vk_0, vk_1, dig_0, dig_1]$.

- * If $i \in [k + 1, s]$, let $\hat{x}_i = (i, (g_i, j_{i,1}, j_{i,2}))$ and

$$\hat{w}_i = (\beta_i, \sigma_i^{(0)}, \sigma_i^{(1)}, (\beta_{j_{i,1}}, \sigma_{j_{i,1}}^{(0)}, \sigma_{j_{i,1}}^{(1)}, \beta_{j_{i,2}}, \sigma_{j_{i,2}}^{(0)}, \sigma_{j_{i,2}}^{(1)})).$$

- * If $i > s$, let $\hat{x}_i = \perp$ and $\hat{w}_i = \perp$.

Essentially, there is an instance \hat{x}_i associated with each wire i of P .

- Compute the BARG proof $\pi_{BARG} \leftarrow \text{Prove}'(crs_{BARG}, C_{aug}, (\hat{x}_1, \dots, \hat{x}_{s_p}), (\hat{w}_1, \dots, \hat{w}_{s_p}))$ and output $\pi = (dig_0, dig_1, \pi_{BARG})$.
- **Verify**($crs, C, P, (x_1, \dots, x_k), \pi$): On input a common reference string $crs = (crs_{BARG}, vk_{BARG}, hk_0, hk_1, vk_0, vk_1)$, a Boolean circuit $C: \{0, 1\}^n \times \{0, 1\}^h \rightarrow \{0, 1\}$, a layered monotone Boolean policy $P: \{0, 1\}^k \rightarrow \{0, 1\}$, statements $x_1, \dots, x_k \in \{0, 1\}^n$, and a proof $\pi = (dig_0, dig_1, \pi_{BARG})$, the verification algorithm does the following:
 - If $H.ValidateDigest(vk_0, dig_0) = 0$ or $H.ValidateDigest(vk_1, dig_1) = 0$, then output 0.
 - Let s be the size of P . Index the wires of P under a canonical topological ordering (Definition 4.1). For each wire $i \in \{k + 1, \dots, s\}$, let $j_{i,1}, j_{i,2} \in \{1, \dots, i - 1\}$ be the indices of the input wires of the gate $g_i \in \{AND, OR\}$ that computes wire i . For each $i \in [s_p]$, construct the statement \hat{x}_i as follows:
 - * If $i \in [k]$, let $\hat{x}_i = (i, x_i)$.
 - * If $i \in \{k + 1, \dots, s\}$, let $\hat{x}_i = (i, (g_i, j_{i,1}, j_{i,2}))$.
 - * If $i > s$, let $\hat{x}_i = \perp$.

- Let C_{aug} be the circuit that computes the relation $\mathcal{R}[C, k, s, \text{vk}_0, \text{vk}_1, \text{dig}_0, \text{dig}_1]$ from Fig. 1.
- Output $\text{Verify}'(\text{vk}_{\text{BARG}}, C_{\text{aug}}, (\hat{x}_1, \dots, \hat{x}_{s_p}), \pi_{\text{BARG}})$.

Theorem 4.5 (Completeness). *If Π'_{BARG} is complete and Π_{H} is correct, then Construction 4.4 is complete.*

Proof. Take any $\lambda, n, s_c, s_p \in \mathbb{N}$, any Boolean circuit $C: \{0, 1\}^n \times \{0, 1\}^h \rightarrow \{0, 1\}$ of size at most s_c , and any monotone Boolean policy $P: \{0, 1\}^k \rightarrow \{0, 1\}$ of size $s \leq s_p$. Let $x_1, \dots, x_k \in \{0, 1\}^n$ and $w_1, \dots, w_k \in \{0, 1\}^h$ be a collection of statements and witnesses such that $P(C(x_1, w_1), \dots, C(x_t, w_t)) = 1$. We start by defining the following quantities:

- For each $i \in [s_p]$, let $\beta_i \in \{0, 1\}$ be the value of wire i for predicate P on input $(C(x_1, w_1), \dots, C(x_t, w_t))$.
- Let $\text{crs} = (\text{crs}_{\text{BARG}}, \text{vk}_{\text{BARG}}, \text{hk}_0, \text{hk}_1, \text{vk}_0, \text{vk}_1) \leftarrow \text{Gen}(1^\lambda, 1^n, 1^{s_c}, 1^{s_p})$. By construction, the hash keys are sampled as $(\text{hk}_b, \text{vk}_b, \text{td}_b) \leftarrow \text{H.Setup}(1^\lambda, 1^{s_p}, \emptyset)$ for each $b \in \{0, 1\}$.
- Let $\pi \leftarrow \text{Prove}(\text{crs}, C, P, (x_1, \dots, x_k), (w_1, \dots, w_k))$. Let

$$\text{dig}_0 = \text{H.Hash}(\text{hk}_0, (\beta_1, \dots, \beta_{s_p})) \quad \text{and} \quad \text{dig}_1 = \text{H.Hash}(\text{hk}_1, (\beta_1, \dots, \beta_{s_p}))$$

be the digests computed by the Prove algorithm. Moreover, by correctness of Π_{H} , for each $b \in \{0, 1\}$, we have $\text{H.ValidateDigest}(\text{vk}_b, \text{dig}_b) = 1$.

Consider now the relation $\mathcal{R}[C, k, s, \text{vk}_0, \text{vk}_1, \text{dig}_0, \text{dig}_1]$ defined in Fig. 1. We show that for all $i \in [s_p]$, the statement (\hat{x}_i, \hat{w}_i) defined in Prove satisfies the relation. First, for all $i \in [s_p]$ and $b \in \{0, 1\}$, the Prove algorithm computes $\sigma_i^{(b)} \leftarrow \text{H.ProveOpen}(\text{hk}_b, (\beta_1, \dots, \beta_{s_p}), i)$. Correspondingly, by correctness of Π_{H} , we conclude that for all $i \in [s_p]$ and $b \in \{0, 1\}$,

$$\text{H.VerOpen}(\text{vk}_b, \text{dig}_b, i, \beta_i, \sigma_i^{(b)}) = 1. \quad (4.1)$$

We now consider each index $i \in [s_p]$:

- If $i \in [k]$, then $\hat{x}_i = (i, x_i)$ and $\hat{w}_i = (\beta_i, \sigma_i^{(0)}, \sigma_i^{(1)}, w_i)$. In this case, the opening validity passes by Eq. (4.1). Moreover, by definition, we have that $\beta_i = C(x_i, w_i)$. Hence, the relation is satisfied.
- If $i \in \{k+1, \dots, s\}$, let $\hat{x}_i = (i, (g_i, j_{i,1}, j_{i,2}))$ and $\hat{w}_i = (\beta_i, \sigma_i^{(0)}, \sigma_i^{(1)}, (\beta_{j_{i,1}}, \sigma_{j_{i,1}}^{(0)}, \sigma_{j_{i,1}}^{(1)}, \beta_{j_{i,2}}, \sigma_{j_{i,2}}^{(0)}, \sigma_{j_{i,2}}^{(1)}))$. Again, the opening validity check passes by Eq. (4.1). Moreover, by definition, $\beta_i = g_i(\beta_{j_{i,1}}, \beta_{j_{i,2}})$ so the wire consistency-check passes. Finally, if $i = s$, then $\beta_s = P(C(x_1, w_1), \dots, C(x_t, w_t)) = 1$ by construction.
- Finally, if $i > s$, the relation is always satisfied.

Thus, we conclude that for all $i \in [s_p]$, the relation \mathcal{R} is always satisfied. By completeness of Π'_{BARG} , this means $\text{Verify}'(\text{vk}_{\text{BARG}}, C_{\text{aug}}, (\hat{x}_1, \dots, \hat{x}_{s_p}), \pi_{\text{BARG}}) = 1$, where $\pi_{\text{BARG}} \leftarrow \text{Prove}'(\text{crs}_{\text{BARG}}, C_{\text{aug}}, (\hat{x}_1, \dots, \hat{x}_{s_p}), (\hat{w}_1, \dots, \hat{w}_{s_p}))$. Letting $\pi = (\text{dig}_0, \text{dig}_1, \pi_{\text{BARG}})$, we conclude that $\text{Verify}(\text{crs}, C, P, (x_1, \dots, x_k), \pi) = 1$, and completeness follows. \square

Theorem 4.6 (Succinctness). *If Π'_{BARG} and Π_{H} satisfy succinctness, then Construction 4.4 has slightly succinct proofs.*

Proof. Fix $\lambda, n, s_c, s_p \in \mathbb{N}$, a Boolean circuit $C: \{0, 1\}^n \times \{0, 1\}^h \rightarrow \{0, 1\}$ of size at most s_c , and a monotone Boolean policy $P: \{0, 1\}^k \rightarrow \{0, 1\}$ of size $s \leq s_p$. Take any $\text{crs} = (\text{crs}_{\text{BARG}}, \text{vk}_{\text{BARG}}, \text{hk}_0, \text{hk}_1, \text{vk}_0, \text{vk}_1)$ in the support of $\text{Gen}(1^\lambda, 1^n, 1^{s_c}, 1^{s_p})$. Consider any collection of statements $x_1, \dots, x_k \in \{0, 1\}^n$ and witnesses $w_1, \dots, w_k \in \{0, 1\}^h$ where $P(C(x_1, w_1), \dots, C(x_t, w_t)) = 1$. We bound the size of $\pi_{\text{BARG}} \leftarrow \text{Prove}(\text{crs}, C, P, (x_1, \dots, x_k), (w_1, \dots, w_k))$:

- First, the Prove algorithm computes dig_0 and dig_1 using the hash keys hk_0 and hk_1 , respectively. Since $(\text{hk}, \text{vk}, \text{td}) \leftarrow \text{H.Setup}(1^\lambda, 1^{s_p}, \emptyset)$, we appeal to succinctness of Π_{H} to conclude that $|\text{dig}_0|, |\text{dig}_1| \leq \text{poly}(\lambda + \log s_p)$.
- Next, consider the size of the BARG proof π_{BARG} . We first bound the size of the circuit C_{aug} for computing the relation $\mathcal{R}[C, k, s, \text{vk}_0, \text{vk}_1, \text{dig}_0, \text{dig}_1]$ from Fig. 1. By construction, C_{aug} performs a constant number of calls to H.VerOpen and also needs to evaluate the underlying circuit C (which has size at most s_c). By succinctness of Π_{H} , each invocation of H.VerOpen can be computed by a circuit of size $\text{poly}(\lambda + \log s_p)$. Hence, the size of the circuit C_{aug} can be bounded by $\text{poly}(\lambda + \log s_p + s_c)$. By succinctness of Π_{BARG} , we conclude that the size of the proof π_{BARG} output by Prove' is bounded by $\text{poly}(\lambda + |C_{\text{aug}}| + \log s_p) \leq \text{poly}(\lambda + s_c + \log s_p)$.

Putting the pieces together, the proof $\pi = (\text{dig}_0, \text{dig}_1, \pi_{\text{BARG}})$ output by Prove is bounded by $\text{poly}(\lambda + s_c + \log s_p)$, and (slight) succinctness follows. \square

Soundness. We now state the soundness theorem, but give the proof in the subsequent section (Section 4.2).

Theorem 4.7 (Non-Adaptive Soundness). *If Π_H satisfies set hiding, index hiding with extracted guess, zero fixing and extractor validity against non-uniform adversaries, and Π_{BARG} is somewhere extractable and satisfies set hiding against non-uniform adversaries, then Construction 4.4 satisfies non-adaptive soundness against non-uniform adversaries.*

4.2 Proof of Theorem 4.7 (Non-Adaptive Soundness)

In this section, we prove non-adaptive soundness of Construction 4.4. Take any efficient non-uniform adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ for the non-adaptive soundness game. Then let $(1^n, 1^{s_c}, 1^{s_p}, C, P, \mathbf{x}, \text{st}_{\mathcal{A}}) \leftarrow \mathcal{A}_1(1^\lambda)$ for any $\lambda \in \mathbb{N}$, where

- $C: \{0, 1\}^n \times \{0, 1\}^h \rightarrow \{0, 1\}$ is a Boolean circuit of size at most s_c ;
- $P: \{0, 1\}^k \rightarrow \{0, 1\}$ is a layered monotone Boolean circuit of size $s \leq s_p$; and
- $\mathbf{x} = (x_1, \dots, x_k) \in \{0, 1\}^{kn}$ where $x_i \in \{0, 1\}^n$ for all $i \in [k]$.

Let d be the depth of P . For each $i \in [k]$, let $\beta_i \in \{0, 1\}$ be equal to 1 if $(C, x_i) \in \mathcal{L}_{\text{CSAT}}$ and 0 otherwise. Extending the definition to all $i \in [s]$, let $\beta_i \in \{0, 1\}$ be the value of wire i in the evaluation of P on input $(\beta_1, \dots, \beta_k)$. For each layer $i \in [d]$, define

$$J_i = \{j \in \text{layer}_i(P) : \beta_j = 0\} \quad (4.2)$$

to be the indices of the zero wires in layer i of $P(\beta_1, \dots, \beta_k)$. We model each J_i as an *ordered* set (ordered in ascending order). We write $J_i[t]$ to denote the t^{th} element in J_i and $J_i[1, \dots, t] = \bigcup_{j \in [t]} J_i[j]$ to denote the first t elements of J_i .

4.2.1 Hybrid Experiment Specification

To prove Theorem 4.7, we start by defining a sequence of hybrid experiments. Each of these hybrids is indexed implicitly by the security parameter λ , but we omit this for ease of exposition.

Outer games. We start by defining a sequence of “outer hybrids.” Here, we provide a general overview of our methodology. The initial hybrid Hyb_0 corresponds to the real non-adaptive soundness game, while Hyb_i corresponds to the hybrid where one of the zero-fixing hash keys is binding on the set J_i (as defined by Eq. (4.2)). We show that the outputs of each adjacent pair of hybrid distributions can only change by a negligible amount, and moreover, that the zero-fixing hash function binding on J_i in Hyb_i outputs Matching. Finally, in hybrid Hyb_d , the following two conditions hold:

- The hash key is zero-fixing on the single output wire (since we know that $P(\beta_1, \dots, \beta_k) = 0$, where β_i is the indicator bit for whether $(C, x_i) \in \mathcal{L}_{\text{CSAT}}$).
- The hash function declares the output bit to be Matching.

Consider the probability that the proof verifies in Hyb_d :

- Suppose the BARG is extractable on the instance associated with the output wire of P . In this case, if the proof verifies in Hyb_d , then somewhere extractability of the BARG allows us to extract an opening to 1 with respect to both zero-fixing hash functions. This follows by definition of the instance \hat{x}_s in Prove and Verify (where s is the size of P).
- Since $P(\beta_1, \dots, \beta_k) = 0$, one of the zero-fixing hash functions will be zero-fixing on the output wire in hybrid Hyb_d . Moreover, this hash function outputs Matching. If we can extract an opening to 1 for the output wire, this breaks zero-fixing security of the hash function.

Thus, when the BARG is extractable on the instance associated with the output wire, the probability that the proof verifies in Hyb_d is negligible. Finally, if the outputs of each adjacent pair of hybrids cannot differ by a non-negligible amount, we conclude the probability that the proof verifies in Hyb_0 is also negligible. This demonstrates non-adaptive soundness. We now define the sequence of games:

- Hyb_0 : This is the non-adaptive soundness game. For ease of exposition, we partition the game into two phases:
 - **Phase 1:** On input the security parameter 1^λ , algorithm \mathcal{A}_1 outputs $1^n, 1^{s_c}, 1^{s_p}$, a Boolean circuit $C: \{0, 1\}^n \times \{0, 1\}^h \rightarrow \{0, 1\}$ of size at most s_c , a monotone Boolean circuit $P: \{0, 1\}^k \rightarrow \{0, 1\}$ of size $s \leq s_p$, an instance $\mathbf{x} = (x_1, \dots, x_k) \in \{0, 1\}^{kn}$, and the state $\text{st}_{\mathcal{A}}$. If $(C, P, \mathbf{x}) \in \mathcal{L}_{\text{MP-CSAT}}$, then the experiment outputs 0.
 - **Phase 2:** The challenger computes $\text{crs} \leftarrow \text{Gen}(1^\lambda, 1^n, 1^{s_c}, 1^{s_p})$. Specifically, the challenger samples the following components:
 - * $(\text{crs}_{\text{BARG}}, \text{vk}_{\text{BARG}}) \leftarrow \text{Gen}'(1^\lambda, 1^{s_p}, 1^{n'}, 1^{s'})$.
 - * $(\text{hk}_0, \text{vk}_0, \text{td}_0) \leftarrow \text{H.Setup}(1^\lambda, 1^{s_p}, \emptyset)$.
 - * $(\text{hk}_1, \text{vk}_1, \text{td}_1) \leftarrow \text{H.Setup}(1^\lambda, 1^{s_p}, \emptyset)$.

The challenger sets $\text{crs} = (\text{crs}_{\text{BARG}}, \text{hk}_0, \text{hk}_1, \text{vk}_0, \text{vk}_1)$ and runs \mathcal{A}_2 on input $(\text{crs}, \text{st}_{\mathcal{A}})$. Algorithm \mathcal{A}_2 outputs a proof string $\pi = (\text{dig}_0, \text{dig}_1, \pi_{\text{BARG}})$. Let $\hat{\mathbf{x}} = (\hat{x}_1, \dots, \hat{x}_{s_p})$ and C_{aug} be as defined in Prove and Verify in [Construction 4.4](#). The output of the experiment is 1 if all of the following conditions hold (and 0 otherwise):

- * $\text{H.ValidateDigest}(\text{vk}_0, \text{dig}_0) = \text{H.ValidateDigest}(\text{vk}_1, \text{dig}_1) = 1$.
- * $\text{Verify}'(\text{vk}_{\text{BARG}}, C_{\text{aug}}, \hat{\mathbf{x}}, \pi_{\text{BARG}}) = 1$.

- Hyb_i for $i \in [d]$: Same as Hyb_0 , but hk_{low} binds on J_i , where $\text{low} = i \bmod 2$ and $\text{high} = 1 - \text{low}$. Specifically, the game proceeds as follows:

- **Phase 1:** On input the security parameter 1^λ , algorithm \mathcal{A}_1 outputs $1^n, 1^{s_c}, 1^{s_p}$, a Boolean circuit $C: \{0, 1\}^n \times \{0, 1\}^h \rightarrow \{0, 1\}$ of size at most s_c , a monotone Boolean circuit $P: \{0, 1\}^k \rightarrow \{0, 1\}$ of size $s \leq s_p$, an instance $\mathbf{x} = (x_1, \dots, x_k) \in \{0, 1\}^{kn}$, and the state $\text{st}_{\mathcal{A}}$. If $(C, P, \mathbf{x}) \in \mathcal{L}_{\text{MP-CSAT}}$, then the experiment outputs 0. In addition, the challenger computes the following quantities:
 - * For $j \in [k]$, let $\beta_j = 1$ if $(C, x_j) \in \mathcal{L}_{\text{CSAT}}$ ([Definition 2.5](#)) and $\beta_j = 0$ otherwise.
 - * For $j \in [k+1, s]$, let β_j be the value of the wire j in the evaluation of P on $(\beta_1, \dots, \beta_k)$.
 - * For each layer $\ell \in [d]$, let $J_\ell = \{j \in \text{layer}_\ell(P) : \beta_j = 0\}$.
- **Phase 2:** The challenger samples the following components:
 - * $(\text{crs}_{\text{BARG}}, \text{vk}_{\text{BARG}}) \leftarrow \text{Gen}'(1^\lambda, 1^{s_p}, 1^{n'}, 1^{s'})$.
 - * $(\text{hk}_{\text{low}}, \text{vk}_{\text{low}}, \text{td}_{\text{low}}) \leftarrow \text{H.Setup}(1^\lambda, 1^{s_p}, J_i)$.
 - * $(\text{hk}_{\text{high}}, \text{vk}_{\text{high}}, \text{td}_{\text{high}}) \leftarrow \text{H.Setup}(1^\lambda, 1^{s_p}, \emptyset)$.

The challenger sets $\text{crs} = (\text{crs}_{\text{BARG}}, \text{hk}_0, \text{hk}_1, \text{vk}_0, \text{vk}_1)$ and runs \mathcal{A}_2 on input $(\text{crs}, \text{st}_{\mathcal{A}})$. Algorithm \mathcal{A}_2 outputs a proof string $\pi = (\text{dig}_0, \text{dig}_1, \pi_{\text{BARG}})$. Let $\hat{\mathbf{x}} = (\hat{x}_1, \dots, \hat{x}_{s_p})$ and C_{aug} be as defined in Prove and Verify in [Construction 4.4](#). The output of the experiment is 1 if all of the following conditions hold (and 0 otherwise):

- * $\text{H.ValidateDigest}(\text{vk}_0, \text{dig}_0) = \text{H.ValidateDigest}(\text{vk}_1, \text{dig}_1) = 1$.
- * $\text{Verify}'(\text{vk}_{\text{BARG}}, C_{\text{aug}}, \hat{\mathbf{x}}, \pi_{\text{BARG}}) = 1$.
- * $\text{H.Extract}(\text{td}_{\text{low}}, \text{dig}_{\text{low}}) = \text{Matching}$.

Inner games. To argue that each pair of outer hybrids is computationally indistinguishable, we now define a sequence of “inner hybrids.” Whereas the outer hybrids advance layer by layer, the inner hybrids advance *across* a layer. In more detail, recall that the difference between Hyb_i and Hyb_{i+1} is that *one* of the two hash keys (i.e., hk_{low}) goes from binding on J_i to binding on J_{i+1} . The idea in the inner hybrids is to program the *other* hash key (i.e., hk_{high}) to be binding on J_{i+1} . Initially, hk_{high} is binding on the empty set. We then step through $|J_{i+1}|$ intermediate hybrids, where on the t^{th} step, the hash key hk_{high} goes from being binding on $J_{i+1}[1, \dots, t-1]$ to being binding on $J_{i+1}[1, \dots, t]$. Each transition relies on the security of the BARG and the zero-fixing hash function. We now define the full sequence of hybrids; each one is indexed by $i \in \{0, \dots, d\}$.

- $\text{Hyb}_{i,t,1}$ for $t \in [|J_{i+1}|]$: Same as Hyb_i , but hk_{high} binds on the first $t - 1$ wires in J_{i+1} .

- **Phase 1:** Same as Hyb_i .

- **Phase 2:** The challenger samples the following components:

- * $(\text{crs}_{\text{BARG}}, \text{vk}_{\text{BARG}}) \leftarrow \text{Gen}'(1^\lambda, 1^{s_p}, 1^{n'}, 1^{s'})$.
- * If $i = 0$, then $(\text{hk}_{\text{low}}, \text{vk}_{\text{low}}, \text{td}_{\text{low}}) \leftarrow \text{H.Setup}(1^\lambda, 1^{s_p}, \emptyset)$. Otherwise, it samples $(\text{hk}_{\text{low}}, \text{vk}_{\text{low}}, \text{td}_{\text{low}}) \leftarrow \text{H.Setup}(1^\lambda, 1^{s_p}, J_i)$.
- * $(\text{hk}_{\text{high}}, \text{vk}_{\text{high}}, \text{td}_{\text{high}}) \leftarrow \text{H.Setup}(1^\lambda, 1^{s_p}, J_{i+1}[1, \dots, t - 1])$.

The challenger sets $\text{crs} = (\text{crs}_{\text{BARG}}, \text{hk}_0, \text{hk}_1, \text{vk}_0, \text{vk}_1)$ and runs \mathcal{A}_2 on input $(\text{crs}, \text{st}_{\mathcal{A}})$. Algorithm \mathcal{A}_2 outputs a proof string $\pi = (\text{dig}_0, \text{dig}_1, \pi_{\text{BARG}})$. Let $\hat{\mathbf{x}} = (\hat{x}_1, \dots, \hat{x}_{s_p})$ and C_{aug} be as defined in Prove and Verify in [Construction 4.4](#). The output of the experiment is 1 if all of the following conditions hold (and 0 otherwise):

- * $\text{H.ValidateDigest}(\text{vk}_0, \text{dig}_0) = \text{H.ValidateDigest}(\text{vk}_1, \text{dig}_1) = 1$.
- * $\text{Verify}'(\text{vk}_{\text{BARG}}, C_{\text{aug}}, \hat{\mathbf{x}}, \pi_{\text{BARG}}) = 1$.
- * If $i \geq 1$, then $\text{H.Extract}(\text{td}_{\text{low}}, \text{dig}_{\text{low}}) = \text{Matching}$.
- * $\text{H.Extract}(\text{td}_{\text{high}}, \text{dig}_{\text{high}}) = \text{Matching}$.

- $\text{Hyb}_{i,t,2}$ for $t \in [|J_{i+1}|]$: Same as $\text{Hyb}_{i,t,1}$, but crs_{BARG} is set to be extractable on index $J_{i+1}[t]$.

- **Phase 1:** Same as Hyb_i .

- **Phase 2:** The challenger samples the following components:

- * $(\text{crs}_{\text{BARG}}, \text{td}_{\text{BARG}}) \leftarrow \text{TrapGen}'(1^\lambda, 1^{s_p}, 1^{n'}, 1^{s'}, J_{i+1}[t])$.
- * If $i = 0$, then $(\text{hk}_{\text{low}}, \text{vk}_{\text{low}}, \text{td}_{\text{low}}) \leftarrow \text{H.Setup}(1^\lambda, 1^{s_p}, \emptyset)$. Otherwise, it samples $(\text{hk}_{\text{low}}, \text{vk}_{\text{low}}, \text{td}_{\text{low}}) \leftarrow \text{H.Setup}(1^\lambda, 1^{s_p}, J_i)$.
- * $(\text{hk}_{\text{high}}, \text{vk}_{\text{high}}, \text{td}_{\text{high}}) \leftarrow \text{H.Setup}(1^\lambda, 1^{s_p}, J_{i+1}[1, \dots, t - 1])$.

The challenger sets $\text{crs} = (\text{crs}_{\text{BARG}}, \text{hk}_0, \text{hk}_1, \text{vk}_0, \text{vk}_1)$ and runs \mathcal{A}_2 on input $(\text{crs}, \text{st}_{\mathcal{A}})$. Algorithm \mathcal{A}_2 outputs a proof string $\pi = (\text{dig}_0, \text{dig}_1, \pi_{\text{BARG}})$. Let $\hat{\mathbf{x}} = (\hat{x}_1, \dots, \hat{x}_{s_p})$ and C_{aug} be as defined in Prove and Verify in [Construction 4.4](#). The output of the experiment is 1 if all of the following conditions hold (and 0 otherwise):

- * $\text{H.ValidateDigest}(\text{vk}_0, \text{dig}_0) = \text{H.ValidateDigest}(\text{vk}_1, \text{dig}_1) = 1$.
- * $\text{Verify}'(\text{vk}_{\text{BARG}}, C_{\text{aug}}, \hat{\mathbf{x}}, \pi_{\text{BARG}}) = 1$.
- * If $i \geq 1$, then $\text{H.Extract}(\text{td}_{\text{low}}, \text{dig}_{\text{low}}) = \text{Matching}$.
- * $\text{H.Extract}(\text{td}_{\text{high}}, \text{dig}_{\text{high}}) = \text{Matching}$.

- $\text{Hyb}_{i,t,3}$ for $t \in [|J_{i+1}|]$: Same as $\text{Hyb}_{i,t,2}$, but the challenger additionally checks that it extracts a valid witness for $\hat{x}_{J_{i+1}[t]}$.

- **Phase 1:** Same as Hyb_i .

- **Phase 2:** The challenger samples the following components:

- * $(\text{crs}_{\text{BARG}}, \text{td}_{\text{BARG}}) \leftarrow \text{TrapGen}'(1^\lambda, 1^{s_p}, 1^{n'}, 1^{s'}, J_{i+1}[t])$.
- * If $i = 0$, then $(\text{hk}_{\text{low}}, \text{vk}_{\text{low}}, \text{td}_{\text{low}}) \leftarrow \text{H.Setup}(1^\lambda, 1^{s_p}, \emptyset)$. Otherwise, it samples $(\text{hk}_{\text{low}}, \text{vk}_{\text{low}}, \text{td}_{\text{low}}) \leftarrow \text{H.Setup}(1^\lambda, 1^{s_p}, J_i)$.
- * $(\text{hk}_{\text{high}}, \text{vk}_{\text{high}}, \text{td}_{\text{high}}) \leftarrow \text{H.Setup}(1^\lambda, 1^{s_p}, J_{i+1}[1, \dots, t - 1])$.

The challenger sets $\text{crs} = (\text{crs}_{\text{BARG}}, \text{hk}_0, \text{hk}_1, \text{vk}_0, \text{vk}_1)$ and runs \mathcal{A}_2 on input $(\text{crs}, \text{st}_{\mathcal{A}})$. Algorithm \mathcal{A}_2 outputs a proof string $\pi = (\text{dig}_0, \text{dig}_1, \pi_{\text{BARG}})$. Let $\hat{\mathbf{x}} = (\hat{x}_1, \dots, \hat{x}_{s_p})$ and C_{aug} be as defined in Prove and Verify in [Construction 4.4](#). The output of the experiment is 1 if all of the following conditions hold (and 0 otherwise):

- * $\text{H.ValidateDigest}(\text{vk}_0, \text{dig}_0) = \text{H.ValidateDigest}(\text{vk}_1, \text{dig}_1) = 1$.
- * $\text{Verify}'(\text{vk}_{\text{BARG}}, C_{\text{aug}}, \hat{\mathbf{x}}, \pi_{\text{BARG}}) = 1$.
- * If $i \geq 1$, then $\text{H.Extract}(\text{td}_{\text{low}}, \text{dig}_{\text{low}}) = \text{Matching}$.

- * $H.Extract(td_{high}, dig_{high}) = Matching$.
 - * $C_{aug}(\hat{x}_{J_{i+1}[t]}, \hat{w}) = 1$ where $\hat{w} = (b, \sigma^{(0)}, \sigma^{(1)}, w) \leftarrow Extract'(td_{BARG}, C_{aug}, \hat{x}, \pi_{BARG}, J_{i+1}[t])$.
- $Hyb_{i,t,4}$ for $t \in [|J_{i+1}|]$: Same as $Hyb_{i,t,3}$, but the challenger additionally checks that the extracted value for wire $J_{i+1}[t]$ is a 0.
 - **Phase 1:** Same as Hyb_i .
 - **Phase 2:** The challenger samples the following components:
 - * $(crs_{BARG}, td_{BARG}) \leftarrow TrapGen'(1^\lambda, 1^{sp}, 1^{n'}, 1^{s'}, J_{i+1}[t])$.
 - * If $i = 0$, then $(hk_{low}, vk_{low}, td_{low}) \leftarrow H.Setup(1^\lambda, 1^{sp}, \emptyset)$. Otherwise, it samples $(hk_{low}, vk_{low}, td_{low}) \leftarrow H.Setup(1^\lambda, 1^{sp}, J_i)$.
 - * $(hk_{high}, vk_{high}, td_{high}) \leftarrow H.Setup(1^\lambda, 1^{sp}, J_{i+1}[1, \dots, t-1])$.
- The challenger sets $crs = (crs_{BARG}, hk_0, hk_1, vk_0, vk_1)$ and runs \mathcal{A}_2 on input $(crs, st_{\mathcal{A}})$. Algorithm \mathcal{A}_2 outputs a proof string $\pi = (dig_0, dig_1, \pi_{BARG})$. Let $\hat{x} = (\hat{x}_1, \dots, \hat{x}_{sp})$ and C_{aug} be as defined in Prove and Verify in [Construction 4.4](#). The output of the experiment is 1 if all of the following conditions hold (and 0 otherwise):
- * $H.ValidateDigest(vk_0, dig_0) = H.ValidateDigest(vk_1, dig_1) = 1$.
 - * $Verify'(vk_{BARG}, C_{aug}, \hat{x}, \pi_{BARG}) = 1$.
 - * If $i \geq 1$, then $H.Extract(td_{low}, dig_{low}) = Matching$.
 - * $H.Extract(td_{high}, dig_{high}) = Matching$.
 - * $C_{aug}(\hat{x}_{J_{i+1}[t]}, \hat{w}) = 1$ and $b = 0$ where $\hat{w} = (b, \sigma^{(0)}, \sigma^{(1)}, w) \leftarrow Extract'(td_{BARG}, C_{aug}, \hat{x}, \pi_{BARG}, J_{i+1}[t])$.
- $Hyb_{i,t,5}$ for $t \in [|J_{i+1}|]$: Same as $Hyb_{i,t,4}$ but hk_{high} now binds on $J_{i+1}[1, \dots, t]$.
 - **Phase 1:** Same as Hyb_i .
 - **Phase 2:** The challenger samples the following components:
 - * $(crs_{BARG}, td_{BARG}) \leftarrow TrapGen'(1^\lambda, 1^{sp}, 1^{n'}, 1^{s'}, J_{i+1}[t])$.
 - * If $i = 0$, then $(hk_{low}, vk_{low}, td_{low}) \leftarrow H.Setup(1^\lambda, 1^{sp}, \emptyset)$. Otherwise, it samples $(hk_{low}, vk_{low}, td_{low}) \leftarrow H.Setup(1^\lambda, 1^{sp}, J_i)$.
 - * $(hk_{high}, vk_{high}, td_{high}) \leftarrow H.Setup(1^\lambda, 1^{sp}, J_{i+1}[1, \dots, t])$.
- The challenger sets $crs = (crs_{BARG}, hk_0, hk_1, vk_0, vk_1)$ and runs \mathcal{A}_2 on input $(crs, st_{\mathcal{A}})$. Algorithm \mathcal{A}_2 outputs a proof string $\pi = (dig_0, dig_1, \pi_{BARG})$. Let $\hat{x} = (\hat{x}_1, \dots, \hat{x}_{sp})$ and C_{aug} be as defined in Prove and Verify in [Construction 4.4](#). The output of the experiment is 1 if all of the following conditions hold (and 0 otherwise):
- * $H.ValidateDigest(vk_0, dig_0) = H.ValidateDigest(vk_1, dig_1) = 1$.
 - * $Verify'(vk_{BARG}, C_{aug}, \hat{x}, \pi_{BARG}) = 1$.
 - * If $i \geq 1$, then $H.Extract(td_{low}, dig_{low}) = Matching$.
 - * $H.Extract(td_{high}, dig_{high}) = Matching$.
 - * $C_{aug}(\hat{x}_{J_{i+1}[t]}, \hat{w}) = 1$ and $b = 0$ where $\hat{w} = (b, \sigma^{(0)}, \sigma^{(1)}, w) \leftarrow Extract'(td_{BARG}, C_{aug}, \hat{x}, \pi_{BARG}, J_{i+1}[t])$.
- $Hyb_{i,t,6}$ for $t \in [|J_{i+1}|]$: Same as $Hyb_{i,t,5}$ but the challenger does not check the extracted witnesses.
 - **Phase 1:** Same as Hyb_i .
 - **Phase 2:** The challenger samples the following components:
 - * $(crs_{BARG}, td_{BARG}) \leftarrow TrapGen'(1^\lambda, 1^{sp}, 1^{n'}, 1^{s'}, J_{i+1}[t])$.
 - * If $i = 0$, then $(hk_{low}, vk_{low}, td_{low}) \leftarrow H.Setup(1^\lambda, 1^{sp}, \emptyset)$. Otherwise, it samples $(hk_{low}, vk_{low}, td_{low}) \leftarrow H.Setup(1^\lambda, 1^{sp}, J_i)$.
 - * $(hk_{high}, vk_{high}, td_{high}) \leftarrow H.Setup(1^\lambda, 1^{sp}, J_{i+1}[1, \dots, t])$.

The challenger sets $\text{crs} = (\text{crs}_{\text{BARG}}, \text{hk}_0, \text{hk}_1, \text{vk}_0, \text{vk}_1)$ and runs \mathcal{A}_2 on input $(\text{crs}, \text{st}_{\mathcal{A}})$. Algorithm \mathcal{A}_2 outputs a proof string $\pi = (\text{dig}_0, \text{dig}_1, \pi_{\text{BARG}})$. Let $\hat{\mathbf{x}} = (\hat{x}_1, \dots, \hat{x}_{s_p})$ and C_{aug} be as defined in Prove and Verify in [Construction 4.4](#). The output of the experiment is 1 if all of the following conditions hold (and 0 otherwise):

- * $\text{H.ValidateDigest}(\text{vk}_0, \text{dig}_0) = \text{H.ValidateDigest}(\text{vk}_1, \text{dig}_1) = 1$.
- * $\text{Verify}'(\text{vk}_{\text{BARG}}, C_{\text{aug}}, \hat{\mathbf{x}}, \pi_{\text{BARG}}) = 1$.
- * If $i \geq 1$, then $\text{H.Extract}(\text{td}_{\text{low}}, \text{dig}_{\text{low}}) = \text{Matching}$.
- * $\text{H.Extract}(\text{td}_{\text{high}}, \text{dig}_{\text{high}}) = \text{Matching}$.

In particular, the challenger does *not* extract a witness $\hat{w} \leftarrow \text{Extract}'(\text{td}_{\text{BARG}}, C_{\text{aug}}, \hat{\mathbf{x}}, \pi_{\text{BARG}}, J_{i+1}[t])$ nor checks any conditions on it.

- $\text{Hyb}_{i,t,7}$ for $t \in [|J_{i+1}|]$: Same as $\text{Hyb}_{i,t,6}$, except the BARG is restored to normal mode.

– **Phase 1:** Same as Hyb_i .

– **Phase 2:** The challenger samples the following components:

- * $\text{crs}_{\text{BARG}} \leftarrow \text{Gen}'(1^\lambda, 1^{s_p}, 1^{n'}, 1^{s'})$.
- * If $i = 0$, then $(\text{hk}_{\text{low}}, \text{vk}_{\text{low}}, \text{td}_{\text{low}}) \leftarrow \text{H.Setup}(1^\lambda, 1^{s_p}, \emptyset)$. Otherwise, it samples $(\text{hk}_{\text{low}}, \text{vk}_{\text{low}}, \text{td}_{\text{low}}) \leftarrow \text{H.Setup}(1^\lambda, 1^{s_p}, J_i)$.
- * $(\text{hk}_{\text{high}}, \text{vk}_{\text{high}}, \text{td}_{\text{high}}) \leftarrow \text{H.Setup}(1^\lambda, 1^{s_p}, J_{i+1}[1, \dots, t])$.

The challenger sets $\text{crs} = (\text{crs}_{\text{BARG}}, \text{hk}_0, \text{hk}_1, \text{vk}_0, \text{vk}_1)$ and runs \mathcal{A}_2 on input $(\text{crs}, \text{st}_{\mathcal{A}})$. Algorithm \mathcal{A}_2 outputs a proof string $\pi = (\text{dig}_0, \text{dig}_1, \pi_{\text{BARG}})$. Let $\hat{\mathbf{x}} = (\hat{x}_1, \dots, \hat{x}_{s_p})$ and C_{aug} be as defined in Prove and Verify in [Construction 4.4](#). The output of the experiment is 1 if all of the following conditions hold (and 0 otherwise):

- * $\text{H.ValidateDigest}(\text{vk}_0, \text{dig}_0) = \text{H.ValidateDigest}(\text{vk}_1, \text{dig}_1) = 1$.
- * $\text{Verify}'(\text{vk}_{\text{BARG}}, C_{\text{aug}}, \hat{\mathbf{x}}, \pi_{\text{BARG}}) = 1$.
- * If $i \geq 1$, then $\text{H.Extract}(\text{td}_{\text{low}}, \text{dig}_{\text{low}}) = \text{Matching}$.
- * $\text{H.Extract}(\text{td}_{\text{high}}, \text{dig}_{\text{high}}) = \text{Matching}$.

- $\text{Hyb}_{i,\text{final}}$: Same as $\text{Hyb}_{i,|J_{i+1}|,7}$, but we no longer checks that dig_{low} matches the binding set J_i .

– **Phase 1:** Same as Hyb_i .

– **Phase 2:** The challenger samples the following components:

- * $\text{crs}_{\text{BARG}} \leftarrow \text{Gen}'(1^\lambda, 1^{s_p}, 1^{n'}, 1^{s'})$.
- * If $i = 0$, then $(\text{hk}_{\text{low}}, \text{vk}_{\text{low}}, \text{td}_{\text{low}}) \leftarrow \text{H.Setup}(1^\lambda, 1^{s_p}, \emptyset)$. Otherwise, it samples $(\text{hk}_{\text{low}}, \text{vk}_{\text{low}}, \text{td}_{\text{low}}) \leftarrow \text{H.Setup}(1^\lambda, 1^{s_p}, J_i)$.
- * $(\text{hk}_{\text{high}}, \text{vk}_{\text{high}}, \text{td}_{\text{high}}) \leftarrow \text{H.Setup}(1^\lambda, 1^{s_p}, J_{i+1})$.

The challenger sets $\text{crs} = (\text{crs}_{\text{BARG}}, \text{hk}_0, \text{hk}_1, \text{vk}_0, \text{vk}_1)$ and runs \mathcal{A}_2 on input $(\text{crs}, \text{st}_{\mathcal{A}})$. Algorithm \mathcal{A}_2 outputs a proof string $\pi = (\text{dig}_0, \text{dig}_1, \pi_{\text{BARG}})$. Let $\hat{\mathbf{x}} = (\hat{x}_1, \dots, \hat{x}_{s_p})$ and C_{aug} be as defined in Prove and Verify in [Construction 4.4](#). The output of the experiment is 1 if all of the following conditions hold (and 0 otherwise):

- * $\text{H.ValidateDigest}(\text{vk}_0, \text{dig}_0) = \text{H.ValidateDigest}(\text{vk}_1, \text{dig}_1) = 1$.
- * $\text{Verify}'(\text{vk}_{\text{BARG}}, C_{\text{aug}}, \hat{\mathbf{x}}, \pi_{\text{BARG}}) = 1$.
- * $\text{H.Extract}(\text{td}_{\text{high}}, \text{dig}_{\text{high}}) = \text{Matching}$.

In particular, the challenger no longer checks if $\text{H.Extract}(\text{td}_{\text{low}}, \text{dig}_{\text{low}}) = \text{Matching}$ when $i > 1$.

4.2.2 Analysis of Hybrid Experiments

We now show that the probability of a hybrid experiment outputting 1 cannot decrease by a non-negligible amount when transitioning from one hybrid to the next. The goal is to eventually show that $\Pr[\text{Hyb}_{i-1}(\mathcal{A}) = 1]$ is negligibly close to $\Pr[\text{Hyb}_i(\mathcal{A}) = 1]$ for all $i \in [d]$. We argue this via a sequence of *non-uniform* reductions to the security properties of the underlying zero-fixing hash function and BARG. Specifically, our reduction algorithms construct

a non-uniform adversary where there is an initial (inefficient) preprocessing phase that outputs an advice string of polynomial size, and a polynomial-time online algorithm that takes the advice as input and interacts with the challenger according to the specifications of the target security game. Our reductions share a common preprocessing phase, which we abstract out as a standalone Preprocess algorithm defined as follows:

- **Preprocess**(C, P, \mathbf{x}): On input a Boolean circuit $C: \{0, 1\}^n \times \{0, 1\}^h \rightarrow \{0, 1\}$, a monotone Boolean policy $P: \{0, 1\}^k \rightarrow \{0, 1\}$ of size s , and an instance $\mathbf{x} = (x_1, \dots, x_k) \in \{0, 1\}^{kn}$, the preprocessing algorithm first checks if $(C, P, \mathbf{x}) \in \mathcal{L}_{\text{MP-CSAT}}$. If so, it outputs \perp . Otherwise, it computes $\beta_i \in \{0, 1\}$ for $i \in [s]$ as follows:
 - For $i \in [k]$, set $\beta_i = 1$ if and only if $(C, x_i) \in \mathcal{L}_{\text{CSAT}}$.
 - For $i \in [k + 1, s]$, set β_i to be the value of the wire i in the evaluation of P on $(\beta_1, \dots, \beta_k)$.

Output $\tau = (C, P, \mathbf{x}, (\beta_1, \dots, \beta_s))$.

We now analyze each pair of adjacent hybrid experiments.

Claim 4.8. *If Π_{H} satisfies extractor validity against efficient non-uniform adversaries, then there exists a negligible function $\text{negl}(\cdot)$ such that for every $i \in \{0, \dots, d - 1\}$, it holds that*

$$|\Pr[\text{Hyb}_i(\mathcal{A}) = 1] - \Pr[\text{Hyb}_{i,1,1}(\mathcal{A}) = 1]| \leq \text{negl}(\lambda).$$

Proof. Take any $i \in \{0, \dots, d - 1\}$ and suppose $|\Pr[\text{Hyb}_i(\mathcal{A}) = 1] - \Pr[\text{Hyb}_{i,1,1}(\mathcal{A}) = 1]| = \varepsilon(\lambda)$ for some non-negligible ε . By construction, the only difference between Hyb_i and $\text{Hyb}_{i,1,1}$ is the *additional* check in $\text{Hyb}_{i,1,1}$:

$$\text{H.Extract}(\text{td}_{\text{high}}, \text{dig}_{\text{high}}) = \text{Matching}.$$

Thus, with probability at least ε , the adversary \mathcal{A} in an execution of $\text{Hyb}_{i,1,1}$ and Hyb_i outputs a proof $\pi = (\text{dig}_0, \text{dig}_1, \pi_{\text{BARG}})$ where $\text{H.ValidateDigest}(\text{vk}_{\text{high}}, \text{dig}_{\text{high}}) = 1$ and $\text{H.Extract}(\text{td}_{\text{high}}, \text{dig}_{\text{high}}) = \text{NotMatching}$. In all other cases, the output of the two experiments are identical. We use \mathcal{A} to construct a non-uniform adversary \mathcal{B} that breaks extract validity of Π_{H} as follows:

- **Preprocessing phase:** On input the security parameter 1^λ , run $(1^n, 1^{s_c}, 1^{s_p}, C, P, \mathbf{x}, \text{st}_{\mathcal{A}}) \leftarrow \mathcal{A}_1(1^\lambda)$. Compute $\tau \leftarrow \text{Preprocess}(C, P, \mathbf{x})$ and output $\text{st}_{\mathcal{B}} = (1^n, 1^{s_c}, 1^{s_p}, \tau, \text{st}_{\mathcal{A}})$.
- **Online phase:** On input the state $\text{st}_{\mathcal{B}} = (1^n, 1^{s_c}, 1^{s_p}, \tau, \text{st}_{\mathcal{A}})$ where $\tau = (C, P, \mathbf{x}, (\beta_1, \dots, \beta_s))$, proceed as follows:
 1. Send 1^{s_p} and the set \emptyset to the challenger. The challenger replies with a hash key hk and a verification key vk .
 2. Sample $(\text{crs}_{\text{BARG}}, \text{vk}_{\text{BARG}}) \leftarrow \text{Gen}'(1^\lambda, 1^{s_p}, 1^{n'}, 1^{s'})$, where n', s' are defined as in [Construction 4.4](#). If $i = 0$, sample $(\text{hk}_{\text{low}}, \text{vk}_{\text{low}}, \text{td}_{\text{low}}) \leftarrow \text{H.Setup}(1^\lambda, 1^{s_p}, \emptyset)$. Otherwise, if $i \neq 0$, sample $(\text{hk}_{\text{low}}, \text{vk}_{\text{low}}, \text{td}_{\text{low}}) \leftarrow \text{H.Setup}(1^\lambda, 1^{s_p}, J_i)$, where $J_i = \{j \in \text{layer}_i(P) : \beta_j = 0\}$. Let $\text{hk}_{\text{high}} \leftarrow \text{hk}$ and $\text{vk}_{\text{high}} \leftarrow \text{vk}$.
 3. Let $\text{crs} = (\text{crs}_{\text{BARG}}, \text{hk}_0, \text{hk}_1, \text{vk}_0, \text{vk}_1)$ and run $\pi \leftarrow \mathcal{A}_2(\text{crs}, \text{st}_{\mathcal{A}})$. Parse $\pi = (\text{dig}_0, \text{dig}_1, \pi_{\text{BARG}})$ and output dig_{high} .

By construction, algorithm \mathcal{B} perfectly simulates an execution of Hyb_i and $\text{Hyb}_{i,1,1}$ for \mathcal{A} , so with probability at least ε , the digest dig_{high} satisfies $\text{H.ValidateDigest}(\text{vk}_{\text{high}}, \text{dig}_{\text{high}}) = 1$ and $\text{H.Extract}(\text{td}_{\text{high}}, \text{dig}_{\text{high}}) = \text{NotMatching}$. Correspondingly, algorithm \mathcal{B} breaks extractor validity with advantage ε . \square

Claim 4.9. *If Π_{BARG} satisfies set hiding against efficient non-uniform adversaries, then there exists a negligible function $\text{negl}(\cdot)$ such that for every $i \in \{0, \dots, d - 1\}$ and $t \in [J_{i+1}]$, it holds that*

$$|\Pr[\text{Hyb}_{i,t,1}(\mathcal{A}) = 1] - \Pr[\text{Hyb}_{i,t,2}(\mathcal{A}) = 1]| \leq \text{negl}(\lambda).$$

Proof. Take any $i \in \{0, \dots, d - 1\}$ and $t \in [J_{i+1}]$. Suppose $|\Pr[\text{Hyb}_{i,t,1}(\mathcal{A}) = 1] - \Pr[\text{Hyb}_{i,t,2}(\mathcal{A}) = 1]| = \varepsilon(\lambda)$ for some non-negligible ε . We construct a non-uniform adversary \mathcal{B} that breaks set hiding of Π_{BARG} :

- **Preprocessing phase:** On input the security parameter 1^λ , run $(1^n, 1^{sc}, 1^{sp}, C, P, \mathbf{x}, \text{st}_{\mathcal{A}}) \leftarrow \mathcal{A}_1(1^\lambda)$. Compute $\tau \leftarrow \text{Preprocess}(C, P, \mathbf{x})$ and output $\text{st}_{\mathcal{B}} = (1^n, 1^{sc}, 1^{sp}, \tau, \text{st}_{\mathcal{A}})$.
- **Online phase:** On input the state $\text{st}_{\mathcal{B}} = (1^n, 1^{sc}, 1^{sp}, \tau, \text{st}_{\mathcal{A}})$ where $\tau = (C, P, \mathbf{x}, (\beta_1, \dots, \beta_s))$, proceed as follows:
 1. Compute $J_i = \{j \in \text{layer}_i(P) : \beta_j = 0\}$ and $J_{i+1} = \{j \in \text{layer}_{i+1}(P) : \beta_j = 0\}$.
 2. Send $1^{sp}, 1^{n'}, 1^{s'}$, and the index $J_{i+1}[t]$ to the challenger, where n', s' are computed as in [Construction 4.4](#). The challenger replies with crs_{BARG} .
 3. If $i = 0$, then $(\text{hk}_{\text{low}}, \text{vk}_{\text{low}}, \text{td}_{\text{low}}) \leftarrow \text{H.Setup}(1^\lambda, 1^{sp}, \emptyset)$. Otherwise, sample $(\text{hk}_{\text{low}}, \text{vk}_{\text{low}}, \text{td}_{\text{low}}) \leftarrow \text{H.Setup}(1^\lambda, 1^{sp}, J_i)$.
 4. Sample $(\text{hk}_{\text{high}}, \text{vk}_{\text{high}}, \text{td}_{\text{high}}) \leftarrow \text{H.Setup}(1^\lambda, 1^{sp}, J_{i+1}[1, \dots, t-1])$.
 5. Let $\text{crs} = (\text{crs}_{\text{BARG}}, \text{hk}_0, \text{hk}_1, \text{vk}_0, \text{vk}_1)$ and run $\pi \leftarrow \mathcal{A}_2(\text{crs}, \text{st}_{\mathcal{A}})$.
 6. Let $\pi = (\text{dig}_0, \text{dig}_1, \pi_{\text{BARG}})$ and let $\hat{\mathbf{x}} = (\hat{x}_1, \dots, \hat{x}_{s_p})$ and C_{aug} be as defined in Prove and Verify in [Construction 4.4](#). Output 1 if all of the following conditions hold (and abort with \perp otherwise):
 - $\text{H.ValidateDigest}(\text{vk}_0, \text{dig}_0) = \text{H.ValidateDigest}(\text{vk}_1, \text{dig}_1) = 1$.
 - $\text{Verify}'(\text{vk}_{\text{BARG}}, C_{\text{aug}}, \hat{\mathbf{x}}, \pi_{\text{BARG}}) = 1$.
 - If $i \geq 1$ then $\text{H.Extract}(\text{td}_{\text{low}}, \text{dig}_{\text{low}}) = \text{Matching}$.
 - $\text{H.Extract}(\text{td}_{\text{high}}, \text{dig}_{\text{high}}) = \text{Matching}$.

We consider two possibilities:

- In $\text{ExptSH}_{\mathcal{B}}(\lambda, 0)$, the challenger samples $\text{crs}_{\text{BARG}} \leftarrow \text{Gen}'(1^\lambda, 1^{sp}, 1^{n'}, 1^{s'})$. In this case, by the construction of \mathcal{B} , it holds that crs is sampled exactly as in $\text{Hyb}_{i,t,1}$. Moreover, \mathcal{B} computes its output exactly as specified by $\text{Hyb}_{i,t,1}$. This means that \mathcal{B} perfectly simulates $\text{Hyb}_{i,t,1}(\mathcal{A})$ and thus

$$\Pr[\text{ExptSH}_{\mathcal{B}}(\lambda, 0) = 1] = \Pr[\text{Hyb}_{i,t,1}(\mathcal{A}) = 1].$$

- In $\text{ExptSH}_{\mathcal{B}}(\lambda, 1)$, the challenger samples $\text{crs}_{\text{BARG}} \leftarrow \text{TrapGen}'(1^\lambda, 1^{sp}, 1^{n'}, 1^{s'}, J_{i+1}[t])$. In this case, by the construction of \mathcal{B} , it holds that crs is sampled exactly as in $\text{Hyb}_{i,t,2}$. Moreover, \mathcal{B} computes its output exactly as specified by $\text{Hyb}_{i,t,2}$. This means that \mathcal{B} perfectly simulates $\text{Hyb}_{i,t,2}(\mathcal{A})$ and thus

$$\Pr[\text{ExptSH}_{\mathcal{B}}(\lambda, 1) = 1] = \Pr[\text{Hyb}_{i,t,2}(\mathcal{A}) = 1].$$

We conclude that algorithm \mathcal{B} breaks the index hiding property of Π_{BARG} with the same advantage ε . \square

Claim 4.10. *If Π_{BARG} satisfies somewhere extractability in trapdoor mode against efficient non-uniform adversaries, then there exists a negligible function $\text{negl}(\cdot)$ such that for every $i \in \{0, \dots, d-1\}, t \in [J_{i+1}]$, it holds that*

$$|\Pr[\text{Hyb}_{i,t,2}(\mathcal{A}) = 1] - \Pr[\text{Hyb}_{i,t,3}(\mathcal{A}) = 1]| \leq \text{negl}(\lambda).$$

Proof. Take any $i \in \{0, \dots, d-1\}$ and $t \in [J_{i+1}]$. Suppose $|\Pr[\text{Hyb}_{i,t,2}(\mathcal{A}) = 1] - \Pr[\text{Hyb}_{i,t,3}(\mathcal{A}) = 1]| = \varepsilon(\lambda)$ for some non-negligible ε . By construction, the only difference between $\text{Hyb}_{i,t,2}$ and $\text{Hyb}_{i,t,3}$ is the *additional* check in $\text{Hyb}_{i,t,3}$:

$$C_{\text{aug}}(\hat{x}_{J_{i+1}[t]}, \hat{w}) = 1 \text{ where } \hat{w} \leftarrow \text{Extract}'(\text{td}_{\text{BARG}}, C_{\text{aug}}, \hat{\mathbf{x}}, \pi_{\text{BARG}}, J_{i+1}[t]). \quad (4.3)$$

Thus, with probability at least ε , the adversary \mathcal{A} in an execution of $\text{Hyb}_{i,t,2}$ and $\text{Hyb}_{i,t,3}$ outputs a proof $\pi = (\text{dig}_0, \text{dig}_1, \pi_{\text{BARG}})$ where $\text{Verify}'(\text{vk}_{\text{BARG}}, C_{\text{aug}}, \hat{\mathbf{x}}, \pi_{\text{BARG}}) = 1$ and [Eq. \(4.3\)](#) does *not* hold (i.e., $C_{\text{aug}}(\hat{x}_{J_{i+1}[t]}, \hat{w}) = 0$). In all other cases, the outputs of $\text{Hyb}_{i,t,2}$ and $\text{Hyb}_{i,t,3}$ are identical. We use \mathcal{A} to construct an adversary \mathcal{B} for the somewhere extractability game of Π_{BARG} :

- **Preprocessing phase:** On input the security parameter 1^λ , run $(1^n, 1^{sc}, 1^{sp}, C, P, \mathbf{x}, \text{st}_{\mathcal{A}}) \leftarrow \mathcal{A}_1(1^\lambda)$. Compute $\tau \leftarrow \text{Preprocess}(C, P, \mathbf{x})$ and output $\text{st}_{\mathcal{B}} = (1^n, 1^{sc}, 1^{sp}, \tau, \text{st}_{\mathcal{A}})$.

• **Online phase:** On input the state $\text{st}_{\mathcal{B}} = (1^n, 1^{s_c}, 1^{s_p}, \tau, \text{st}_{\mathcal{A}})$ where $\tau = (C, P, \mathbf{x}, (\beta_1, \dots, \beta_s))$, proceed as follows:

1. Compute $J_i = \{j \in \text{layer}_i(P) : \beta_j = 0\}$ and $J_{i+1} = \{j \in \text{layer}_{i+1}(P) : \beta_j = 0\}$.
2. Send $1^{s_p}, 1^{n'}, 1^{s'}$ and the index $J_{i+1}[t]$ to the challenger, where n', s' are computed as in [Construction 4.4](#). The challenger replies with crs_{BARG} .
3. If $i = 0$, then $(\text{hk}_{\text{low}}, \text{vk}_{\text{low}}, \text{td}_{\text{low}}) \leftarrow \text{H.Setup}(1^\lambda, 1^{s_p}, \emptyset)$. Otherwise, sample $(\text{hk}_{\text{low}}, \text{vk}_{\text{low}}, \text{td}_{\text{low}}) \leftarrow \text{H.Setup}(1^\lambda, 1^{s_p}, J_i)$.
4. Sample $(\text{hk}_{\text{high}}, \text{vk}_{\text{high}}, \text{td}_{\text{high}}) \leftarrow \text{H.Setup}(1^\lambda, 1^{s_p}, J_{i+1}[1, \dots, t-1])$.
5. Let $\text{crs} = (\text{crs}_{\text{BARG}}, \text{hk}_0, \text{hk}_1, \text{vk}_0, \text{vk}_1)$ and run $\pi \leftarrow \mathcal{A}_2(\text{crs}, \text{st}_{\mathcal{A}})$. Let $\hat{\mathbf{x}} = (\hat{x}_1, \dots, \hat{x}_{s_p})$ and C_{aug} be as defined in Prove and Verify in [Construction 4.4](#). Parse $\pi = (\text{dig}_0, \text{dig}_1, \pi_{\text{BARG}})$ and output the circuit C_{aug} , the statements $\hat{\mathbf{x}}$, and the proof π_{BARG} .

By construction, algorithm \mathcal{B} perfectly simulates an execution of $\text{Hyb}_{i,t,2}$ and $\text{Hyb}_{i,t,3}$ for \mathcal{A} , so with probability ε , it outputs π_{BARG} such that $\text{Verify}'(\text{vk}_{\text{BARG}}, C_{\text{aug}}, \hat{\mathbf{x}}, \pi_{\text{BARG}}) = 1$ and [Eq. \(4.3\)](#) does not hold. In particular, this means the proof π_{BARG} verifies with respect to crs_{BARG} and yet $C_{\text{aug}}(\hat{x}_{J_{i+1}[t]}, \hat{w}) = 0$ where $\hat{w} \leftarrow \text{Extract}'(\text{td}_{\text{BARG}}, C_{\text{aug}}, \hat{\mathbf{x}}, \pi_{\text{BARG}}, J_{i+1}[t])$, and td_{BARG} is the trapdoor associated with crs_{BARG} that the challenger sampled. This means \mathcal{B} wins the somewhere extractability game with the same advantage ε . \square

Claim 4.11. *If Π_{H} satisfies zero-fixing against efficient non-uniform adversaries, then there exists a negligible function $\text{negl}(\cdot)$ such that for every $i \in \{0, \dots, d-1\}$, $t \in [|J_{i+1}|]$, it holds that*

$$|\Pr[\text{Hyb}_{i,t,3}(\mathcal{A}) = 1] - \Pr[\text{Hyb}_{i,t,4}(\mathcal{A}) = 1]| \leq \text{negl}(\lambda).$$

Proof. By construction, the only difference between $\text{Hyb}_{i,t,3}$ and $\text{Hyb}_{i,t,4}$ is the *additional* check in $\text{Hyb}_{i,t,4}$ that the extracted bit b satisfies $b = 0$. We consider two cases in our analysis:

- Suppose $i = 0$. In this case, $J_1[t]$ refers to an *input* wire in P , which means $J_1[t] \leq k$. Suppose $\text{Hyb}_{i,t,3}(\mathcal{A})$ outputs 1. This means that $C_{\text{aug}}(\hat{x}_{J_1[k]}, \hat{w}) = 1$ where

$$\hat{w} = (b, \sigma^{(0)}, \sigma^{(1)}, w) \leftarrow \text{Extract}'(\text{td}_{\text{BARG}}, C_{\text{aug}}, \hat{\mathbf{x}}, \pi_{\text{BARG}}, J_1[t]).$$

Since $J_1[k] \leq k$ and by construction of $\hat{\mathbf{x}}$, we have that $\hat{x}_{J_1[k]} = (J_1[k], x_{J_1[k]})$. By definition of C_{aug} (see [Fig. 1](#)), we have that $C_{\text{aug}}(\hat{x}_{J_1[k]}, \hat{w}) = 1$ only if $C(x_{J_1[k]}, w) = b$. However, by definition of J_1 , it must be the case that $(C, x_{J_1[k]}) \notin \mathcal{L}_{\text{CSAT}}$. This means $C(x_{J_1[k]}, w) = 0 = b$. In this case, $\text{Hyb}_{i,t,4}(\mathcal{A})$ also outputs 1. Conversely, since the verification conditions in $\text{Hyb}_{i,t,4}$ are a *superset* of the conditions in $\text{Hyb}_{i,t,3}$, if $\text{Hyb}_{i,t,4}(\mathcal{A}) = 1$, then $\text{Hyb}_{i,t,3}(\mathcal{A}) = 1$. We conclude that in this case

$$\Pr[\text{Hyb}_{i,t,3}(\mathcal{A}) = 1] = \Pr[\text{Hyb}_{i,t,4} = 1].$$

- Suppose $i > 0$. In this case, security reduces to the zero-fixing security of Π_{H} . We give this proof below.

To argue the second case, take any $i \in \{1, \dots, d-1\}$ and $t \in [|J_{i+1}|]$, and suppose that

$$|\Pr[\text{Hyb}_{i,t,3}(\mathcal{A}) = 1] - \Pr[\text{Hyb}_{i,t,4}(\mathcal{A}) = 1]| = \varepsilon(\lambda)$$

for some non-negligible ε . By construction, the only difference between $\text{Hyb}_{i,t,3}$ and $\text{Hyb}_{i,t,4}$ is the *additional* check in $\text{Hyb}_{i,t,4}$ that $b = 0$ where $\hat{w} = (b, \sigma^{(0)}, \sigma^{(1)}, w) \leftarrow \text{Extract}'(\text{td}_{\text{BARG}}, C_{\text{aug}}, \hat{\mathbf{x}}, \pi_{\text{BARG}}, J_{i+1}[t])$. Thus, with probability at least ε , the adversary \mathcal{A} in an execution of $\text{Hyb}_{i,t,3}$ and $\text{Hyb}_{i,t,4}$ will output a proof $\pi = (\text{dig}_0, \text{dig}_1, \pi_{\text{BARG}})$ where

$$\text{H.Extract}(\text{td}_{\text{low}}, \text{dig}_{\text{low}}) = \text{Matching} \quad \text{and} \quad C_{\text{aug}}(\hat{x}_{J_{i+1}[t]}, \hat{w}) = 1 \quad \text{and} \quad b \neq 0. \quad (4.4)$$

In all other cases, the outputs of $\text{Hyb}_{i,t,3}$ and $\text{Hyb}_{i,t,4}$ are identical. We use \mathcal{A} to construct an adversary \mathcal{B} for the zero-fixing game for Π_{H} :

- **Preprocessing phase:** On input the security parameter 1^λ , run $(1^n, 1^{s_c}, 1^{s_p}, C, P, \mathbf{x}, \text{st}_{\mathcal{A}}) \leftarrow \mathcal{A}_1(1^\lambda)$. Compute $\tau \leftarrow \text{Preprocess}(C, P, \mathbf{x})$ and output $\text{st}_{\mathcal{B}} = (1^n, 1^{s_c}, 1^{s_p}, \tau, \text{st}_{\mathcal{A}})$.
- **Online phase:** On input the state $\text{st}_{\mathcal{B}} = (1^n, 1^{s_c}, 1^{s_p}, \tau, \text{st}_{\mathcal{A}})$ where $\tau = (C, P, \mathbf{x}, (\beta_1, \dots, \beta_s))$, proceed as follows:
 1. Compute $J_i = \{j \in \text{layer}_i(P) : \beta_j = 0\}$ and $J_{i+1} = \{j \in \text{layer}_{i+1}(P) : \beta_j = 0\}$.
 2. Send 1^{s_p} and the set J_i to the challenger. The challenger replies with a hash key hk and a verification key vk .
 3. Set $(\text{hk}_{\text{high}}, \text{vk}_{\text{high}}, \text{td}_{\text{high}}) \leftarrow \text{H.Setup}(1^\lambda, 1^{s_p}, J_{i+1}[1, \dots, t-1])$, $\text{hk}_{\text{low}} \leftarrow \text{hk}$, and $\text{vk}_{\text{low}} \leftarrow \text{vk}$. Finally, sample $(\text{crs}_{\text{BARG}}, \text{td}_{\text{BARG}}) \leftarrow \text{TrapGen}'(1^\lambda, 1^{s_p}, 1^{n'}, 1^{s'}, J_{i+1}[t])$, where s' is defined as in [Construction 4.4](#).
 4. Let $\text{crs} = (\text{crs}_{\text{BARG}}, \text{hk}_0, \text{hk}_1, \text{vk}_0, \text{vk}_1)$ and run $\pi \leftarrow \mathcal{A}_2(\text{crs}, \text{st}_{\mathcal{A}})$.
 5. Let $\pi = (\text{dig}_0, \text{dig}_1, \pi_{\text{BARG}})$ and suppose $\hat{\mathbf{x}} = (\hat{x}_1, \dots, \hat{x}_{s_p})$ and C_{aug} be as defined in Prove and Verify in [Construction 4.4](#). Compute

$$\hat{w} = (b, \sigma^{(0)}, \sigma^{(1)}, w) \leftarrow \text{Extract}'(\text{td}_{\text{BARG}}, C_{\text{aug}}, \hat{\mathbf{x}}, \pi_{\text{BARG}}, J_{i+1}[t]),$$

and parse the extracted witness w as $w = (b_1, \sigma_1^{(0)}, \sigma_1^{(1)}, b_2, \sigma_2^{(0)}, \sigma_2^{(1)})$.

6. Parse $\hat{x}_{J_{i+1}[t]} = (g, j_1, j_2)$. If there exists $\alpha \in \{1, 2\}$ such that $j_\alpha \in J_i$ and $b_\alpha = 1$, output the digest dig_{low} , the index j_α , and the opening $\sigma_\alpha^{(\text{low})}$.

By construction, algorithm \mathcal{B} perfectly simulates an execution of $\text{Hyb}_{i,t,3}$ and $\text{Hyb}_{i,t,4}$ for \mathcal{A} . Thus, with probability at least ε , algorithm \mathcal{A} will output a proof $\pi = (\text{dig}_0, \text{dig}_1, \pi_{\text{BARG}})$ where [Eq. \(4.4\)](#) holds. Since $C_{\text{aug}}(\hat{x}_{J_{i+1}[t]}, \hat{w}) = 1$, the following properties hold:

- $b = g(b_1, b_2)$, where $g = g_{J_{i+1}[t]} \in \{\text{AND}, \text{OR}\}$ is the gate in the circuit P that computes wire $J_{i+1}[t]$.
- $\text{H.VerOpen}(\text{vk}_{\text{low}}, \text{dig}_{\text{low}}, j_1, b_1, \sigma_1^{(\text{low})}) = 1$ and $\text{H.VerOpen}(\text{vk}_{\text{low}}, \text{dig}_{\text{low}}, j_2, b_2, \sigma_2^{(\text{low})}) = 1$.

By definition, j_1, j_2 are the indices of the input wires to the gate whose output wire is $J_{i+1}[t]$. We consider two possibilities:

- Suppose $b_1 \leq \beta_{j_1}$ and $b_2 \leq \beta_{j_2}$. By definition (see the details of the Preprocess algorithm), $\beta_{J_{i+1}[t]} = g(\beta_{j_1}, \beta_{j_2})$. Since $J_{i+1}[t] \in J_{i+1}$, this means $\beta_{J_{i+1}[t]} = 0$. Since g is a monotone gate and $b_1 \leq \beta_{j_1}$ and $b_2 \leq \beta_{j_2}$, we have that $b = g(b_1, b_2) \leq g(\beta_{j_1}, \beta_{j_2}) = 0$. Since $b \in \{0, 1\}$, this means that $b = 0$. However, if [Eq. \(4.4\)](#) holds, then $b \neq 0$, so this case does not happen.
- Suppose there exist $\alpha \in \{1, 2\}$ such that $b_\alpha > \beta_{j_\alpha}$. This means that $\beta_{j_\alpha} = 0$ and $b_\alpha = 1$. Since P is a layered monotone circuit, this means $j_1, j_2 \in \text{layer}_i(P)$. Since $\beta_{j_\alpha} = 0$, this means that $j_\alpha \in J_i$. In conjunction with [Eq. \(4.4\)](#), this means

$$\text{H.Extract}(\text{td}_{\text{low}}, \text{dig}_{\text{low}}) = \text{Matching} \text{ and } \text{H.VerOpen}(\text{vk}_{\text{low}}, \text{dig}_{\text{low}}, j_\alpha, 1, \sigma_\alpha^{(\text{low})}) = 1 \text{ and } j_\alpha \in J_i,$$

where $(\text{hk}_{\text{low}}, \text{vk}_{\text{low}}, \text{td}_{\text{low}}) \leftarrow \text{H.Setup}(1^\lambda, 1^{s_p}, J_i)$ is the hash function parameters sampled by the zero-fixing challenger. In this case, algorithm \mathcal{B} wins the zero-fixing game.

We conclude that if the proof π output by \mathcal{A} satisfies [Eq. \(4.4\)](#) with probability ε , then algorithm \mathcal{B} wins the zero-fixing game with advantage at least ε . \square

Claim 4.12. *If Π_{H} satisfies one-sided index hiding with extracted guess security against efficient non-uniform adversaries, then there exists a negligible function $\text{negl}(\cdot)$ such that for every $i \in \{0, \dots, d-1\}$, $t \in [|\text{layer}_{i+1}|]$, it holds that*

$$\Pr[\text{Hyb}_{i,t,5}(\mathcal{A}) = 1] \geq \Pr[\text{Hyb}_{i,t,4}(\mathcal{A}) = 1] - \text{negl}(\lambda).$$

Proof. Take any $i \in \{0, \dots, d-1\}$ and $t \in [J_{i+1}]$. Suppose $\Pr[\text{Hyb}_{i,t,5}(\mathcal{A}) = 1] \leq \Pr[\text{Hyb}_{i,t,4}(\mathcal{A})] - \varepsilon(\lambda)$ for some non-negligible ε . We use \mathcal{A} to construct a non-uniform adversary \mathcal{B} for the index hiding with extracted guess game of Π_H :

- **Preprocessing phase:** On input the security parameter 1^λ , run $(1^n, 1^{sc}, 1^{sp}, C, P, \mathbf{x}, \text{st}_{\mathcal{A}}) \leftarrow \mathcal{A}_1(1^\lambda)$. Compute $\tau \leftarrow \text{Preprocess}(C, P, \mathbf{x})$ and output $\text{st}_{\mathcal{B}} = (1^n, 1^{sc}, 1^{sp}, \tau, \text{st}_{\mathcal{A}})$.
- **Online phase:** On input the state $\text{st}_{\mathcal{B}} = (1^n, 1^{sc}, 1^{sp}, \tau, \text{st}_{\mathcal{A}})$ where $\tau = (C, P, \mathbf{x}, (\beta_1, \dots, \beta_s))$, proceed as follows:
 1. Compute $J_i = \{j \in \text{layer}_i(P) : \beta_j = 0\}$ and $J_{i+1} = \{j \in \text{layer}_{i+1}(P) : \beta_j = 0\}$.
 2. Send 1^{sp} , the set $J_{i+1}[1, \dots, t]$, and the index $J_{i+1}[t]$ to the challenger. The challenger replies with hk and vk .
 3. If $i = 0$, then $(\text{hk}_{\text{low}}, \text{vk}_{\text{low}}, \text{td}_{\text{low}}) \leftarrow \text{H.Setup}(1^\lambda, 1^{sp}, \emptyset)$. Otherwise, sample $(\text{hk}_{\text{low}}, \text{vk}_{\text{low}}, \text{td}_{\text{low}}) \leftarrow \text{H.Setup}(1^\lambda, 1^{sp}, J_i)$. Let $\text{hk}_{\text{high}} \leftarrow \text{hk}$ and $\text{vk}_{\text{high}} \leftarrow \text{vk}$.
 4. Sample $(\text{crs}_{\text{BARG}}, \text{td}_{\text{BARG}}) \leftarrow \text{TrapGen}'(1^\lambda, 1^{sp}, 1^{n'}, 1^{s'}, J_{i+1}[t])$, where s' is defined as in [Construction 4.4](#).
 5. Let $\text{crs} = (\text{crs}_{\text{BARG}}, \text{hk}_0, \text{hk}_1, \text{vk}_0, \text{vk}_1)$ and run $\pi \leftarrow \mathcal{A}_2(\text{crs}, \text{st}_{\mathcal{A}})$. Let $\hat{\mathbf{x}} = (\hat{x}_1, \dots, \hat{x}_{s_p})$ and C_{aug} be as defined in Prove and Verify in [Construction 4.4](#). Parse $\pi = (\text{dig}_0, \text{dig}_1, \pi_{\text{BARG}})$ and compute $\hat{w} = (b, \sigma^{(0)}, \sigma^{(1)}, w) \leftarrow \text{Extract}'(\text{td}_{\text{BARG}}, C_{\text{aug}}, \hat{\mathbf{x}}, \pi_{\text{BARG}}, J_{i+1}[t])$.
 6. Check each of the following conditions:
 - $\text{H.ValidateDigest}(\text{vk}_0, \text{dig}_0) = \text{H.ValidateDigest}(\text{vk}_1, \text{dig}_1) = 1$.
 - $\text{Verify}'(\text{vk}_{\text{BARG}}, C_{\text{aug}}, \hat{\mathbf{x}}, \pi_{\text{BARG}}) = 1$.
 - If $i \geq 1$, then $\text{H.Extract}(\text{td}_{\text{low}}, \text{dig}_{\text{low}}) = \text{Matching}$.
 - $C_{\text{aug}}(\hat{x}_{J_{i+1}[t]}, \hat{w}) = 1$ and $b = 0$.

If any condition fails to verify, then output \perp . Otherwise, output the digest dig_{high} and the opening $\sigma^{(\text{high})}$.

We consider the two possibilities:

- Suppose the challenger responds according to the specification of $\text{ExptIHE}_{\mathcal{B}}(\lambda, 0)$. In this case, the challenger samples $(\text{hk}_{\text{high}}, \text{vk}_{\text{high}}, \text{td}_{\text{high}}) \leftarrow \text{H.Setup}(1^\lambda, 1^{sp}, J_{i+1}[1, \dots, t-1])$. By construction, algorithm \mathcal{B} perfectly simulates an execution of $\text{Hyb}_{i,t,4}$ for \mathcal{A} . The output of $\text{ExptIHE}_{\mathcal{B}}(\lambda, 0)$ is 1 if and only if *all* of the following events occur:
 - $\text{H.ValidateDigest}(\text{vk}_0, \text{dig}_0) = \text{H.ValidateDigest}(\text{vk}_1, \text{dig}_1) = 1$.
 - $\text{Verify}'(\text{vk}_{\text{BARG}}, C_{\text{aug}}, \hat{\mathbf{x}}, \pi_{\text{BARG}}) = 1$.
 - If $i \geq 1$, then $\text{H.Extract}(\text{td}_{\text{low}}, \text{dig}_{\text{low}}) = \text{Matching}$.
 - $C_{\text{aug}}(\hat{x}_{J_{i+1}[t]}, \hat{w}) = 1$ and $b = 0$.
 - $\text{H.VerOpen}(\text{vk}_{\text{high}}, \text{dig}_{\text{high}}, J_{i+1}[t], 0, \sigma^{(\text{high})}) = 1$ and $\text{H.Extract}(\text{td}_{\text{high}}, \text{dig}_{\text{high}}) = \text{Matching}$.

We now argue that

$$\text{ExptIHE}_{\mathcal{B}}(\lambda, 0) = 1 \iff \text{Hyb}_{i,t,4}(\mathcal{A}) = 1.$$

The forward direction is immediate since the set of conditions under which $\text{ExptIHE}_{\mathcal{B}}(\lambda, 0)$ outputs 1 is a strict superset of the conditions under which $\text{Hyb}_{i,t,4}(\mathcal{A})$ outputs 1. For the backward direction, we show that if $C_{\text{aug}}(\hat{x}_{J_{i+1}[t]}, \hat{w}) = 1$ and $b = 0$, then $\text{H.VerOpen}(\text{vk}_{\text{high}}, \text{dig}_{\text{high}}, J_{i+1}[t], 0, \sigma^{(\text{high})}) = 1$. By construction of C_{aug} , we have that $C_{\text{aug}}(\hat{x}_{J_{i+1}[t]}, \hat{w}) = 1$ implies that for all $\alpha \in \{1, 2\}$, it holds that

$$\text{H.VerOpen}(\text{vk}_\alpha, \text{dig}_\alpha, J_{i+1}[t], 0, \sigma^{(\alpha)}) = 1.$$

In particular, this holds for $\alpha = \text{high}$, so the claim holds. We conclude then that

$$\Pr[\text{ExptIHE}_{\mathcal{B}}(\lambda, 0) = 1] = \Pr[\text{Hyb}_{i,t,4}(\mathcal{A}) = 1].$$

- Suppose the challenger responds according to the specification of $\text{ExptIHE}_{\mathcal{B}}(\lambda, 1)$. In this case, the challenger samples $(\text{hk}_{\text{high}}, \text{vk}_{\text{high}}, \text{td}_{\text{high}}) \leftarrow \text{H.Setup}(1^\lambda, 1^{s_p}, J_{i+1}[1, \dots, t])$. Thus, algorithm \mathcal{B} perfectly simulates an execution of $\text{Hyb}_{i,t,5}$ for \mathcal{A} . Since $\text{Hyb}_{i,t,4}$ and $\text{Hyb}_{i,t,5}$ share *identical* verification conditions, we can appeal to the same argument as before to argue that

$$\text{ExptIHE}_{\mathcal{B}}(\lambda, 1) = 1 \iff \text{Hyb}_{i,t,5}(\mathcal{A}) = 1.$$

Correspondingly,

$$\Pr[\text{ExptIHE}_{\mathcal{B}}(\lambda, 1) = 1] = \Pr[\text{Hyb}_{i,t,5}(\mathcal{A}) = 1].$$

We conclude that \mathcal{B} breaks the one-sided index hiding with extracted guess with the same advantage ϵ . \square

Claim 4.13. *For every $i \in \{0, \dots, d-1\}$, $t \in [|J_{i+1}|]$, it holds that*

$$\Pr[\text{Hyb}_{i,t,6}(\mathcal{A}) = 1] \geq \Pr[\text{Hyb}_{i,t,5}(\mathcal{A}) = 1].$$

Proof. The only difference between $\text{Hyb}_{i,t,5}$ and $\text{Hyb}_{i,t,6}$ is that $\text{Hyb}_{i,t,5}$ performs an additional check that the extracted witness \hat{w} satisfies certain properties. Thus, whenever $\text{Hyb}_{i,t,6}$ outputs 1, hybrid $\text{Hyb}_{i,t,5}$ also outputs 1 and the claim follows. \square

Claim 4.14. *If Π_{BARG} satisfies set hiding against efficient non-uniform adversaries, then there exists a negligible function $\text{negl}(\cdot)$ such that for every $i \in \{0, \dots, d-1\}$ and $t \in [|J_{i+1}|]$, it holds that*

$$|\Pr[\text{Hyb}_{i,t,6}(\mathcal{A}) = 1] - \Pr[\text{Hyb}_{i,t,7}(\mathcal{A}) = 1]| \leq \text{negl}(\lambda).$$

Proof. This follows by a similar argument as the proof of [Claim 4.9](#). \square

Claim 4.15. *For every $i \in \{0, \dots, d-1\}$, it holds that*

$$\Pr[\text{Hyb}_{i,\text{final}}(\mathcal{A}) = 1] \geq \Pr[\text{Hyb}_{i,|J_{i+1}|,7}(\mathcal{A}) = 1].$$

Proof. The only difference between $\text{Hyb}_{i,|J_{i+1}|,7}$ and $\text{Hyb}_{i,\text{final}}$ is that $\text{Hyb}_{i,|J_{i+1}|,7}$ performs an additional check that dig_{low} is Matching. Thus, whenever $\text{Hyb}_{i,\text{final}}$ outputs 1, hybrid $\text{Hyb}_{i,|J_{i+1}|,7}$ also outputs 1 and the claim follows. \square

Claim 4.16. *If Π_{H} satisfies set hiding property against efficient non-uniform adversaries then there exists a negligible function $\text{negl}(\cdot)$ such that for every $i \in \{0, \dots, d-1\}$, it holds that*

$$|\Pr[\text{Hyb}_{i,\text{final}}(\mathcal{A}) = 1] - \Pr[\text{Hyb}_{i+1}(\mathcal{A}) = 1]| \leq \text{negl}(\lambda).$$

Proof. We consider two cases in our analysis:

- Suppose $i = 0$. Then, in hybrid $\text{Hyb}_{0,\text{final}}$, the challenger samples $(\text{hk}_0, \text{vk}_0, \text{td}_0) \leftarrow \text{H.Setup}(1^\lambda, 1^{s_p}, \emptyset)$ and $(\text{hk}_1, \text{vk}_1, \text{td}_1) \leftarrow \text{H.Setup}(1^\lambda, 1^{s_p}, J_1)$. This is identical to how the challenger samples hk_0 and hk_1 in Hyb_1 . Similarly, the verification conditions in the two experiments are identical, so we conclude that

$$\Pr[\text{Hyb}_{0,\text{final}}(\mathcal{A}) = 1] = \Pr[\text{Hyb}_{1}(\mathcal{A}) = 1].$$

- Suppose $i > 0$. Let $\alpha = i \bmod 2$. By construction, in hybrid $\text{Hyb}_{i,\text{final}}$, the challenger samples

$$(\text{hk}_\alpha, \text{vk}_\alpha, \text{td}_\alpha) \leftarrow \text{H.Setup}(1^\lambda, 1^{s_p}, J_i) \quad \text{and} \quad (\text{hk}_{1-\alpha}, \text{vk}_{1-\alpha}, \text{td}_{1-\alpha}) \leftarrow \text{H.Setup}(1^\lambda, 1^{s_p}, J_{i+1}).$$

In hybrid Hyb_{i+1} , the challenger samples

$$(\text{hk}_\alpha, \text{vk}_\alpha, \text{td}_\alpha) \leftarrow \text{H.Setup}(1^\lambda, 1^{s_p}, \emptyset) \quad \text{and} \quad (\text{hk}_{1-\alpha}, \text{vk}_{1-\alpha}, \text{td}_{1-\alpha}) \leftarrow \text{H.Setup}(1^\lambda, 1^{s_p}, J_{i+1}). \quad (4.5)$$

Both experiments check $\text{H.Extract}(\text{td}_{1-\alpha}, \text{dig}_{1-\alpha}) = \text{Matching}$. Thus, the only difference between $\text{Hyb}_{i,\text{final}}$ and Hyb_{i+1} is the distribution of hk_α . In this case, security reduces to the set hiding security of Π_{H} . We give this proof below.

To argue the second case, take any $i \in \{1, \dots, d-1\}$ and suppose that $|\Pr[\text{Hyb}_{i,\text{final}}(\mathcal{A}) = 1] - \Pr[\text{Hyb}_{i+1}(\mathcal{A}) = 1]| = \varepsilon(\lambda)$ for some non-negligible ε . We use \mathcal{A} to build an efficient non-uniform adversary \mathcal{B} that breaks set hiding of Π_H as follows:

- **Preprocessing phase:** On input the security parameter 1^λ , run $(1^n, 1^{s_c}, 1^{s_p}, C, P, \mathbf{x}, \text{st}_{\mathcal{A}}) \leftarrow \mathcal{A}_1(1^\lambda)$. Compute $\tau \leftarrow \text{Preprocess}(C, P, \mathbf{x})$ and output $\text{st}_{\mathcal{B}} = (1^n, 1^{s_c}, 1^{s_p}, \tau, \text{st}_{\mathcal{A}})$.
- **Online phase:** On input the state $\text{st}_{\mathcal{B}} = (1^n, 1^{s_c}, 1^{s_p}, \tau, \text{st}_{\mathcal{A}})$ where $\tau = (C, P, \mathbf{x}, (\beta_1, \dots, \beta_s))$, proceed as follows:
 1. Compute $J_i = \{j \in \text{layer}_i(P) : \beta_j = 0\}$ and $J_{i+1} = \{j \in \text{layer}_{i+1}(P) : \beta_j = 0\}$.
 2. Send 1^{s_p} and the set J_i to the challenger. The challenger replies with hk and vk .
 3. Sample $(\text{crs}_{\text{BARG}}, \text{vk}_{\text{BARG}}) \leftarrow \text{Gen}'(1^\lambda, 1^{s_p}, 1^{n'}, 1^{s'})$, where s' is defined as in [Construction 4.4](#). Let $\alpha = i \bmod 2$. Sample $(\text{hk}_{1-\alpha}, \text{vk}_{1-\alpha}, \text{td}_{1-\alpha}) \leftarrow \text{H.Setup}(1^\lambda, 1^{s_p}, J_{i+1})$ and set $\text{hk}_\alpha \leftarrow \text{hk}$, $\text{vk}_\alpha \leftarrow \text{vk}$.
 4. Let $\text{crs} = (\text{crs}_{\text{BARG}}, \text{hk}_0, \text{hk}_1, \text{vk}_0, \text{vk}_1)$ and run $\pi \leftarrow \mathcal{A}_2(\text{crs}, \text{st}_{\mathcal{A}})$.
 5. Let $\pi = (\text{dig}_0, \text{dig}_1, \pi_{\text{BARG}})$. Let $\hat{\mathbf{x}} = (\hat{x}_1, \dots, \hat{x}_{s_p})$ and C_{aug} be as defined in Prove and Verify in [Construction 4.4](#). Output 1 if all of the following conditions hold (and 0 otherwise):
 - $\text{H.ValidateDigest}(\text{vk}_0, \text{dig}_0) = \text{H.ValidateDigest}(\text{vk}_1, \text{dig}_1) = 1$.
 - $\text{Verify}'(\text{vk}_{\text{BARG}}, C_{\text{aug}}, \hat{\mathbf{x}}, \pi_{\text{BARG}}) = 1$.
 - $\text{H.Extract}(\text{td}_{1-\alpha}, \text{dig}_{1-\alpha}) = \text{Matching}$.

We now consider two possibilities:

- In $\text{ExptSH}_{\mathcal{B}}(\lambda, 0)$, the challenger samples $(\text{hk}, \text{vk}, \text{td}) \leftarrow \text{H.Setup}(1^\lambda, 1^{s_p}, \emptyset)$. In this case, algorithm \mathcal{B} samples crs according to the specification of Hyb_{i+1} . Moreover, algorithm \mathcal{B} computes its output exactly as described in Hyb_{i+1} . This means that \mathcal{B} perfectly simulates $\text{Hyb}_{i+1}(\mathcal{A})$ and thus

$$\Pr[\text{ExptSH}_{\mathcal{B}}(\lambda, 0) = 1] = \Pr[\text{Hyb}_{i+1}(\mathcal{A}) = 1].$$

- In $\text{ExptSH}_{\mathcal{B}}(\lambda, 1)$, the challenger samples $(\text{hk}, \text{vk}, \text{td}) \leftarrow \text{H.Setup}(1^\lambda, 1^{s_p}, J_i)$. In this case, algorithm \mathcal{B} samples crs according to the specification of $\text{Hyb}_{i,\text{final}}$. Moreover, algorithm \mathcal{B} computes the output exactly as described in $\text{Hyb}_{i,\text{final}}$. This means that \mathcal{B} perfectly simulates $\text{Hyb}_{i,\text{final}}(\mathcal{A})$ and thus:

$$\Pr[\text{ExptSH}_{\mathcal{B}}(\lambda, 1) = 1] = \Pr[\text{Hyb}_{i,\text{final}}(\mathcal{A}) = 1].$$

We conclude that algorithm \mathcal{B} breaks the set hiding property with advantage ε in this case and the claim follows. \square

Completing the proof. Combining [Claims 4.8 to 4.16](#), we conclude that there exists a negligible function $\mu(\cdot)$ such that for all $i \in [d]$,

$$\Pr[\text{Hyb}_i(\mathcal{A}) = 1] \geq \Pr[\text{Hyb}_{i-1}(\mathcal{A}) = 1] - O(1) \cdot |J_i| \cdot \mu(\lambda).$$

Moreover, by the same sequence, we conclude that

$$\Pr[\text{Hyb}_{d-1,1,4}(\mathcal{A}) = 1] \geq \Pr[\text{Hyb}_{d-1}(\mathcal{A}) = 1] - O(1) \cdot \mu(\lambda).$$

Putting the pieces together,

$$\Pr[\text{Hyb}_{d-1,1,4}(\mathcal{A}) = 1] \geq \Pr[\text{Hyb}_0(\mathcal{A}) = 1] - O(1) \cdot d \cdot |J_i| \cdot \mu(\lambda) = \Pr[\text{Hyb}_0(\mathcal{A}) = 1] - \text{negl}(\lambda),$$

since $d \cdot |J_i| \leq s = \text{poly}(\lambda)$. Note that we take $\text{Hyb}_{d-1,1,4}$ to be our final hybrid since it imposes the *most* constraints (subsequent hybrids remove requirements from the experiment). To complete the proof we show that for all adversaries \mathcal{A} , $\Pr[\text{Hyb}_{d-1,1,4}(\mathcal{A}) = 1] = \text{negl}(\lambda)$.

Claim 4.17. For all adversaries \mathcal{A} , $\Pr[\text{Hyb}_{d-1,1,4}(\mathcal{A}) = 1] = 0$.

Proof. Fix an adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ and let $(1^n, 1^{s_c}, 1^{s_p}, C, P, \mathbf{x}, \text{st}_{\mathcal{A}})$ be the output of $\mathcal{A}_1(1^\lambda)$. Let $s = |P|$ and d be the depth of P . As usual, for $i \in [k]$, let $\beta_i = 1$ if $(C, x_i) \in \mathcal{L}_{\text{CSAT}}$ and $\beta_i = 0$ otherwise. For $i \in [k+1, s]$, let β_i be the value of wire i in the evaluation of P on $(\beta_1, \dots, \beta_k)$. By construction, all hybrids require that $\beta_s = P(\beta_1, \dots, \beta_k) = 0$, and therefore $J_d[1] = \{\beta_s\}$. However, the conditions for $\widetilde{\text{Hyb}}_{d-1,1,4}$ to output 1 cannot hold simultaneously:

- On the one hand, there must exist a witness $\hat{w} = (b, \sigma^{(0)}, \sigma^{(1)}, w)$ for instance \hat{x}_s (of the relation in Fig. 1) where $b = 0$.
- On the other hand, by definition of instance \hat{x}_s , since s is the output wire, it must be that $b = 1$.

Therefore $\widetilde{\text{Hyb}}_{d-1,1,4}$ is unwinnable. □

Combining Eq. (4.5) and Claim 4.17, we conclude that there exists a negligible function $\text{negl}(\cdot)$ such that

$$0 \geq \Pr[\text{Hyb}_0(\mathcal{A}) = 1] - \text{negl}(\lambda).$$

This means that $\Pr[\text{Hyb}_0(\mathcal{A}) = 1] \leq \text{negl}(\lambda)$, which proves Theorem 4.7. □

5 Generic Construction of Zero-Fixing Hash Functions

In this section, we show how to construct a zero-fixing hash function by combining an index BARG (Definition 2.9), an additively homomorphic encryption scheme with bounded support (Definition 2.1), and a vector encryption scheme with succinct local openings (Definition 2.3).

Binary tree indexing. In the following construction, we will work with complete binary trees. We will use the following procedure to associate a unique index with each node in the binary tree:

Definition 5.1 (Binary Tree Indexing). Let \mathcal{T} be a complete binary tree with $n = 2^k$ leaves. Then \mathcal{T} contains exactly $2n - 1$ nodes. We associate a unique index $i \in [2n - 1]$ via the following procedure:

- First, associate the value $v = 1$ to the root node.
- If v is the value associated with a node, then associate values $2v$ and $2v+1$ with its left and right child. Recursively apply this process to assign a value to every node in the tree.
- The index i associated with a node is defined to be $2n - v$, where v is the value associated with the node.

By design, Definition 5.1 has the following properties:

- The leaf nodes are indexed 1 through n and the root node is indexed $2n - 1$.
- The index of every non-leaf node is greater than the index of its children.
- Given the index of any non-leaf node, we can efficiently compute the indices of its left and right child.

Construction 5.2 (Zero-Fixing Hash Function). Our construction will rely on the following building blocks:

- Let $\Pi'_{\text{BARG}} = (\text{Gen}', \text{Prove}', \text{Verify}', \text{TrapGen}', \text{Extract}')$ be a somewhere extractable *index* BARG (Definition 2.9).
- Let $\Pi_{\text{HE}} = (\text{HE.Gen}, \text{HE.Enc}, \text{HE.Dec}, \text{HE.Add})$ be an additively homomorphic encryption scheme with bounded support (Definition 2.1). For a security parameter λ and a range parameter n , let $\ell_{\text{ct}}(\lambda, n)$ be a bound on the length of the ciphertexts output by either $\text{HE.Enc}(\text{pk}, \cdot)$ or $\text{HE.Add}(\text{pk}, \cdot, \cdot)$ for any (sk, pk) in the support of $\text{HE.Gen}(1^\lambda, 1^n)$.
- Let $\Pi_{\text{Com}} = (\text{Com.Setup}, \text{Com.Commit}, \text{Com.Verify})$ be a vector commitment scheme with succinct local openings (Definition 2.3).

We construct a zero-fixing hash $\Pi_H = (\text{Setup}, \text{Hash}, \text{ProveOpen}, \text{VerOpen}, \text{Extract}, \text{ValidateDigest})$. In the following description, we assume without loss of generality that the bound on the input length $n \in \mathbb{N}$ is a power of two (i.e., $n = 2^k$ for some integer $k \in \mathbb{N}$). Next, we define the following NP relation which we will be using in our construction:

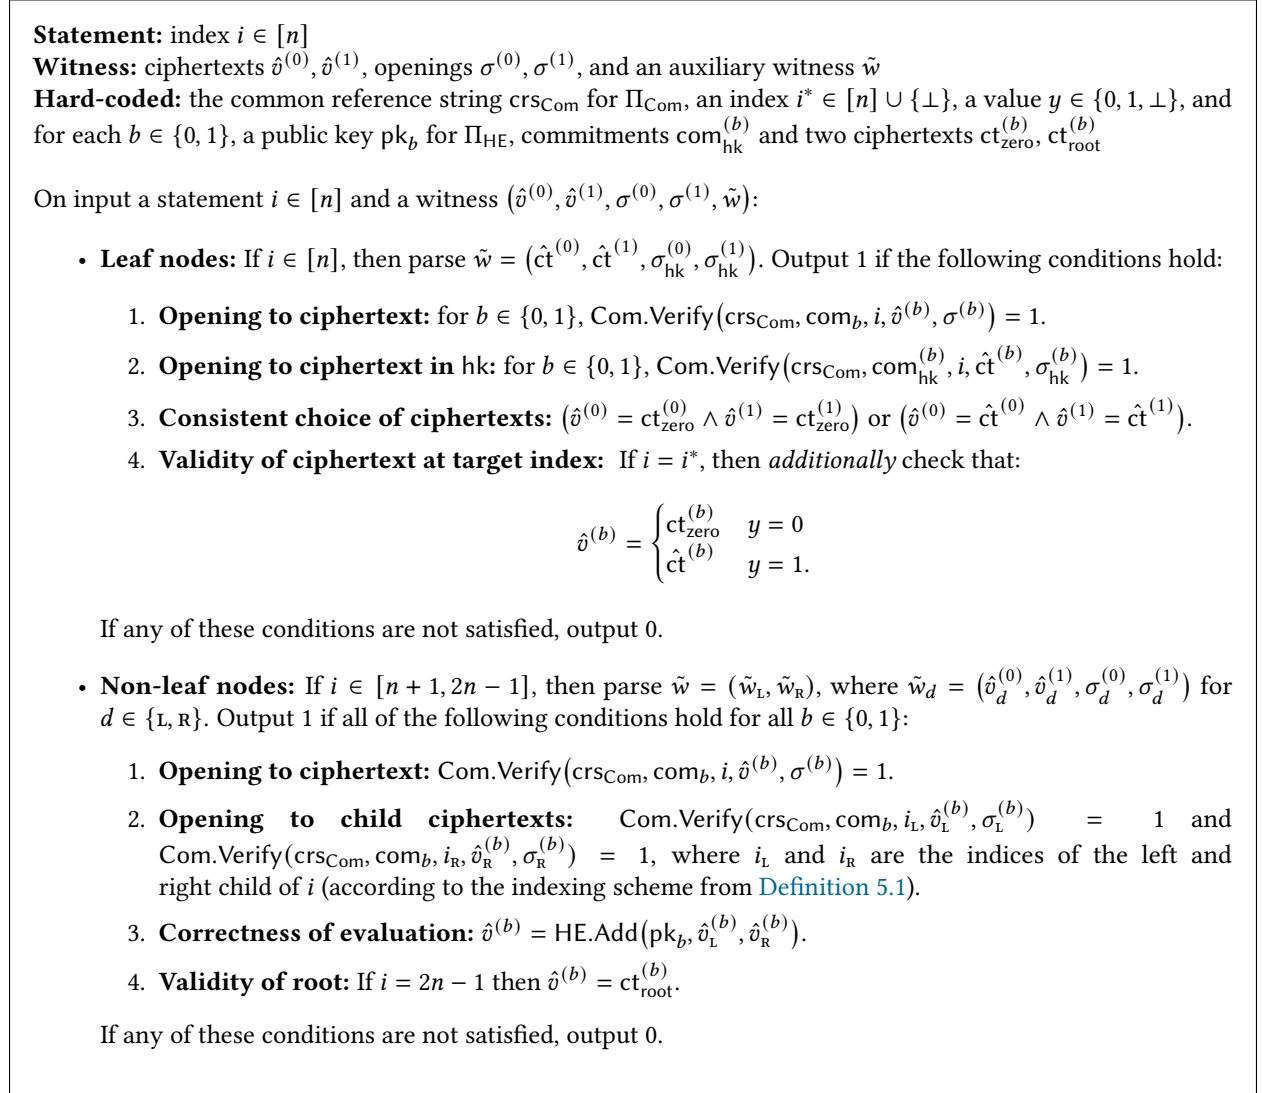


Figure 2: The index relation $\mathcal{R} \left[\text{crs}_{\text{Com}}, \{\text{pk}_b, \text{com}_{\text{hk}}^{(b)}, \text{com}_b, \text{ct}_{\text{zero}}^{(b)}, \text{ct}_{\text{root}}^{(b)}\}_{b \in \{0,1\}}, i^*, y \right]$.

We describe our construction below:

- **Setup**($1^\lambda, 1^n, S$): On input a security parameter λ , the input length $n = 2^k$, and a set $S \subseteq [n]$, the setup algorithm starts by sampling the following:
 - Sample two key pairs: $(\text{sk}_0, \text{pk}_0) \leftarrow \text{HE.Gen}(1^\lambda, 1^n)$ and $(\text{sk}_1, \text{pk}_1) \leftarrow \text{HE.Gen}(1^\lambda, 1^n)$.
 - Sample the CRS for the commitment scheme with block length $\ell_{\text{ct}}(\lambda, n)$ and up to $2n-1$ blocks: $\text{crs}_{\text{Com}} \leftarrow \text{Com.Setup}(1^\lambda, 1^{\ell_{\text{ct}}(\lambda, n)}, 2n-1)$.
 - Sample the CRS for an index BARG: $(\text{crs}_{\text{BARG}}, \text{vk}_{\text{BARG}}) \leftarrow \text{Gen}'(1^\lambda, 1^{2n-1}, 1^s, 1^3)$, where s is a bound on the size of the circuit computing the index relation from [Fig. 2](#). Here, the CRS is extractable on up to 3 positions. Note that since Π_{BARG} is an *index* BARG, Gen' does *not* separately take the statement length as input ([Definition 2.9](#)).

Next, for each $b \in \{0, 1\}$, construct an encryption of 0: $\text{ct}_{\text{zero}}^{(b)} \leftarrow \text{HE.Enc}(\text{pk}_b, 0)$. Next, for each $i \in S$ and $b \in \{0, 1\}$, construct the hash key ciphertexts as follows:

- If $i \in S$, compute $\text{ct}_i^{(b)} \leftarrow \text{HE.Enc}(\text{pk}_b, 1)$.
- If $i \notin S$, compute $\text{ct}_i^{(b)} \leftarrow \text{HE.Enc}(\text{pk}_b, 0)$.

Next, the setup algorithm constructs a commitment to the ciphertexts associated with the hash key. Specifically, for each $b \in \{0, 1\}$, it computes

$$(\text{com}_{\text{hk}}^{(b)}, \sigma_{\text{hk},1}^{(b)}, \dots, \sigma_{\text{hk},n}^{(b)}) \leftarrow \text{Com.Commit}(\text{crs}_{\text{Com}}, (\text{ct}_1^{(b)}, \dots, \text{ct}_n^{(b)})).$$

Finally, the setup algorithm constructs the hash key hk , the verification key vk , and the trapdoor td as follows:

$$\text{hk} = \left(\text{crs}_{\text{Com}}, \text{crs}_{\text{BARG}}, \{ \text{pk}_b, \text{ct}_{\text{zero}}^{(b)}, \text{ct}_1^{(b)}, \dots, \text{ct}_n^{(b)}, \sigma_{\text{hk},1}^{(b)}, \dots, \sigma_{\text{hk},n}^{(b)} \}_{b \in \{0,1\}} \right) \quad (5.1)$$

$$\text{vk} = (\text{crs}_{\text{Com}}, \text{vk}_{\text{BARG}}, \text{pk}_0, \text{pk}_1, \text{ct}_{\text{zero}}^{(0)}, \text{ct}_{\text{zero}}^{(1)}, \text{com}_{\text{hk}}^{(0)}, \text{com}_{\text{hk}}^{(1)}) \quad (5.2)$$

$$\text{td} = (\text{sk}_0, \text{sk}_1). \quad (5.3)$$

- $\text{Hash}(\text{hk}, x)$: On input a hash key hk (parsed as in Eq. (5.1)) and a string $x \in \{0, 1\}^n$, the hashing algorithm proceeds as follows:

- Construct two complete binary trees $\mathcal{T}_0, \mathcal{T}_1$, each with n leaves. For each tree \mathcal{T}_b , we assign a ciphertext $v_i^{(b)}$ to each node $i \in [2s - 1]$ in the tree as follows (where the nodes are indexed using Definition 5.1):
 - * If $i \in [n]$, let $v_i^{(b)} \leftarrow \text{ct}_{\text{zero}}^{(b)}$ if $x_i = 0$ and $v_i^{(b)} \leftarrow \text{ct}_i^{(b)}$ if $x_i = 1$.
 - * For each internal node $i \in [n + 1, 2n - 1]$, compute $v_i^{(b)} \leftarrow \text{HE.Add}(\text{pk}_b, v_{i_L}^{(b)}, v_{i_R}^{(b)})$, where i_L and i_R are the indices associated with the left and right child of node i under the canonical tree indexing scheme (Definition 5.1).
- For $b \in \{0, 1\}$, construct commitments $(\text{com}_b, \sigma_1^{(b)}, \dots, \sigma_{2n-1}^{(b)}) \leftarrow \text{Com.Commit}(\text{crs}_{\text{Com}}, (v_1^{(b)}, \dots, v_{2n-1}^{(b)}))$ to the ciphertexts associated with \mathcal{T}_b .
- For $b \in \{0, 1\}$, let $\text{ct}_{\text{root}}^{(b)} = v_{2n-1}^{(b)}$ (i.e., the ciphertext associated with the root of \mathcal{T}_b). Let C_{\perp} be the circuit that computes the following instantiation of the relation from Fig. 2:

$$\mathcal{R} \left[\text{crs}_{\text{Com}}, \{ \text{pk}_b, \text{com}_{\text{hk}}^{(b)}, \text{com}_b, \text{ct}_{\text{zero}}^{(b)}, \text{ct}_{\text{root}}^{(b)} \}_{b \in \{0,1\}}, \perp, \perp \right].$$

- For each $i \in [2n - 1]$, let $\tau_i = (v_i^{(0)}, v_i^{(1)}, \sigma_i^{(0)}, \sigma_i^{(1)})$ be the opening for the ciphertexts associated with node i in \mathcal{T}_0 and \mathcal{T}_1 . Then, for each $i \in [2s - 1]$, define the auxiliary witness \tilde{w}_i to be
 - * If $i \in [n]$ then $\tilde{w}_i = (\text{ct}_i^{(0)}, \text{ct}_i^{(1)}, \sigma_{\text{hk},i}^{(0)}, \sigma_{\text{hk},i}^{(1)})$.
 - * If $i \in [n + 1, 2n - 1]$ then $\tilde{w}_i = (\tau_{i_L}, \tau_{i_R})$ where i_L, i_R are the indices of the left and right child of node i , respectively.

Finally, let $w_i = (\tau_i, \tilde{w}_i)$ for each $i \in [2n - 1]$. Compute the BARG proof $\pi_{\text{dig}} \leftarrow \text{Prove}'(\text{crs}_{\text{BARG}}, C_{\perp}, 2n - 1, (w_1, \dots, w_{2n-1}))$.

- Output the digest

$$\text{dig} = \left(\text{ct}_{\text{root}}^{(0)}, \text{ct}_{\text{root}}^{(1)}, \text{com}_0, \text{com}_1, \pi_{\text{dig}} \right).$$

- $\text{ProveOpen}(\text{hk}, x, i^*)$: On input a hash key hk (parsed as in Eq. (5.1)), a string $x \in \{0, 1\}^n$ and an index $i^* \in [n]$, the opening algorithm proceeds as follows:

- Let $C_{i^*, x_{i^*}}$ be the circuit that computes the following instantiation of the relation from Fig. 2:

$$\mathcal{R} \left[\text{crs}_{\text{Com}}, \{ \text{pk}_b, \text{com}_{\text{hk}}^{(b)}, \text{com}_b, \text{ct}_{\text{zero}}^{(b)}, \text{ct}_{\text{root}}^{(b)} \}_{b \in \{0,1\}}, i^*, x_{i^*} \right].$$

- Compute the witnesses w_i for each $i \in [2n - 1]$ using the same procedure as in the Hash algorithm.
- Output the opening $\sigma \leftarrow \text{Prove}'(\text{crs}_{\text{BARG}}, C_{i^*, x_{i^*}}, 2n - 1, (w_1, \dots, w_{2n-1}))$
- **VerOpen**(vk, dig, i, b, σ): On input the verification key vk (parsed according to Eq. (5.2)), a digest dig = $(\text{ct}_{\text{root}}^{(0)}, \text{ct}_{\text{root}}^{(1)}, \text{com}_0, \text{com}_1, \pi_{\text{dig}})$, an index $i^* \in [n]$, a bit $b \in \{0, 1\}$ and an opening σ , the verification algorithm outputs $\text{Verify}'(\text{vk}_{\text{BARG}}, C_{i^*, b}, 2n - 1, \sigma)$ where $C_{i^*, b}$ is the circuit computing the following relation from Fig. 2:

$$\mathcal{R}[\text{crs}_{\text{Com}}, \{\text{pk}_b, \text{com}_{\text{hk}}^{(b)}, \text{com}_b, \text{ct}_{\text{zero}}^{(b)}, \text{ct}_{\text{root}}^{(b)}\}_{b \in \{0,1\}}, i^*, b].$$

- **Extract**(td, dig): On input a trapdoor td = $(\text{sk}_0, \text{sk}_1)$ and a digest dig = $(\text{ct}_{\text{root}}^{(0)}, \text{ct}_{\text{root}}^{(1)}, \text{com}_0, \text{com}_1, \pi_{\text{dig}})$, the extraction algorithm outputs Matching if $\text{HE.Dec}(\text{sk}_0, \text{ct}_{\text{root}}^{(0)}) = 0$. Otherwise, the algorithm outputs NotMatching.
- **ValidateDigest**(vk, dig): On input the verification key vk (parsed according to Eq. (5.2)) and a digest dig = $(\text{ct}_{\text{root}}^{(0)}, \text{ct}_{\text{root}}^{(1)}, \text{com}_0, \text{com}_1, \pi_{\text{dig}})$, the digest-validation algorithm outputs $\text{Verify}'(\text{vk}_{\text{BARG}}, C_{\perp}, 2n - 1, \pi_{\text{dig}})$ where C_{\perp} is the circuit computing the following relation from Fig. 2:

$$\mathcal{R}[\text{crs}_{\text{Com}}, \{\text{pk}_b, \text{com}_{\text{hk}}^{(b)}, \text{com}_b, \text{ct}_{\text{zero}}^{(b)}, \text{ct}_{\text{root}}^{(b)}\}_{b \in \{0,1\}}, \perp, \perp].$$

Theorem 5.3 (Correctness). *Suppose Π_{Com} is correct and Π'_{BARG} is complete. Then, Construction 5.2 is correct.*

Proof. Take any $\lambda, n \in \mathbb{N}$ and $x \in \{0, 1\}^n$. Suppose $(\text{hk}, \text{vk}, \text{td}) \leftarrow \text{Setup}(1^\lambda, 1^n, \emptyset)$. Parse

$$\begin{aligned} \text{hk} &= (\text{crs}_{\text{Com}}, \text{crs}_{\text{BARG}}, \{\text{pk}_b, \text{ct}_{\text{zero}}^{(b)}, \text{ct}_1^{(b)}, \dots, \text{ct}_n^{(b)}, \sigma_{\text{hk},1}^{(b)}, \dots, \sigma_{\text{hk},n}^{(b)}\}_{b \in \{0,1\}}) \\ \text{vk} &= (\text{crs}_{\text{Com}}, \text{vk}_{\text{BARG}}, \text{pk}_0, \text{pk}_1, \text{ct}_{\text{zero}}^{(0)}, \text{ct}_{\text{zero}}^{(1)}, \text{com}_{\text{hk}}^{(0)}, \text{com}_{\text{hk}}^{(1)}) \\ \text{td} &= (\text{sk}_0, \text{sk}_1). \end{aligned}$$

We now show each property individually.

Opening correctness. Take any index $i^* \in [n]$ and let $\sigma \leftarrow \text{ProveOpen}(\text{hk}, x, i^*)$. By definition, this means $\sigma \leftarrow \text{Prove}'(\text{crs}_{\text{BARG}}, C_{i^*, x_{i^*}}, 2n - 1, (w_1, \dots, w_{2n-1}))$, where $C_{i^*, x_{i^*}}$ is the circuit that computes the index relation

$$\mathcal{R}[\text{crs}_{\text{Com}}, \{\text{pk}_b, \text{com}_{\text{hk}}^{(b)}, \text{com}_b, \text{ct}_{\text{zero}}^{(b)}, \text{ct}_{\text{root}}^{(b)}\}_{b \in \{0,1\}}, i^*, x_{i^*}]$$

from Fig. 2. By construction of ProveOpen (and by correspondence, Hash), $w_i = (\tau_i, \tilde{w}_i)$ and $\tau_i = (v_i^{(0)}, v_i^{(1)}, \sigma_i^{(0)}, \sigma_i^{(1)})$. We now show that $C_{i^*, x_{i^*}}(i, w_i) = 1$ for all $i \in [2n - 1]$:

- **Leaf nodes:** Suppose $i \in [n]$. Then, $\tilde{w}_i = (\text{ct}_i^{(0)}, \text{ct}_i^{(1)}, \sigma_{\text{hk},i}^{(0)}, \sigma_{\text{hk},i}^{(1)})$. Consider each of the conditions:
 1. **Opening to ciphertext:** By construction of Hash, for $b \in \{0, 1\}$, the commitment com_b is a vector commitment to $(v_1^{(b)}, \dots, v_{2n-1}^{(b)})$, and the opening for position i is $\sigma_i^{(b)}$. This check follows by correctness of the vector commitment scheme.
 2. **Opening to ciphertext in hk:** By construction of Setup, for $b \in \{0, 1\}$, the commitment $\text{com}_{\text{hk}}^{(b)}$ is a vector commitment to $(\text{ct}_1^{(b)}, \dots, \text{ct}_n^{(b)})$ with opening $\sigma_{\text{hk},i}^{(b)}$. This check follows by correctness of the vector commitment scheme.
 3. **Consistent choice of ciphertexts:** By construction of Hash, for $b \in \{0, 1\}$, we have that depending on the value of x_i , either $v_i^{(b)} = \text{ct}_{\text{zero}}^{(b)}$ or $v_i^{(b)} = \text{ct}_i^{(b)}$, and the check passes.
 4. **Validity of ciphertext at target index:** By construction of Hash, $v_{i^*}^{(b)} = \text{ct}_{\text{zero}}^{(b)}$ if $x_{i^*} = 0$ and $v_{i^*}^{(b)} = \text{ct}_i^{(b)}$ if $x_{i^*} = 1$. As such, this check passes.

- **Non-leaf nodes:** Suppose $i \in [n + 1, 2n - 1]$. Then, $\tilde{w}_i = (\tau_{i_L}, \tau_{i_R})$. Consider each of the conditions:
 1. **Opening to ciphertext:** This follows by the same reason as above.
 2. **Opening to child ciphertexts:** This also follows by construction of com_b (namely, com_b is a vector commitment to $(\hat{v}_1^{(b)}, \dots, \hat{v}_{2n-1}^{(b)})$ with openings $\sigma_1^{(b)}, \dots, \sigma_{2n-1}^{(b)}$).
 3. **Correctness of evaluation:** By construction of Hash, for all non-leaf nodes $i \in [n + 1, 2n - 1]$, it holds that $v_i^{(b)} \leftarrow \text{HE.Add}(\text{pk}_b, v_{i_L}^{(b)}, v_{i_R}^{(b)})$, and the checks passes (since HE.Add is deterministic).
 4. **Validity of root:** The Hash algorithm defines $\text{ct}_{\text{root}}^{(b)} = v_{2n-1}^{(b)}$, so this condition is trivially satisfied.

Since $C_{i^*, x_{i^*}}(i, w_i) = 1$ for all $i \in [2n - 1]$, correctness follows by completeness of Π'_{BARG} .

Digest correctness. This follows by an analogous argument as that used to argue opening correctness, with the one difference being the circuit C_{\perp} computes the the index relation

$$\mathcal{R} \left[\text{crs}_{\text{Com}}, \{ \text{pk}_b, \text{com}_{\text{hk}}^{(b)}, \text{com}_b, \text{ct}_{\text{zero}}^{(b)}, \text{ct}_{\text{root}}^{(b)} \}_{b \in \{0,1\}}, \perp, \perp \right].$$

In other words, C_{\perp} (as defined in Hash and in ValidateDigest) does not define a target index i^* or value $b \in \{0, 1\}$, and thus checks a *strict* subset of the conditions as $C_{i^*, x_{i^*}}$ defined in VerOpen. Finally, the witness (w_1, \dots, w_{2n-1}) is defined in an identical manner as before, so all of the required conditions checked by C_{\perp} are satisfied. \square

Theorem 5.4 (Succinctness). *If Π_{HE} is compact and $\Pi_{\text{Com}}, \Pi_{\text{BARG}}$ are succinct, then [Construction 5.2](#) is succinct.*

Proof. Take any $\lambda, n \in \mathbb{N}$ and any $(\text{hk}, \text{vk}, \text{td})$ in the support of $\text{Setup}(1^\lambda, 1^n, \emptyset)$. Take any input $x \in \{0, 1\}^n$ and index $i \in [n]$, and let $\text{dig} \leftarrow \text{Hash}(\text{hk}, x)$, $\pi_{\text{ProveOpen}} \leftarrow \text{ProveOpen}(\text{hk}, x, i)$. Parse

$$\begin{aligned} \text{hk} &= \left(\text{crs}_{\text{Com}}, \text{crs}_{\text{BARG}}, \{ \text{pk}_b, \text{ct}_{\text{zero}}^{(b)}, \text{ct}_1^{(b)}, \dots, \text{ct}_n^{(b)}, \sigma_{\text{hk},1}^{(b)}, \dots, \sigma_{\text{hk},n}^{(b)} \}_{b \in \{0,1\}} \right) \\ \text{vk} &= (\text{crs}_{\text{Com}}, \text{vk}_{\text{BARG}}, \text{pk}_0, \text{pk}_1, \text{ct}_{\text{zero}}^{(0)}, \text{ct}_{\text{zero}}^{(1)}, \text{com}_{\text{hk}}^{(0)}, \text{com}_{\text{hk}}^{(1)}) \\ \text{td} &= (\text{sk}_0, \text{sk}_1) \\ \text{dig} &= (\text{ct}_{\text{root}}^{(0)}, \text{ct}_{\text{root}}^{(1)}, \text{com}_0, \text{com}_1, \pi_{\text{Hash}}) \end{aligned}$$

By compactness of Π_{HE} , the lengths of the public keys pk_0, pk_1 as well as the ciphertexts $\text{ct}_{\text{zero}}^{(b)}, \text{ct}_i^{(b)}$ for all $i \in [n]$ and $b \in \{0, 1\}$ are bounded by $\text{poly}(\lambda + \log n)$. By succinctness of Π_{Com} , it holds that $\text{crs}_{\text{Com}}, \text{com}_{\text{hk}}^{(0)}, \text{com}_{\text{hk}}^{(1)}, \text{com}_0$ and com_1 all have length $\text{poly}(\lambda + \log n)$. Next, let s be a bound on the size of the circuits computing the relation in [Fig. 2](#). The relation in [Fig. 2](#) requires a constant number of opening of ciphertext checks, each of which can be implemented by a circuit of size $\text{poly}(\lambda + \log n)$. Similarly, the correctness of homomorphic evaluation check and the constant number of ciphertext comparisons also require a circuit of size $\text{poly}(\lambda + \log n)$. Thus, the size s of the circuit in [Fig. 2](#) is bounded by $\text{poly}(\lambda + \log n)$. By succinctness of Π_{BARG} , it holds that the length of the verification key vk_{BARG} and the proofs π_{Hash} and $\pi_{\text{ProveOpen}}$ have size $\text{poly}(\lambda + \log n)$. In total, everything is polynomial in $\text{poly}(\lambda + \log n)$ and therefore all of the succinctness requirements ([Definition 3.1](#)) are satisfied by [Construction 5.2](#). \square

Security. In the subsequent sections, we prove each of the required security properties on [Construction 5.2](#). Instantiating the underlying batch argument ([Fact 2.7](#)), the additively homomorphic encryption ([Fact 2.2](#)), and the vector commitment scheme with existing constructions ([Fact 2.4](#)), we obtain the following corollary:

Corollary 5.5 (Zero-Fixing Hash Functions). *Assuming any of (1) the plain LWE assumption, (2) the k -Lin assumption over pairing groups for any constant k , or (3) the (sub-exponential) DDH assumption in pairing-free groups, there exists a zero-fixing hash function.*

[Theorem 1.1](#) now follows in conjunction with our generic construction ([Construction 4.4](#)).

5.1 Security Analysis of Construction 5.2

In this section, we prove that [Construction 5.2](#) satisfies the security requirements on a zero-fixing hash function.

5.1.1 Additive Invariants on Ciphertexts and Predicate Propagation

At a high level, the different security properties of the zero-fixing hash function (zero fixing, extractor validity, and index hiding with extracted guess) will rely on reasoning about various properties on the ciphertext associated with the root node in our tree of ciphertexts (i.e., the hash digest). The analysis of each of these properties follow a similar strategy where we first establish that a certain predicate holds for the ciphertexts in the leaves (i.e., the *honestly-generated* ciphertexts in the hash key). Then, we appeal to the security of the BARG to “propagate” the invariants to the root ciphertext. In this section, we describe a general abstraction for this predicate-propagation strategy that will help unify the analysis of the different security requirements. This construction exploits the fact that the ciphertext tree is perfectly balanced and has depth $\log n$ (where n is the input length); as such, we can rely on a similar type of inductive analysis as that in [\[BBK⁺23\]](#) for arguing soundness of a monotone policy BARG for *log-depth* predicates. We start by formally defining the type of invariants we consider in our security analysis.

Definition 5.6 (Tree-Based Additive Invariant on Ciphertexts). Let n be a power of two and let $\Pi_{\text{HE}} = (\text{Gen}, \text{Enc}, \text{Dec}, \text{Add})$ be a homomorphic encryption scheme. We say that an efficiently-computable predicate $P: \{0, 1\}^* \rightarrow \{0, 1\}$ is a tree-based additive invariant for Π_{HE} if for all $\lambda, n \in \mathbb{N}$, all indices $i^* \in [n] \cup \{\perp\}$, all key-pairs $(\text{sk}_0, \text{pk}_0), (\text{sk}_1, \text{pk}_1)$ in the support of $\text{Gen}(1^\lambda, 1^n)$, all indices $j, j_L, j_R \in [2n - 1]$ where j_L and j_R are the children of j according to the indexing scheme in [Definition 5.1](#), and all ciphertexts $(\text{ct}_L^{(0)}, \text{ct}_L^{(1)}), (\text{ct}_R^{(0)}, \text{ct}_R^{(1)})$ where

$$P(i^*, \text{ct}_L^{(0)}, \text{ct}_L^{(1)}, \text{sk}, \text{sk}', j_L) = 1 \quad \text{and} \quad P(i^*, \text{ct}_R^{(0)}, \text{ct}_R^{(1)}, \text{sk}, \text{sk}', j_R) = 1,$$

it holds that

$$P(i^*, \text{ct}_{\text{sum}}, \text{ct}'_{\text{sum}}, \text{sk}, \text{sk}', j) = 1,$$

where $\text{ct}_{\text{sum}}^{(0)} = \text{Add}(\text{pk}_0, \text{ct}_L^{(0)}, \text{ct}_R^{(0)})$ and $\text{ct}_{\text{sum}}^{(1)} = \text{Add}(\text{pk}_1, \text{ct}_L^{(1)}, \text{ct}_R^{(1)})$. This implies that if P holds for the two children of a node, then it also holds for the parent node.

Predicate propagation experiment. We now define the general predicate propagation experiment we use in the analysis of [Construction 5.2](#). This is a general experiment specification that captures the structure of the security definitions for a zero-fixing hash function.

Definition 5.7 (Predicate Propagation Experiment). The predicate propagation experiment for [Construction 5.2](#) is parameterized by the following two components:

- A tree-based additive invariant P ([Definition 5.6](#)) for the homomorphic encryption scheme Π_{HE} .
- An efficiently-computable “challenge-derivation” function $\text{DeriveChal}(S, i)$ that takes as input a set $S \subseteq [n]$ and an index $i \in [n]$ and outputs two sets $S_0, S_1 \subseteq [n]$ and an index idx that is either a pair (i^*, y^*) or \perp . In the predicate propagation experiment, the sets S_0 and S_1 will determine the distribution of the ciphertexts in the common reference string. The index idx will determine the verification check. Each of the security properties (i.e., zero fixing, extractor validity, and index hiding with extracted guess) will induce a different choice of DeriveChal (to be specified in their respective proofs).

We now define the predicate propagation experiment $\text{Expt}[P, \text{DeriveChal}]$ between a challenger and an adversary \mathcal{A} :

1. On input the security parameter 1^λ , algorithm \mathcal{A} outputs the input length 1^n , a set $S \subseteq [n]$, and an index $i^* \in S$ (or a special symbol \perp).
2. The challenger computes $(S_0, S_1, \text{idx}) \leftarrow \text{DeriveChal}(S, i^*)$.
3. The challenger now samples the following quantities as in Setup:

- Sample $(sk_0, pk_0) \leftarrow \text{HE.Gen}(1^\lambda, 1^n)$ and $(sk_1, pk_1) \leftarrow \text{HE.Gen}(1^\lambda, 1^n)$.
- Sample $\text{crs}_{\text{Com}} \leftarrow \text{Com.Setup}(1^\lambda, 1^{\ell_{\text{ct}}(\lambda, n)}, 2n - 1)$.
- Sample $(\text{crs}_{\text{BARG}}, \text{vk}_{\text{BARG}}) \leftarrow \text{Gen}'(1^\lambda, 1^{2n-1}, 1^s, 1^3)$, where s is a bound on the size of the circuit computing the index relation from Fig. 2.
- For each $b \in \{0, 1\}$, sample $\text{ct}_{\text{zero}}^{(b)} \leftarrow \text{HE.Enc}(pk_b, 0)$. Then, for each $i \in [n]$ and $b \in \{0, 1\}$, if $i \in S_b$, sample $\text{ct}_i^{(b)} \leftarrow \text{HE.Enc}(pk_b, 1)$; otherwise, if $i \notin S_b$, sample $\text{ct}_i^{(b)} \leftarrow \text{HE.Enc}(pk_b, 0)$.
- For each $b \in \{0, 1\}$, let $(\text{com}_{\text{hk}}^{(b)}, \sigma_{\text{hk},1}^{(b)}, \dots, \sigma_{\text{hk},n}^{(b)}) \leftarrow \text{Com.Commit}(\text{crs}_{\text{Com}}, (\text{ct}_1^{(b)}, \dots, \text{ct}_n^{(b)}))$

4. The challenger constructs hk and vk according to Eqs. (5.1) and (5.2):

$$\begin{aligned} \text{hk} &= \left(\text{crs}_{\text{Com}}, \text{crs}_{\text{BARG}}, \left\{ pk_b, \text{ct}_{\text{zero}}^{(b)}, \text{ct}_1^{(b)}, \dots, \text{ct}_n^{(b)}, \sigma_{\text{hk},1}^{(b)}, \dots, \sigma_{\text{hk},n}^{(b)} \right\}_{b \in \{0,1\}} \right) \\ \text{vk} &= (\text{crs}_{\text{Com}}, \text{vk}_{\text{BARG}}, pk_0, pk_1, \text{ct}_{\text{zero}}^{(0)}, \text{ct}_{\text{zero}}^{(1)}, \text{com}_{\text{hk}}^{(0)}, \text{com}_{\text{hk}}^{(1)}) \end{aligned}$$

The challenger gives (hk, vk) to \mathcal{A} .

5. Algorithm \mathcal{A} outputs a digest $\text{dig} = (\text{ct}_{\text{root}}^{(0)}, \text{ct}_{\text{root}}^{(1)}, \text{com}_0, \text{com}_1, \pi_{\text{dig}})$ and a proof π .
6. The output of the experiment is 1 if

$$\text{Verify}'(\text{vk}_{\text{BARG}}, C_{\text{idx}}, 2n - 1, \pi) = 1 \quad \text{and} \quad P(i^*, \text{ct}_{\text{root}}^{(0)}, \text{ct}_{\text{root}}^{(1)}, sk_0, sk_1, 2n - 1) = 0.$$

Otherwise, the output is 0. Here, the circuit C_{idx} computes the relation from Fig. 2:

- If $\text{idx} = (i, y)$, then C_{idx} computes the relation $\mathcal{R}[\text{crs}_{\text{Com}}, \{pk_b, \text{com}_{\text{hk}}^{(b)}, \text{com}_b, \text{ct}_{\text{zero}}^{(b)}, \text{ct}_{\text{root}}^{(b)}\}_{b \in \{0,1\}}, i, y]$ as in `VerOpen`.
- If $\text{idx} = \perp$, then C_{idx} computes the relation $\mathcal{R}[\text{crs}_{\text{Com}}, \{pk_b, \text{com}_{\text{hk}}^{(b)}, \text{com}_b, \text{ct}_{\text{zero}}^{(b)}, \text{ct}_{\text{root}}^{(b)}\}_{b \in \{0,1\}}, \perp, \perp]$ as in `ValidateDigest`.

In words, the adversary “wins” the game if it produces a proof π that verifies, but the digest does *not* satisfy the tree-based additive invariant P .

Proof strategy. As we show in the subsequent sections (Sections 5.1.2 to 5.1.5), most of the security properties for the zero-fixing hash function (zero fixing, extractor validity, and index hiding with extracted guess) are a special case of the general predicate propagation experiment (with a suitable choice of the tree-based additive invariant P and the challenge-derivation function `DeriveChal`). Our goal below is to show that if specific “pre-conditions” hold, then for all efficient adversaries \mathcal{A} , the probability that $\text{Expt}[P, \text{DeriveChal}]$ outputs 1 is negligible. In turn, this will imply the desired security properties on the zero-fixing hash function.

Predicate propagation hybrid experiment. The proof π the adversary outputs is a BARG on $2n - 1$ statements. We can associate these $2n - 1$ statements with the nodes of a complete binary tree with n leaves. For each $j \in [2n - 1]$, we now define an intermediate predicate propagation experiment $\text{Expt}_j[P, \text{DeriveChal}]$ where instead of checking the tree-based additive invariant holds for the values $\text{ct}_{\text{root}}^{(0)}, \text{ct}_{\text{root}}^{(1)}$ from the digest dig , the challenger instead checks the invariant for the value associated with node j in the tree obtained by *extracting* a witness from the BARG. In the subsequent analysis (Theorem 5.9), we show (inductively) that if the invariant holds for the values extracted from the children of a node j , then it also holds for the values extracted from node j itself. In this way, if the invariant P holds for all the values associated with the leaves of the tree, then the invariant also holds for the values associated with the root of the tree. Finally, the relation in Fig. 2 from Construction 5.2 enforces that the adversarially-chosen values $\text{ct}_{\text{root}}^{(0)}, \text{ct}_{\text{root}}^{(1)}$ are consistent with the values that would be extracted from the root of the tree. This allows us to reason about properties of the adversarially-chosen values $\text{ct}_{\text{root}}^{(0)}, \text{ct}_{\text{root}}^{(1)}$. We define this experiment below and then state the main predicate propagation theorem (Theorem 5.9).

Definition 5.8 (Predicate Propagation Hybrid Experiment). Let $j \in \mathbb{N}$ be an index. For a tree-based additive invariant P and a challenge-derivation function DeriveChal , we define the predicate propagation hybrid experiment $\text{Expt}_j[P, \text{DeriveChal}]$ between a challenger and an adversary \mathcal{A} as follows:

1. On input the security parameter 1^λ , algorithm \mathcal{A} outputs the input length 1^n , a set $S \subseteq [n]$, and an index $i^* \in S$.
2. The challenger computes $(S_0, S_1, \text{idx}) \leftarrow \text{DeriveChal}(S, i^*)$.
3. The challenger now samples the following quantities as in Setup:
 - Sample $(\text{sk}_0, \text{pk}_0) \leftarrow \text{HE.Gen}(1^\lambda, 1^n)$ and $(\text{sk}_1, \text{pk}_1) \leftarrow \text{HE.Gen}(1^\lambda, 1^n)$.
 - Sample $\text{crs}_{\text{Com}} \leftarrow \text{Com.Setup}(1^\lambda, 1^{\ell_{\text{ct}}(\lambda, n)}, 2n - 1)$.
 - Sample $(\text{crs}_{\text{BARG}}, \text{vk}_{\text{BARG}}, \text{td}_{\text{BARG}}) \leftarrow \text{TrapGen}'(1^\lambda, 1^{2n-1}, 1^s, 1^3, \{j\})$, where s is a bound on the size of the circuit computing the index relation from Fig. 2.
 - For each $b \in \{0, 1\}$, sample $\text{ct}_{\text{zero}}^{(b)} \leftarrow \text{HE.Enc}(\text{pk}_b, 0)$. Then, for each $i \in [n]$ and $b \in \{0, 1\}$, if $i \in S_b$, sample $\text{ct}_i^{(b)} \leftarrow \text{HE.Enc}(\text{pk}_b, 1)$; otherwise, if $i \notin S_b$, sample $\text{ct}_i^{(b)} \leftarrow \text{HE.Enc}(\text{pk}_b, 0)$.
 - For each $b \in \{0, 1\}$, let $(\text{com}_{\text{hk}}^{(b)}, \sigma_{\text{hk},1}^{(b)}, \dots, \sigma_{\text{hk},n}^{(b)}) \leftarrow \text{Com.Commit}(\text{crs}_{\text{Com}}, (\text{ct}_1^{(b)}, \dots, \text{ct}_n^{(b)}))$
4. The challenger constructs hk and vk according to Eqs. (5.1) and (5.2):

$$\begin{aligned} \text{hk} &= \left(\text{crs}_{\text{Com}}, \text{crs}_{\text{BARG}}, \{ \text{pk}_b, \text{ct}_{\text{zero}}^{(b)}, \text{ct}_1^{(b)}, \dots, \text{ct}_n^{(b)}, \sigma_{\text{hk},1}^{(b)}, \dots, \sigma_{\text{hk},n}^{(b)} \}_{b \in \{0,1\}} \right) \\ \text{vk} &= \left(\text{crs}_{\text{Com}}, \text{vk}_{\text{BARG}}, \text{pk}_0, \text{pk}_1, \text{ct}_{\text{zero}}^{(0)}, \text{ct}_{\text{zero}}^{(1)}, \text{com}_{\text{hk}}^{(0)}, \text{com}_{\text{hk}}^{(1)} \right) \end{aligned}$$

The challenger gives (hk, vk) to \mathcal{A} .

5. Algorithm \mathcal{A} outputs a digest $\text{dig} = (\text{ct}_{\text{root}}^{(0)}, \text{ct}_{\text{root}}^{(1)}, \text{com}_0, \text{com}_1, \pi_{\text{dig}})$ and a proof π .
6. The challenger computes $(\hat{\nu}_j^{(0)}, \hat{\nu}_j^{(1)}, \sigma_j^{(0)}, \sigma_j^{(1)}, \tilde{w}_j) \leftarrow \text{Extract}'(\text{td}_{\text{BARG}}, \pi, j)$.
7. The output of the experiment is 1 if the following conditions hold:
 - $\text{Verify}'(\text{vk}_{\text{BARG}}, C_{\text{idx}}, 2n - 1, \pi) = 1$.
 - $C_{\text{idx}}(j, (\hat{\nu}_j^{(0)}, \hat{\nu}_j^{(1)}, \sigma_j^{(0)}, \sigma_j^{(1)}, \tilde{w}_j)) = 1$.
 - $P(i^*, \hat{\nu}_j^{(0)}, \hat{\nu}_j^{(1)}, \text{sk}_0, \text{sk}_1, j) = 0$.

Otherwise, the output is 0. As in Definition 5.7, the circuit C_{idx} computes the relation from Fig. 2:

- If $\text{idx} = (i, y)$, then C_{idx} computes the relation $\mathcal{R}[\text{crs}_{\text{Com}}, \{ \text{pk}_b, \text{com}_{\text{hk}}^{(b)}, \text{com}_b, \text{ct}_{\text{zero}}^{(b)}, \text{ct}_{\text{root}}^{(b)} \}_{b \in \{0,1\}}, i, y]$.
- If $\text{idx} = \perp$, then C_{idx} computes the relation $\mathcal{R}[\text{crs}_{\text{Com}}, \{ \text{pk}_b, \text{com}_{\text{hk}}^{(b)}, \text{com}_b, \text{ct}_{\text{zero}}^{(b)}, \text{ct}_{\text{root}}^{(b)} \}_{b \in \{0,1\}}, \perp, \perp]$.

Otherwise, the output is 0.

Theorem 5.9 (Predicate Propagation). *Let P be a tree-based additive invariant and let DeriveChal be a challenge-derivation function. Suppose Π_{Com} satisfies computational binding and Π_{BARG} satisfies set hiding with extraction, set hiding, and somewhere extractability. Let \mathcal{A} be any efficient adversary for the predicate propagation experiment. Suppose that for every index $j \in [n]$ (where $n = n(\lambda)$ is the input length chosen by \mathcal{A}), there exists a negligible function $\varepsilon_j(\cdot)$ such that*

$$\Pr[\text{Expt}_j[P, \text{DeriveChal}](\mathcal{A}) = 1] = \varepsilon_j(\lambda).$$

Then there exists a negligible function $\text{negl}(\cdot)$ such that

$$\Pr[\text{Expt}[P, \text{DeriveChal}](\mathcal{A}) = 1] = \text{negl}(\lambda).$$

Proof. To simplify notation, we write $\text{Expt} := \text{Expt}[P, \text{DeriveChal}]$ and $\text{Expt}_j := \text{Expt}_j[P, \text{DeriveChal}]$ in the following proof. Fix an adversary \mathcal{A} and let n be the input length chosen by \mathcal{A} . We proceed by induction on the index $j \in [2n-1]$. In the following, we will view the index j as an index of a node in a (complete) binary tree with n leaves (indexed according to [Definition 5.1](#)). As such, we can refer to the “height” of an index j . Then, we show the following lemma:

Lemma 5.10. *Suppose the conditions of [Theorem 5.9](#) hold. Take any index $j \in [2n-1]$ and let h be the height of node j (where the leaf nodes have height 0). Then, there exists a negligible function $\varepsilon_j(\lambda)$ such that*

$$\Pr[\text{Expt}_j(\mathcal{A}) = 1] = 2^h \cdot \varepsilon_j(\lambda).$$

Proof. Suppose the conditions of [Theorem 5.9](#) hold. We prove the lemma by induction on the height h of the index $j \in [2n-1]$.

Base case. For the indices $j \in [n]$ of height 0 (i.e., the leaves of the tree), the lemma follows by assumption.

Inductive step. Suppose the inductive hypothesis holds for every index $j' \in [2n-1]$ of height h . Let $j \in [2n-1]$ be an index with height $h+1$. Let $j_L, j_R \in [2n-1]$ be the indices of the left and right child of node j (as defined in [Definition 5.1](#)). By construction, j_L and j_R have height h . The inductive hypothesis now asserts that for $j^* \in \{j_L, j_R\}$,

$$\Pr[\text{Expt}_{j^*}(\mathcal{A}) = 1] = 2^h \cdot \varepsilon_{j^*}(\lambda), \quad (5.4)$$

for some negligible function $\varepsilon_{j^*}(\lambda)$. We now define an intermediate experiment Expt'_j for each node j of height $h > 0$:

1. On input the security parameter 1^λ , algorithm \mathcal{A} outputs the input length 1^n , a set $S \subseteq [n]$, and an index $i^* \in S$.
2. The challenger computes $(S_0, S_1, \text{idX}) \leftarrow \text{DeriveChal}(S, i^*)$.
3. The challenger now samples the following quantities as in Setup:
 - Sample $(\text{sk}_0, \text{pk}_0) \leftarrow \text{HE.Gen}(1^\lambda, 1^n)$ and $(\text{sk}_1, \text{pk}_1) \leftarrow \text{HE.Gen}(1^\lambda, 1^n)$.
 - Sample $\text{crs}_{\text{Com}} \leftarrow \text{Com.Setup}(1^\lambda, 1^{\text{ct}(\lambda, n)}, 2n-1)$.
 - Sample $(\text{crs}_{\text{BARG}}, \text{vk}_{\text{BARG}}, \text{td}_{\text{BARG}}) \leftarrow \text{TrapGen}'(1^\lambda, 1^{2n-1}, 1^s, 1^3, \{j, j_L, j_R\})$, where s is a bound on the size of the circuit computing the index relation from [Fig. 2](#).
 - For each $b \in \{0, 1\}$, sample $\text{ct}_{\text{zero}}^{(b)} \leftarrow \text{HE.Enc}(\text{pk}_b, 0)$. Then, for each $i \in [n]$ and $b \in \{0, 1\}$, if $i \in S_b$, sample $\text{ct}_i^{(b)} \leftarrow \text{HE.Enc}(\text{pk}_b, 1)$; otherwise, if $i \notin S_b$, sample $\text{ct}_i^{(b)} \leftarrow \text{HE.Enc}(\text{pk}_b, 0)$.
 - For each $b \in \{0, 1\}$, let $(\text{com}_{\text{hk}}^{(b)}, \sigma_{\text{hk},1}^{(b)}, \dots, \sigma_{\text{hk},n}^{(b)}) \leftarrow \text{Com.Commit}(\text{crs}_{\text{Com}}, (\text{ct}_1^{(b)}, \dots, \text{ct}_n^{(b)}))$
4. The challenger constructs hk and vk according to [Eqs. \(5.1\) and \(5.2\)](#):

$$\begin{aligned} \text{hk} &= \left(\text{crs}_{\text{Com}}, \text{crs}_{\text{BARG}}, \left\{ \text{pk}_b, \text{ct}_{\text{zero}}^{(b)}, \text{ct}_1^{(b)}, \dots, \text{ct}_n^{(b)}, \sigma_{\text{hk},1}^{(b)}, \dots, \sigma_{\text{hk},n}^{(b)} \right\}_{b \in \{0,1\}} \right) \\ \text{vk} &= (\text{crs}_{\text{Com}}, \text{vk}_{\text{BARG}}, \text{pk}_0, \text{pk}_1, \text{ct}_{\text{zero}}^{(0)}, \text{ct}_{\text{zero}}^{(1)}, \text{com}_{\text{hk}}^{(0)}, \text{com}_{\text{hk}}^{(1)}) \end{aligned}$$

The challenger gives (hk, vk) to \mathcal{A} .

5. Algorithm \mathcal{A} outputs a digest $\text{dig} = (\text{ct}_{\text{root}}^{(0)}, \text{ct}_{\text{root}}^{(1)}, \text{com}_0, \text{com}_1, \pi_{\text{dig}})$ and a proof π .
6. The challenger computes $(\hat{v}_j^{(0)}, \hat{v}_j^{(1)}, \sigma_j^{(0)}, \sigma_j^{(1)}, \tilde{w}_j) \leftarrow \text{Extract}'(\text{td}_{\text{BARG}}, \pi, j)$.
7. The output of the experiment is 1 if the following conditions hold:
 - $\text{Verify}'(\text{vk}_{\text{BARG}}, C_{\text{idX}}, 2n-1, \pi) = 1$.
 - $C_{\text{idX}}(j, (\hat{v}_j^{(0)}, \hat{v}_j^{(1)}, \sigma_j^{(0)}, \sigma_j^{(1)}, \tilde{w}_j)) = 1$.

- $P(i^*, \hat{v}_j^{(0)}, \hat{v}_j^{(1)}, \text{sk}_0, \text{sk}_1, j) = 0$.

Otherwise, the output is 0.

In our analysis below, we define an additional set of events in an execution of Expt'_j with \mathcal{A} . First, define the following two quantities:

- $(\hat{v}_{j_L}^{(0)}, \hat{v}_{j_L}^{(1)}, \sigma_{j_L}^{(0)}, \sigma_{j_L}^{(1)}, \tilde{w}_{j_L}) \leftarrow \text{Extract}'(\text{td}_{\text{BARG}}, \pi, j_L)$.
- $(\hat{v}_{j_R}^{(0)}, \hat{v}_{j_R}^{(1)}, \sigma_{j_R}^{(0)}, \sigma_{j_R}^{(1)}, \tilde{w}_{j_R}) \leftarrow \text{Extract}'(\text{td}_{\text{BARG}}, \pi, j_R)$.

Now, define the following events:

- $E_{\text{Verify}}^{(j)}$: This is the event that $\text{Verify}'(\text{vk}_{\text{BARG}}, C_{\text{idx}}, 2n - 1, \pi) = 1$.
- $E_{P, j^*}^{(j)}$ for each $j^* \in \{j, j_L, j_R\}$: This is the event where $P(i^*, \hat{v}_{j^*}^{(0)}, \hat{v}_{j^*}^{(1)}, \text{sk}_0, \text{sk}_1, j^*) = 1$.
- $E_{\text{ValidCom}, j^*}^{(j)}$ for each $j^* \in \{j_L, j_R\}$: This is the event

$$\text{Com.Verify}(\text{crs}_{\text{Com}}, \text{com}_0, j^*, \hat{v}_{j^*}^{(0)}, \sigma_{j^*}^{(0)}) = 1 = \text{Com.Verify}(\text{crs}_{\text{Com}}, \text{com}_1, j^*, \hat{v}_{j^*}^{(1)}, \sigma_{j^*}^{(1)}).$$

- $E_{\text{SAT}, j^*}^{(j)}$ for each $j^* \in \{j, j_L, j_R\}$: This is the event $C_{\text{idx}}(j^*, (\hat{v}_{j^*}^{(0)}, \hat{v}_{j^*}^{(1)}, \sigma_{j^*}^{(0)}, \sigma_{j^*}^{(1)}, \tilde{w}_{j^*})) = 1$.

We now relate the probability that $\text{Expt}_j(\mathcal{A})$ outputs 1 to the probability that $\text{Expt}_{j_L}(\mathcal{A})$ and $\text{Expt}_{j_R}(\mathcal{A})$ outputs 1. To do so, we first program the BARG to be extracting on the set $\{j, j_L, j_R\}$. We then argue via somewhere extractability of the BARG and computational binding of the commitment scheme that if the values associated with the nodes j_L and j_R satisfy the predicate P and the proof verifies, then the value associated with j must also satisfy the predicate P . In this case, the output of $\text{Expt}_j(\mathcal{A})$ is guaranteed to be 0.

Claim 5.11. *If Π_{BARG} satisfies set hiding with extraction, then there exists a negligible function $\text{negl}(\cdot)$ such that for all $j^* \in \{j, j_L, j_R\}$, it holds that*

$$\left| \Pr[\text{Expt}_{j^*}(\mathcal{A}) = 1] - \Pr[E_{\text{Verify}}^{(j)} \wedge E_{\text{SAT}, j^*}^{(j)} \wedge \neg E_{P, j^*}^{(j)}] \right| = \text{negl}(\lambda).$$

Proof. Take any $j^* \in \{j, j_L, j_R\}$ and suppose

$$\left| \Pr[\text{Expt}_{j^*}(\mathcal{A}) = 1] - \Pr[E_{\text{Verify}}^{(j)} \wedge E_{\text{SAT}, j^*}^{(j)} \wedge \neg E_{P, j^*}^{(j)}] \right| \geq \varepsilon(\lambda),$$

for some non-negligible ε . Importantly, note that the events $E_{\text{Verify}}^{(j)}$, $E_{\text{SAT}, j^*}^{(j)}$, and $E_{P, j^*}^{(j)}$ are defined for Expt'_j and not Expt_{j^*} . We use \mathcal{A} to construct an adversary \mathcal{B} for the set hiding with extraction game of Π_{BARG} :

1. On input the security parameter 1^λ , algorithm \mathcal{B} runs algorithm \mathcal{A} to obtain the input length 1^n , the set $S \subseteq [n]$, and an index $i^* \in S$.
2. Algorithm \mathcal{B} outputs 1^{2n-1} , 1^s , 1^3 , the challenge set $J = \{j, j_L, j_R\}$, and the challenge index $j^* \in J$ to the challenger, where s is the bound on the size of the circuit in Fig. 2. The challenger responds with $(\text{crs}_{\text{BARG}}, \text{vk}_{\text{BARG}})$.
3. Algorithm \mathcal{B} computes $(S_0, S_1, \text{idx}) \leftarrow \text{DeriveChal}(S, i^*)$. It then samples the following components:
 - Sample $(\text{sk}_0, \text{pk}_0) \leftarrow \text{HE.Gen}(1^\lambda, 1^n)$ and $(\text{sk}_1, \text{pk}_1) \leftarrow \text{HE.Gen}(1^\lambda, 1^n)$.
 - Sample $\text{crs}_{\text{Com}} \leftarrow \text{Com.Setup}(1^\lambda, 1^{\ell_{\text{ct}}(\lambda, n)}, 2n - 1)$.
 - For each $b \in \{0, 1\}$, sample $\text{ct}_{\text{zero}}^{(b)} \leftarrow \text{HE.Enc}(\text{pk}_b, 0)$. Then, for each $i \in [n]$ and $b \in \{0, 1\}$, if $i \in S_b$, sample $\text{ct}_i^{(b)} \leftarrow \text{HE.Enc}(\text{pk}_b, 1)$; otherwise, if $i \notin S_b$, sample $\text{ct}_i^{(b)} \leftarrow \text{HE.Enc}(\text{pk}_b, 0)$.

- For each $b \in \{0, 1\}$, let $(\text{com}_{\text{hk}}^{(b)}, \sigma_{\text{hk},1}^{(b)}, \dots, \sigma_{\text{hk},n}^{(b)}) \leftarrow \text{Com.Commit}(\text{crs}_{\text{Com}}, (\text{ct}_1^{(b)}, \dots, \text{ct}_n^{(b)}))$

4. Algorithm \mathcal{B} constructs hk and vk according to Eqs. (5.1) and (5.2):

$$\begin{aligned} \text{hk} &= \left(\text{crs}_{\text{Com}}, \text{crs}_{\text{BARG}}, \{ \text{pk}_b, \text{ct}_{\text{zero}}^{(b)}, \text{ct}_1^{(b)}, \dots, \text{ct}_n^{(b)}, \sigma_{\text{hk},1}^{(b)}, \dots, \sigma_{\text{hk},n}^{(b)} \}_{b \in \{0,1\}} \right) \\ \text{vk} &= (\text{crs}_{\text{Com}}, \text{vk}_{\text{BARG}}, \text{pk}_0, \text{pk}_1, \text{ct}_{\text{zero}}^{(0)}, \text{ct}_{\text{zero}}^{(1)}, \text{com}_{\text{hk}}^{(0)}, \text{com}_{\text{hk}}^{(1)}) \end{aligned}$$

Algorithm \mathcal{B} gives (hk, vk) to \mathcal{A} .

5. Algorithm \mathcal{A} outputs a digest $\text{dig} = (\text{ct}_{\text{root}}^{(0)}, \text{ct}_{\text{root}}^{(1)}, \text{com}_0, \text{com}_1, \pi_{\text{dig}})$ and a proof π .

6. Let C_{idx} be the circuit as defined in Definition 5.7. Algorithm \mathcal{B} first checks

$$\text{Verify}'(\text{vk}_{\text{BARG}}, C_{\text{idx}}, 2n-1, \pi) = 1.$$

If the check fails, algorithm \mathcal{B} aborts with output \perp . Otherwise, algorithm \mathcal{B} sends the circuit C_{idx} , the instance number $2n-1$, and the proof π to the challenger. The challenger replies with a string which \mathcal{B} parses as $(\hat{v}_{j^*}^{(0)}, \hat{v}_{j^*}^{(1)}, \sigma_{j^*}^{(0)}, \sigma_{j^*}^{(1)}, \tilde{w}_{j^*})$.

7. Algorithm \mathcal{B} outputs 1 if

$$C_{\text{idx}}(j^*, (\hat{v}_{j^*}^{(0)}, \hat{v}_{j^*}^{(1)}, \sigma_{j^*}^{(0)}, \sigma_{j^*}^{(1)}, \tilde{w}_{j^*})) = 1 \quad \text{and} \quad P(i^*, \hat{v}_{j^*}^{(0)}, \hat{v}_{j^*}^{(1)}, \text{sk}_0, \text{sk}_1, j) = 0.$$

Otherwise, algorithm \mathcal{B} outputs 0.

Let $(\text{crs}_{\text{BARG}}, \text{vk}_{\text{BARG}}, \text{td}_{\text{BARG}})$ be the parameters sampled by the challenger in the set hiding with extraction game. In the game, after \mathcal{B} outputs $(C_{\text{idx}}, 2n-1, \pi)$, the challenger checks $\text{Verify}'(\text{vk}_{\text{BARG}}, C_{\text{idx}}, 2n-1, \pi) = 1$. If the check passes, it replies with $(\hat{v}_{j^*}^{(0)}, \hat{v}_{j^*}^{(1)}, \sigma_{j^*}^{(0)}, \sigma_{j^*}^{(1)}, \tilde{w}_{j^*}) \leftarrow \text{Extract}'(\text{td}_{\text{BARG}}, \pi, j^*)$. We now consider the two possibilities:

- Suppose the challenger responds according to the specification of $\text{ExptIHE}_{\mathcal{A}}(\lambda, 0)$. In this case, the challenger samples $(\text{crs}_{\text{BARG}}, \text{vk}_{\text{BARG}}, \text{td}_{\text{BARG}}) \leftarrow \text{TrapGen}'(1^\lambda, 1^{2n-1}, 1^s, 1^3, \{j, j_L, j_R\})$. Thus, algorithm \mathcal{B} perfectly simulates for \mathcal{A} an execution of Expt'_j . We claim that algorithm \mathcal{B} outputs 1 if and only if the event $E_{\text{Verify}}^{(j)} \wedge E_{\text{SAT}, j^*}^{(j)} \wedge \neg E_{P, j^*}^{(j)}$ occurs. This event corresponds to the following set of conditions:

$$\text{Verify}'(\text{vk}_{\text{BARG}}, C_{\text{idx}}, 2n-1, \pi) = 1 \text{ and } C_{\text{idx}}(j^*, (\hat{v}_{j^*}^{(0)}, \hat{v}_{j^*}^{(1)}, \sigma_{j^*}^{(0)}, \sigma_{j^*}^{(1)}, \tilde{w}_{j^*})) = 1 \text{ and } P(i^*, \hat{v}_{j^*}^{(0)}, \hat{v}_{j^*}^{(1)}, \text{sk}_0, \text{sk}_1, j^*) = 0.$$

where $(\hat{v}_{j^*}^{(0)}, \hat{v}_{j^*}^{(1)}, \sigma_{j^*}^{(0)}, \sigma_{j^*}^{(1)}, \tilde{w}_{j^*}) \leftarrow \text{Extract}'(\text{td}_{\text{BARG}}, \pi, j^*)$. This is the same set of conditions that algorithm \mathcal{B} checks, so algorithm \mathcal{B} outputs 1 with probability $\Pr[E_{\text{Verify}}^{(j)} \wedge E_{\text{SAT}, j^*}^{(j)} \wedge \neg E_{P, j^*}^{(j)}]$ in this case.

- Suppose the challenger responds according to the specification of $\text{ExptIHE}_{\mathcal{A}}(\lambda, 1)$. In this case, the challenger samples $(\text{crs}_{\text{BARG}}, \text{vk}_{\text{BARG}}, \text{td}_{\text{BARG}}) \leftarrow \text{TrapGen}'(1^\lambda, 1^{2n-1}, 1^s, 1^3, \{j^*\})$. Thus, algorithm \mathcal{B} simulates for \mathcal{A} an execution of Expt_{j^*} . We claim that algorithm \mathcal{B} outputs 1 if and only if $\text{Expt}_{j^*}(\mathcal{A})$ outputs 1. The latter corresponds to the following conditions being satisfied:

$$\text{Verify}'(\text{vk}_{\text{BARG}}, C_{\text{idx}}, 2n-1, \pi) = 1 \text{ and } C_{\text{idx}}(j^*, (\hat{v}_{j^*}^{(0)}, \hat{v}_{j^*}^{(1)}, \sigma_{j^*}^{(0)}, \sigma_{j^*}^{(1)}, \tilde{w}_{j^*})) = 1 \text{ and } P(i^*, \hat{v}_{j^*}^{(0)}, \hat{v}_{j^*}^{(1)}, \text{sk}_0, \text{sk}_1, j^*) = 0.$$

where $(\hat{v}_{j^*}^{(0)}, \hat{v}_{j^*}^{(1)}, \sigma_{j^*}^{(0)}, \sigma_{j^*}^{(1)}, \tilde{w}_{j^*}) \leftarrow \text{Extract}'(\text{td}_{\text{BARG}}, \pi, j^*)$. Once again, this is the same set of conditions that \mathcal{B} checks. Thus, in this case algorithm \mathcal{B} outputs 1 with probability $\Pr[\text{Expt}_{j^*}(\mathcal{A}) = 1]$.

We conclude that the distinguishing advantage of \mathcal{B} is precisely

$$\left| \Pr[\text{Expt}_{j^*}(\mathcal{A}) = 1] - \Pr[E_{\text{Verify}}^{(j)} \wedge E_{\text{SAT}, j^*}^{(j)} \wedge \neg E_{P, j^*}^{(j)}] \right| = \varepsilon,$$

which completes the proof. \square

Claim 5.12. *If Π_{BARG} is somewhere extractable then there exists a negligible function $\text{negl}(\cdot)$ such that for all $j^* \in \{j, j_L, j_R\}$, it holds that $\Pr [E_{\text{Verify}}^{(j)} \wedge \neg E_{\text{SAT}, j^*}^{(j)}] = \text{negl}(\lambda)$.*

Proof. Take any $j^* \in \{j, j_L, j_R\}$ and suppose $\Pr [E_{\text{Verify}}^{(j)} \wedge \neg E_{\text{SAT}, j^*}^{(j)}] \geq \varepsilon(\lambda)$ for some non-negligible ε . We use \mathcal{A} to construct an adversary \mathcal{B} that breaks somewhere extractability of Π_{BARG} :

1. On input the security parameter 1^λ , algorithm \mathcal{B} runs algorithm \mathcal{A} to obtain the input length 1^n , the set $S \subseteq [n]$, and an index $i^* \in S$.
2. Algorithm \mathcal{B} outputs $1^{2n-1}, 1^s, 1^3$, the challenge set $J = \{j, j_L, j_R\}$, and the challenge index $j^* \in J$ to the challenger, where s is the bound on the size of the circuit in Fig. 2. The challenger responds with $(\text{crs}_{\text{BARG}}, \text{vk}_{\text{BARG}})$.
3. Algorithm \mathcal{B} computes $(S_0, S_1, \text{idx}) \leftarrow \text{DeriveChal}(S, i^*)$. It then samples the following components:
 - Sample $(\text{sk}_0, \text{pk}_0) \leftarrow \text{HE.Gen}(1^\lambda, 1^n)$ and $(\text{sk}_1, \text{pk}_1) \leftarrow \text{HE.Gen}(1^\lambda, 1^n)$.
 - Sample $\text{crs}_{\text{Com}} \leftarrow \text{Com.Setup}(1^\lambda, 1^{\ell_{\text{ct}}(\lambda, n)}, 2n - 1)$.
 - For each $b \in \{0, 1\}$, sample $\text{ct}_{\text{zero}}^{(b)} \leftarrow \text{HE.Enc}(\text{pk}_b, 0)$. Then, for each $i \in [n]$ and $b \in \{0, 1\}$, if $i \in S_b$, sample $\text{ct}_i^{(b)} \leftarrow \text{HE.Enc}(\text{pk}_b, 1)$; otherwise, if $i \notin S_b$, sample $\text{ct}_i^{(b)} \leftarrow \text{HE.Enc}(\text{pk}_b, 0)$.
 - For each $b \in \{0, 1\}$, let $(\text{com}_{\text{hk}}^{(b)}, \sigma_{\text{hk}, 1}^{(b)}, \dots, \sigma_{\text{hk}, n}^{(b)}) \leftarrow \text{Com.Commit}(\text{crs}_{\text{Com}}, (\text{ct}_1^{(b)}, \dots, \text{ct}_n^{(b)}))$
4. Algorithm \mathcal{B} constructs hk and vk according to Eqs. (5.1) and (5.2):

$$\begin{aligned} \text{hk} &= \left(\text{crs}_{\text{Com}}, \text{crs}_{\text{BARG}}, \{ \text{pk}_b, \text{ct}_{\text{zero}}^{(b)}, \text{ct}_1^{(b)}, \dots, \text{ct}_n^{(b)}, \sigma_{\text{hk}, 1}^{(b)}, \dots, \sigma_{\text{hk}, n}^{(b)} \}_{b \in \{0, 1\}} \right) \\ \text{vk} &= (\text{crs}_{\text{Com}}, \text{vk}_{\text{BARG}}, \text{pk}_0, \text{pk}_1, \text{ct}_{\text{zero}}^{(0)}, \text{ct}_{\text{zero}}^{(1)}, \text{com}_{\text{hk}}^{(0)}, \text{com}_{\text{hk}}^{(1)}) \end{aligned}$$

Algorithm \mathcal{B} gives (hk, vk) to \mathcal{A} .

5. Algorithm \mathcal{A} outputs a digest $\text{dig} = (\text{ct}_{\text{root}}^{(0)}, \text{ct}_{\text{root}}^{(1)}, \text{com}_0, \text{com}_1, \pi_{\text{dig}})$ and a proof π .
6. Let C_{idx} be the circuit as defined in Definition 5.7. Algorithm \mathcal{B} outputs the circuit C_{idx} , the instance number $2n - 1$, and the proof π .

By construction, algorithm \mathcal{B} perfectly simulates an execution of Expt_j . Thus, with probability at least ε , the digest dig and proof π output by \mathcal{A} satisfies $E_{\text{Verify}}^{(j)}$ but not $E_{\text{SAT}, j^*}^{(j)}$. This means

$$\text{Verify}'(\text{vk}_{\text{BARG}}, C_{\text{idx}}, 2n - 1, \pi) = 1 \quad \text{and} \quad C_{\text{idx}}(j^*, (\hat{v}_{j^*}^{(0)}, \hat{v}_{j^*}^{(1)}, \sigma_{j^*}^{(0)}, \sigma_{j^*}^{(1)}, \tilde{w}_{j^*})) = 0.$$

This means algorithm \mathcal{B} successfully breaks somewhere extractability of Π_{BARG} and the claim holds. \square

Claim 5.13. *Suppose the conditions in Claims 5.11 and 5.12 hold. Then, there exists a negligible function $\text{negl}(\cdot)$ such that*

$$\Pr [\text{Expt}'_j(\mathcal{A}) = 1 \wedge (\neg E_{\text{ValidCom}, j_L}^{(j)} \vee \neg E_{P, j_L}^{(j)} \vee \neg E_{\text{ValidCom}, j_R}^{(j)} \vee \neg E_{P, j_R}^{(j)})] \leq 2^{h+1} \cdot \varepsilon_j(\lambda) + \text{negl}(\lambda),$$

where $\varepsilon_j(\lambda) = \max(\varepsilon_{j_L}(\lambda), \varepsilon_{j_R}(\lambda))$.

Proof. By Claim 5.11 there exists a negligible function $\text{negl}_1(\cdot)$ such that for all $j^* \in \{j_L, j_R\}$, it holds that:

$$\left| \Pr [\text{Expt}_{j^*}(\mathcal{A}) = 1] - \Pr [E_{\text{Verify}}^{(j)} \wedge E_{\text{SAT}, j^*}^{(j)} \wedge \neg E_{P, j^*}^{(j)}] \right| \leq \text{negl}_1(\lambda). \quad (5.5)$$

By Claim 5.12 there exists a negligible function $\text{negl}_2(\cdot)$ such that for all $j^* \in \{j_L, j_R\}$ it holds that

$$\Pr [E_{\text{Verify}}^{(j)} \wedge \neg E_{\text{SAT}, j^*}^{(j)}] \leq \text{negl}_2(\lambda). \quad (5.6)$$

By definition, if $\text{Expt}'_j(\mathcal{A}) = 1$, then event $E_{\text{Verify}}^{(j)}$ also occurs. Thus, for all events E , it holds that

$$\Pr[\text{Expt}'_j(\mathcal{A}) = 1 \wedge E] \leq \Pr[E_{\text{Verify}}^{(j)} \wedge E]. \quad (5.7)$$

Similarly, by construction of the circuit C_{idx} , the event $\neg E_{\text{ValidCom},j^*}^{(j)}$ implies event $\neg E_{\text{SAT},j^*}^{(j)}$. Thus, for any event E , it holds that

$$\Pr[\neg E_{\text{ValidCom},j^*}^{(j)} \wedge E] \leq \Pr[\neg E_{\text{SAT},j^*}^{(j)} \wedge E]. \quad (5.8)$$

Take any $j^* \in \{j_L, j_R\}$. Since the height of j^* is h , the inductive hypothesis applies and Eq. (5.4) holds. We first show that

$$\Pr[\text{Expt}'_j(\mathcal{A}) = 1 \wedge \neg E_{P,j^*}^{(j)}] \leq 2^h \cdot \varepsilon_{j^*}(\lambda) + \text{negl}_1(\lambda) + \text{negl}_2(\lambda). \quad (5.9)$$

This follows by the following sequence of calculations:

$$\begin{aligned} \Pr[\text{Expt}'_j(\mathcal{A}) = 1 \wedge \neg E_{P,j^*}^{(j)}] &\leq \Pr[E_{\text{Verify}}^{(j)} \wedge \neg E_{P,j^*}^{(j)}] && \text{by Eq. (5.7)} \\ &= \Pr[E_{\text{Verify}}^{(j)} \wedge E_{\text{SAT},j^*}^{(j)} \wedge \neg E_{P,j^*}^{(j)}] + \Pr[E_{\text{Verify}}^{(j)} \wedge \neg E_{\text{SAT},j^*}^{(j)} \wedge \neg E_{P,j^*}^{(j)}] \\ &\leq \Pr[E_{\text{Verify}}^{(j)} \wedge E_{\text{SAT},j^*}^{(j)} \wedge \neg E_{P,j^*}^{(j)}] + \text{negl}_2(\lambda) && \text{by Eq. (5.6)} \\ &\leq \Pr[\text{Expt}_{j^*}(\mathcal{A}) = 1] + \text{negl}_1(\lambda) + \text{negl}_2(\lambda) && \text{by Eq. (5.5)} \\ &\leq 2^h \cdot \varepsilon_{j^*}(\lambda) + \text{negl}_1(\lambda) + \text{negl}_2(\lambda) && \text{by Eq. (5.4)}. \end{aligned}$$

Next, we have

$$\begin{aligned} \Pr[\text{Expt}'_j(\mathcal{A}) = 1 \wedge \neg E_{\text{ValidCom},j^*}] &\leq \Pr[E_{\text{Verify}}^{(j)} \wedge \neg E_{\text{ValidCom},j^*}] && \text{by Eq. (5.7)} \\ &\leq \Pr[E_{\text{Verify}}^{(j)} \wedge \neg E_{\text{SAT},j^*}] && \text{by Eq. (5.8)} \\ &\leq \text{negl}_2(\lambda) && \text{by Eq. (5.6)}. \end{aligned}$$

Combined with Eq. (5.9) and applying a union bound, we have

$$\begin{aligned} \Pr[\text{Expt}'_j(\mathcal{A}) = 1 \wedge (\neg E_{\text{ValidCom},j_L}^{(j)} \vee \neg E_{P,j_L}^{(j)} \vee \neg E_{\text{ValidCom},j_R}^{(j)} \vee \neg E_{P,j_R}^{(j)})] &\leq 2^h \cdot (\varepsilon_{j_L}(\lambda) + \varepsilon_{j_R}(\lambda)) + \delta(\lambda) \\ &\leq 2^{h+1} \cdot \varepsilon_j(\lambda) + \delta(\lambda), \end{aligned}$$

where $\delta(\lambda) = 2\text{negl}_1(\lambda) + 4\text{negl}_2(\lambda) = \text{negl}(\lambda)$ and $\varepsilon_j(\lambda) = \max(\varepsilon_{j_L}(\lambda), \varepsilon_{j_R}(\lambda))$. \square

Claim 5.14. *If P is a tree-based additive invariant and Π_{Com} is computationally binding, then there exists a negligible function $\text{negl}(\cdot)$ such that*

$$\Pr[\text{Expt}'_j(\mathcal{A}) = 1 \wedge E_{\text{ValidCom},j_L}^{(j)} \wedge E_{P,j_L}^{(j)} \wedge E_{\text{ValidCom},j_R}^{(j)} \wedge E_{P,j_R}^{(j)}] \leq \text{negl}(\lambda).$$

Proof. Suppose

$$\Pr[\text{Expt}'_j(\mathcal{A}) = 1 \wedge E_{\text{ValidCom},j_L}^{(j)} \wedge E_{P,j_L}^{(j)} \wedge E_{\text{ValidCom},j_R}^{(j)} \wedge E_{P,j_R}^{(j)}] \geq \varepsilon(\lambda),$$

for some non-negligible ε . We use \mathcal{A} to construct an adversary \mathcal{B} for the binding game for Π_{Com} as follows:

1. On input the security parameter 1^λ , algorithm \mathcal{B} runs algorithm \mathcal{A} to obtain the input length 1^n , the set $S \subseteq [n]$, and an index $i^* \in S$.
2. Algorithm \mathcal{B} outputs the block length $1^{\ell_{\text{ct}}(\lambda,n)}$ and the vector length $2n - 1$ to the challenger. The challenger responds with crs_{Com} .
3. Algorithm \mathcal{B} computes $(S_0, S_1, \text{idx}) \leftarrow \text{DeriveChal}(S, i^*)$. It then samples the following components:
 - Sample $(\text{sk}_0, \text{pk}_0) \leftarrow \text{HE.Gen}(1^\lambda, 1^n)$ and $(\text{sk}_1, \text{pk}_1) \leftarrow \text{HE.Gen}(1^\lambda, 1^n)$.

- Sample $(\text{crs}_{\text{BARG}}, \text{vk}_{\text{BARG}}, \text{td}_{\text{BARG}}) \leftarrow \text{TrapGen}'(1^\lambda, 1^{2n-1}, 1^s, 1^3, \{j, j_L, j_R\})$.
- For each $b \in \{0, 1\}$, sample $\text{ct}_{\text{zero}}^{(b)} \leftarrow \text{HE.Enc}(\text{pk}_b, 0)$. Then, for each $i \in [n]$ and $b \in \{0, 1\}$, if $i \in S_b$, sample $\text{ct}_i^{(b)} \leftarrow \text{HE.Enc}(\text{pk}_b, 1)$; otherwise, if $i \notin S_b$, sample $\text{ct}_i^{(b)} \leftarrow \text{HE.Enc}(\text{pk}_b, 0)$.
- For each $b \in \{0, 1\}$, let $(\text{com}_{\text{hk}}^{(b)}, \sigma_{\text{hk},1}^{(b)}, \dots, \sigma_{\text{hk},n}^{(b)}) \leftarrow \text{Com.Commit}(\text{crs}_{\text{Com}}, (\text{ct}_1^{(b)}, \dots, \text{ct}_n^{(b)}))$

4. Algorithm \mathcal{B} constructs hk and vk according to Eqs. (5.1) and (5.2):

$$\begin{aligned} \text{hk} &= \left(\text{crs}_{\text{Com}}, \text{crs}_{\text{BARG}}, \{ \text{pk}_b, \text{ct}_{\text{zero}}^{(b)}, \text{ct}_1^{(b)}, \dots, \text{ct}_n^{(b)}, \sigma_{\text{hk},1}^{(b)}, \dots, \sigma_{\text{hk},n}^{(b)} \}_{b \in \{0,1\}} \right) \\ \text{vk} &= (\text{crs}_{\text{Com}}, \text{vk}_{\text{BARG}}, \text{pk}_0, \text{pk}_1, \text{ct}_{\text{zero}}^{(0)}, \text{ct}_{\text{zero}}^{(1)}, \text{com}_{\text{hk}}^{(0)}, \text{com}_{\text{hk}}^{(1)}) \end{aligned}$$

Algorithm \mathcal{B} gives (hk, vk) to \mathcal{A} .

5. Algorithm \mathcal{A} outputs a digest $\text{dig} = (\text{ct}_{\text{root}}^{(0)}, \text{ct}_{\text{root}}^{(1)}, \text{com}_0, \text{com}_1, \pi_{\text{dig}})$ and a proof π .

6. Algorithm \mathcal{B} computes the following:

- $(\hat{v}_j^{(0)}, \hat{v}_j^{(1)}, \sigma_j^{(0)}, \sigma_j^{(1)}, \tilde{w}_j) \leftarrow \text{Extract}'(\text{td}_{\text{BARG}}, \pi, j)$.
- $(\hat{v}_{j_L}^{(0)}, \hat{v}_{j_L}^{(1)}, \sigma_{j_L}^{(0)}, \sigma_{j_L}^{(1)}, \tilde{w}_{j_L}) \leftarrow \text{Extract}'(\text{td}_{\text{BARG}}, \pi, j_L)$.
- $(\hat{v}_{j_R}^{(0)}, \hat{v}_{j_R}^{(1)}, \sigma_{j_R}^{(0)}, \sigma_{j_R}^{(1)}, \tilde{w}_{j_R}) \leftarrow \text{Extract}'(\text{td}_{\text{BARG}}, \pi, j_R)$.

In addition, it parses $\tilde{w}_j = (\tilde{w}_{j_L}, \tilde{w}_{j_R})$ and $\tilde{w}_{j_L} = (\hat{v}_{j_L}^{(0)}, \hat{v}_{j_L}^{(1)}, \sigma_{j_L}^{(0)}, \sigma_{j_L}^{(1)})$ and $\tilde{w}_{j_R} = (\hat{v}_{j_R}^{(0)}, \hat{v}_{j_R}^{(1)}, \sigma_{j_R}^{(0)}, \sigma_{j_R}^{(1)})$.

7. Algorithm \mathcal{B} checks if there exists $b \in \{0, 1\}$ and $d \in \{L, R\}$ such that $\hat{v}_{j_d}^{(b)} \neq \hat{v}_{j,d}^{(b)}$ and

$$\text{Com.Verify}(\text{crs}_{\text{Com}}, \text{com}_b, j_d, \hat{v}_{j_d}^{(b)}, \sigma_{j_d}^{(b)}) = 1 \quad \text{and} \quad \text{Com.Verify}(\text{crs}_{\text{Com}}, \text{com}_b, j_d, \hat{v}_{j,d}^{(b)}, \sigma_{j,d}^{(b)}) = 1.$$

If so, it outputs the commitment com_b , the index $j_d \in [2n - 1]$, and the value-opening pairs $(\hat{v}_{j,d}^{(b)}, \sigma_{j,d}^{(b)})$ and $(\hat{v}_{j_d}^{(b)}, \sigma_{j_d}^{(b)})$. Otherwise, algorithm \mathcal{B} aborts with output \perp .

By construction, algorithm \mathcal{B} perfectly simulates an execution of Expt'_j for adversary \mathcal{A} . By assumption, with probability at least ϵ , algorithm \mathcal{A} will output a digest dig and a proof π such that the following conditions hold:

- $\text{Expt}'_j(\mathcal{A}) = 1$: This means $\text{Verify}'(\text{vk}_{\text{BARG}}, C_{\text{idX}}, 2n - 1, \pi) = 1$, $C_{\text{idX}}(j, (\hat{v}_j^{(0)}, \hat{v}_j^{(1)}, \sigma_j^{(0)}, \sigma_j^{(1)}, \tilde{w}_j)) = 1$, and $P(\hat{v}_j^{(0)}, \hat{v}_j^{(1)}, \text{sk}_0, \text{sk}_1, j) = 0$.
- $E_{\text{ValidCom}, j^*}^{(j)}$ for $j^* \in \{j_L, j_R\}$: This means

$$\text{Com.Verify}(\text{crs}_{\text{Com}}, \text{com}_0, j^*, \hat{v}_{j^*}^{(0)}, \sigma_{j^*}^{(0)}) = 1 = \text{Com.Verify}(\text{crs}_{\text{Com}}, \text{com}_1, j^*, \hat{v}_{j^*}^{(1)}, \sigma_{j^*}^{(1)}).$$

- $E_{P, j^*}^{(j)}$ for $j^* \in \{j_L, j_R\}$: This means $P(\hat{v}_{j^*}^{(0)}, \hat{v}_{j^*}^{(1)}, \text{sk}_0, \text{sk}_1, j^*) = 1$.

We consider two possibilities:

- Suppose for all $b \in \{0, 1\}$, we have $\hat{v}_{j_L}^{(b)} = \hat{v}_{j_L}^{(b)}$ and $\hat{v}_{j_R}^{(b)} = \hat{v}_{j_R}^{(b)}$. Since $C_{i^*, y}(j, (\hat{v}_j^{(0)}, \hat{v}_j^{(1)}, \sigma_j^{(0)}, \sigma_j^{(1)}, \tilde{w}_j)) = 1$, this means that $\hat{v}_j^{(b)} = \text{HE.Add}(\text{pk}_b, \hat{v}_{j_L}^{(b)}, \hat{v}_{j_R}^{(b)})$ for all $b \in \{0, 1\}$. Combined with the third condition, this means

$$\begin{aligned} P(i^*, \hat{v}_{j_L}^{(0)}, \hat{v}_{j_L}^{(1)}, \text{sk}_0, \text{sk}_1, j_L) &= P(i^*, \hat{v}_{j_L}^{(0)}, \hat{v}_{j_L}^{(1)}, \text{sk}_0, \text{sk}_1, j_L) = 1 \\ P(i^*, \hat{v}_{j_R}^{(0)}, \hat{v}_{j_R}^{(1)}, \text{sk}_0, \text{sk}_1, j_R) &= P(i^*, \hat{v}_{j_R}^{(0)}, \hat{v}_{j_R}^{(1)}, \text{sk}_0, \text{sk}_1, j_R) = 1. \end{aligned}$$

Since P is a tree-based additive invariant, this means that

$$P(i^*, \hat{v}_j^{(0)}, \hat{v}_j^{(1)}, \text{sk}_0, \text{sk}_1, j) = 1.$$

However, this contradicts the condition that $P(i^*, \hat{v}_j^{(0)}, \hat{v}_j^{(1)}, \text{sk}_0, \text{sk}_1, j) = 0$, so this case does not occur.

- Suppose there exists $b \in \{0, 1\}$ and $d \in \{L, R\}$ where $\hat{v}_{j_d}^{(b)} \neq \hat{v}_{j_d}^{(b)}$. From Fig. 2, since

$$C_{i^*, y}(j, (\hat{v}_j^{(0)}, \hat{v}_j^{(1)}, \sigma_j^{(0)}, \sigma_j^{(1)}, \tilde{w}_j)) = 1,$$

this means that $\text{Com.Verify}(\text{crs}_{\text{Com}}, \text{com}_b, j_d, \hat{v}_{j_d}^{(b)}, \sigma_{j_d}^{(b)}) = 1$. By the second condition, we also have

$$\text{Com.Verify}(\text{crs}_{\text{Com}}, \text{com}_b, j_d, \hat{v}_{j_d}^{(b)}, \sigma_{j_d}^{(b)}) = 1.$$

In this case, algorithm \mathcal{B} outputs the commitment com_b , the index j_d , and the value-opening pairs $(\hat{v}_{j_d}^{(b)}, \sigma_{j_d}^{(b)})$ and $(\hat{v}_{j_d}^{(b)}, \sigma_{j_d}^{(b)})$. This is a pair of valid openings for com_b so algorithm \mathcal{B} wins the binding game.

We conclude that algorithm \mathcal{B} succeeds with the same advantage ε and the claim follows. \square

Claim 5.15. *Suppose the conditions of Claims 5.13 and 5.14 hold. Then there exists a negligible function $\text{negl}(\cdot)$ such that*

$$\Pr[\text{Expt}'_j(\mathcal{A}) = 1] \leq 2^{h+1} \cdot \varepsilon_j(\lambda) + \text{negl}(\lambda),$$

where $\varepsilon_j(\lambda) = \max(\varepsilon_{j_L}(\lambda), \varepsilon_{j_R}(\lambda))$.

Proof. By the law of total probability, we have

$$\begin{aligned} \Pr[\text{Expt}'_j(\mathcal{A}) = 1] &\leq \Pr[\text{Expt}'_j(\mathcal{A}) = 1 \wedge E_{\text{ValidCom}, j_L}^{(j)} \wedge E_{P, j_L}^{(j)} \wedge E_{\text{ValidCom}, j_R}^{(j)} \wedge E_{P, j_R}^{(j)}] + \\ &\quad \Pr[\text{Expt}'_j(\mathcal{A}) = 1 \wedge (\neg E_{\text{ValidCom}, j_L}^{(j)} \vee \neg E_{P, j_L}^{(j)} \vee \neg E_{\text{ValidCom}, j_R}^{(j)} \vee \neg E_{P, j_R}^{(j)})]. \end{aligned}$$

By Claims 5.13 and 5.14, there exist negligible functions $\text{negl}_1(\cdot)$ and $\text{negl}_2(\cdot)$ such that:

$$\begin{aligned} \Pr[\text{Expt}'_j(\mathcal{A}) = 1 \wedge (\neg E_{\text{ValidCom}, j_L}^{(j)} \vee \neg E_{P, j_L}^{(j)} \vee \neg E_{\text{ValidCom}, j_R}^{(j)} \vee \neg E_{P, j_R}^{(j)})] &\leq 2^{h+1} \cdot \varepsilon_j(\lambda) + \text{negl}_1(\lambda) \\ \Pr[\text{Expt}'_j(\mathcal{A}) = 1 \wedge E_{\text{ValidCom}, j_L}^{(j)} \wedge E_{P, j_L}^{(j)} \wedge E_{\text{ValidCom}, j_R}^{(j)} \wedge E_{P, j_R}^{(j)}] &\leq \text{negl}_2(\lambda). \end{aligned}$$

where $\varepsilon_j(\lambda) = \max(\varepsilon_{j_L}(\lambda), \varepsilon_{j_R}(\lambda))$. The claim follows. \square

Completing the proof of Lemma 5.10. To complete the proof of the inductive step (for Lemma 5.10), we first appeal to Claim 5.15 to conclude that there exists negligible function $\text{negl}_1(\cdot)$ such that

$$\Pr[\text{Expt}'_j(\mathcal{A}) = 1] \leq 2^{h+1} \cdot \varepsilon_j(\lambda) + \text{negl}_1(\lambda),$$

where $\varepsilon_j(\lambda) = \max(\varepsilon_{j_L}(\lambda), \varepsilon_{j_R}(\lambda))$. From the inductive hypothesis, $\varepsilon_{j_L}(\lambda)$ and $\varepsilon_{j_R}(\lambda)$ are both negligible functions. By definition of Expt'_j , we have that

$$\Pr[\text{Expt}'_j(\mathcal{A}) = 1] = \Pr[E_{\text{Verify}}^{(j)} \wedge E_{\text{SAT}, j}^{(j)} \wedge \neg E_{P, j}^{(j)}].$$

By Claim 5.11, there exists a negligible function $\text{negl}_2(\cdot)$ such that

$$\left| \Pr[\text{Expt}_j(\mathcal{A}) = 1] - \Pr[\text{Expt}'_j(\mathcal{A}) = 1] \right| \leq \text{negl}_2(\lambda).$$

We conclude that

$$\Pr[\text{Expt}_j(\mathcal{A}) = 1] \leq 2^{h+1} \cdot \varepsilon_j(\lambda) + \text{negl}_1(\lambda) + \text{negl}_2(\lambda).$$

Setting $\varepsilon'_j(\lambda) = \max(\varepsilon_j(\lambda), (\text{negl}_1(\lambda) + \text{negl}_2(\lambda))/2^{h+1})$, we have that $\Pr[\text{Expt}_j(\mathcal{A}) = 1] \leq 2^{h+1} \cdot \varepsilon'_j(\lambda)$, where $\varepsilon'_j(\lambda)$ is a negligible function. Lemma 5.10 now follows by induction on the height h . \square

Completing the proof of Theorem 5.9. We now use Lemma 5.10 to complete the proof of Theorem 5.9. Suppose the conditions of Theorem 5.9 hold. Noting that the index $2n - 1$ has height $h = \log n$ in a complete binary tree with n leaves, we appeal to Lemma 5.10 and conclude that there exists a negligible function $\text{negl}(\cdot)$ such that

$$\Pr[\text{Expt}_{2n-1}(\mathcal{A}) = 1] \leq n \cdot \text{negl}(\lambda). \quad (5.10)$$

To complete the proof, we define a sequence of hybrid experiments:

- Hyb_0 : This is the experiment $\text{Expt}_{2n-1}[P, \text{DeriveChal}]$ with adversary \mathcal{A} .
- Hyb_1 : Same as Hyb_0 , except the output of the experiment is 1 if the following properties hold:
 - $\text{Verify}'(\text{vk}_{\text{BARG}}, C_{\text{idx}}, 2n - 1, \pi) = 1$;
 - $C_{\text{idx}}(2n - 1, (\hat{v}_{2n-1}^{(0)}, \hat{v}_{2n-1}^{(1)}, \sigma_{2n-1}^{(0)}, \sigma_{2n-1}^{(1)}, \tilde{w}_{2n-1})) = 1$; and
 - $P(i^*, \text{ct}_{\text{root}}^{(0)}, \text{ct}_{\text{root}}^{(1)}, \text{sk}_0, \text{sk}_1, 2n - 1) = 1$.
- Hyb_2 : Same as Hyb_1 , except the output of the experiment is 1 if the following properties hold:
 - $\text{Verify}'(\text{vk}_{\text{BARG}}, C_{\text{idx}}, 2n - 1, \pi) = 1$; and
 - $P(i^*, \text{ct}_{\text{root}}^{(0)}, \text{ct}_{\text{root}}^{(1)}, \text{sk}_0, \text{sk}_1, 2n - 1) = 1$.

In particular, the challenger no longer checks the value of C_{idx} . Note that in this experiment, the challenger's behavior no longer depends on the BARG trapdoor td_{BARG} .

- Hyb_3 : Same as Hyb_2 , except when sampling the BARG parameters at the beginning of the experiment, the challenger now samples $(\text{crs}_{\text{BARG}}, \text{vk}_{\text{BARG}}) \leftarrow \text{Gen}'(1^\lambda, 1^{2n-1}, 1^s, 1^3)$. This corresponds to the experiment $\text{Expt}[P, \text{DeriveChal}]$ with adversary \mathcal{A} .

For an adversary \mathcal{A} , we write $\text{Hyb}_i(\mathcal{A}) = 1$ to denote the output of Hyb_i with adversary \mathcal{A} . We now analyze each pair of adjacent experiments.

Claim 5.16. *It holds that $\Pr[\text{Hyb}_1(\mathcal{A}) = 1] = \Pr[\text{Hyb}_0(\mathcal{A}) = 1]$.*

Proof. These experiments are identical. Specifically, by definition of C_{idx} (and specifically, the relation in Fig. 2), if $C_{\text{idx}}(2n - 1, (\hat{v}_{2n-1}^{(0)}, \hat{v}_{2n-1}^{(1)}, \sigma_{2n-1}^{(0)}, \sigma_{2n-1}^{(1)}, \tilde{w}_{2n-1})) = 1$, then $\hat{v}_{2n-1}^{(b)} = \text{ct}_{\text{root}}^{(b)}$ for $b \in \{0, 1\}$. This means that

$$P(i^*, \hat{v}_{2n-1}^{(0)}, \hat{v}_{2n-1}^{(1)}, \text{sk}_0, \text{sk}_1, 2n - 1) = P(i^*, \text{ct}_{\text{root}}^{(0)}, \text{ct}_{\text{root}}^{(1)}, \text{sk}_0, \text{sk}_1, 2n - 1).$$

Thus, the output of $\text{Hyb}_0(\mathcal{A})$ is identical to that of $\text{Hyb}_1(\mathcal{A})$. \square

Claim 5.17. *If Π_{BARG} is somewhere extractable, then there exists a negligible function $\text{negl}(\cdot)$ such that*

$$|\Pr[\text{Hyb}_2(\mathcal{A}) = 1] - \Pr[\text{Hyb}_1(\mathcal{A}) = 1]| = \text{negl}(\lambda).$$

Proof. Suppose $\Pr[\text{Hyb}_2(\mathcal{A}) = 1] - \Pr[\text{Hyb}_1(\mathcal{A}) = 1] \geq \varepsilon(\lambda)$, for some non-negligible ε . Since the only difference between Hyb_1 and Hyb_2 is the conditions the challenger checks at the very end of the experiment, this means that with probability at least ε , the adversary in Hyb_1 will output a digest dig and a proof π such that the following conditions hold:

- $\text{Verify}'(\text{vk}_{\text{BARG}}, C_{\text{idx}}, 2n - 1, \pi) = 1$.
- $P(i^*, \text{ct}_{\text{root}}^{(0)}, \text{ct}_{\text{root}}^{(1)}, \text{sk}_0, \text{sk}_1, 2n - 1) = 1$.
- $C_{\text{idx}}(2n - 1, (\hat{v}_{2n-1}^{(0)}, \hat{v}_{2n-1}^{(1)}, \sigma_{2n-1}^{(0)}, \sigma_{2n-1}^{(1)}, \tilde{w}_{2n-1})) = 0$.

In all other settings, the output of the two experiments are identical. We use \mathcal{A} to construct an adversary \mathcal{B} that breaks somewhere extractability of Π_{BARG} (similar to the proof of Claim 5.12):

1. On input the security parameter 1^λ , algorithm \mathcal{B} runs algorithm \mathcal{A} to obtain the input length 1^n , the set $S \subseteq [n]$, and an index $i^* \in S$.
2. Let $j = 2n - 1$ and j_L, j_R be the indices of the input wires that determine the value of the output wire j . Algorithm \mathcal{B} outputs $1^{2n-1}, 1^s, 1^3$, the challenge set $J = \{j, j_R, j_L\}$, and the challenge index $j = 2n - 1$ to the challenger. Here, s is the bound on the size of the circuit in Fig. 2. The challenger responds with $(\text{crs}_{\text{BARG}}, \text{vk}_{\text{BARG}})$.
3. Algorithm \mathcal{B} computes $(S_0, S_1, \text{idx}) \leftarrow \text{DeriveChal}(S, i^*)$. It then samples the following components:
 - Sample $(\text{sk}_0, \text{pk}_0) \leftarrow \text{HE.Gen}(1^\lambda, 1^n)$ and $(\text{sk}_1, \text{pk}_1) \leftarrow \text{HE.Gen}(1^\lambda, 1^n)$.
 - Sample $\text{crs}_{\text{Com}} \leftarrow \text{Com.Setup}(1^\lambda, 1^{\ell_{\text{ct}}(\lambda, n)}, 2n - 1)$.
 - For each $b \in \{0, 1\}$, sample $\text{ct}_{\text{zero}}^{(b)} \leftarrow \text{HE.Enc}(\text{pk}_b, 0)$. Then, for each $i \in [n]$ and $b \in \{0, 1\}$, if $i \in S_b$, sample $\text{ct}_i^{(b)} \leftarrow \text{HE.Enc}(\text{pk}_b, 1)$; otherwise, if $i \notin S_b$, sample $\text{ct}_i^{(b)} \leftarrow \text{HE.Enc}(\text{pk}_b, 0)$.
 - For each $b \in \{0, 1\}$, let $(\text{com}_{\text{hk}}^{(b)}, \sigma_{\text{hk},1}^{(b)}, \dots, \sigma_{\text{hk},n}^{(b)}) \leftarrow \text{Com.Commit}(\text{crs}_{\text{Com}}, (\text{ct}_1^{(b)}, \dots, \text{ct}_n^{(b)}))$
4. Algorithm \mathcal{B} constructs hk and vk according to Eqs. (5.1) and (5.2):

$$\begin{aligned} \text{hk} &= \left(\text{crs}_{\text{Com}}, \text{crs}_{\text{BARG}}, \{ \text{pk}_b, \text{ct}_{\text{zero}}^{(b)}, \text{ct}_1^{(b)}, \dots, \text{ct}_n^{(b)}, \sigma_{\text{hk},1}^{(b)}, \dots, \sigma_{\text{hk},n}^{(b)} \}_{b \in \{0,1\}} \right) \\ \text{vk} &= (\text{crs}_{\text{Com}}, \text{vk}_{\text{BARG}}, \text{pk}_0, \text{pk}_1, \text{ct}_{\text{zero}}^{(0)}, \text{ct}_{\text{zero}}^{(1)}, \text{com}_{\text{hk}}^{(0)}, \text{com}_{\text{hk}}^{(1)}) \end{aligned}$$

Algorithm \mathcal{B} gives (hk, vk) to \mathcal{A} .

5. Algorithm \mathcal{A} outputs a digest $\text{dig} = (\text{ct}_{\text{root}}^{(0)}, \text{ct}_{\text{root}}^{(1)}, \text{com}_0, \text{com}_1, \pi_{\text{dig}})$ and a proof π .
6. Let C_{idx} be the circuit as defined in Definition 5.7. Algorithm \mathcal{B} outputs the circuit C_{idx} , the instance number $2n - 1$, and the proof π .

By definition, the challenger samples $(\text{crs}_{\text{BARG}}, \text{vk}_{\text{BARG}}, \text{td}_{\text{BARG}}) \leftarrow \text{TrapGen}'(1^\lambda, 1^{2n-1}, 1^s, 1^3, \{j, j_L, j_R\})$. This means algorithm \mathcal{B} perfectly simulates an execution of Hyb_1 . Thus, with probability at least ε , the digest dig and proof π output by \mathcal{A} satisfies

$$\text{Verify}'(\text{vk}_{\text{BARG}}, C_{\text{idx}}, 2n - 1, \pi) = 1 \quad \text{and} \quad C_{\text{idx}}(2n - 1, (\hat{v}_{2n-1}^{(0)}, \hat{v}_{2n-1}^{(1)}, \sigma_{2n-1}^{(0)}, \sigma_{2n-1}^{(1)}, \tilde{w}_{2n-1})) = 0,$$

where $(\hat{v}_{2n-1}^{(0)}, \hat{v}_{2n-1}^{(1)}, \sigma_{2n-1}^{(0)}, \sigma_{2n-1}^{(1)}, \tilde{w}_{2n-1}) \leftarrow \text{Extract}'(\text{td}_{\text{BARG}}, \pi, 2n - 1)$. This means algorithm \mathcal{B} successfully breaks somewhere extractability of Π_{BARG} and the claim holds. \square

Claim 5.18. *If Π_{BARG} satisfies set hiding then there exists a negligible function $\text{negl}(\cdot)$ such that*

$$\left| \Pr[\text{Hyb}_3(\mathcal{A}) = 1] - \Pr[\text{Hyb}_2(\mathcal{A}) = 1] \right| = \text{negl}(\lambda).$$

Proof. Suppose $\left| \Pr[\text{Hyb}_3(\mathcal{A}) = 1] - \Pr[\text{Hyb}_2(\mathcal{A}) = 1] \right| \geq \varepsilon(\lambda)$ for some non-negligible ε . We use \mathcal{A} to construct an adversary \mathcal{B} that breaks set hiding of Π_{BARG} :

1. On input the security parameter 1^λ , algorithm \mathcal{B} runs algorithm \mathcal{A} to obtain the input length 1^n , the set $S \subseteq [n]$, and an index $i^* \in S$.
2. Let $j = 2n - 1$ and j_L, j_R be the indices of the input wires that determine the value of the output wire j . Algorithm \mathcal{B} outputs $1^{2n-1}, 1^s, 1^3$ and the challenge set $J = \{j, j_L, j_R\}$ to the challenger. Here, s is the bound on the size of the circuit in Fig. 2. The challenger responds with $(\text{crs}_{\text{BARG}}, \text{vk}_{\text{BARG}})$.
3. Algorithm \mathcal{B} computes $(S_0, S_1) \leftarrow \text{DeriveChal}(S, i^*)$. It then samples the following components:
 - Sample $(\text{sk}_0, \text{pk}_0) \leftarrow \text{HE.Gen}(1^\lambda, 1^n)$ and $(\text{sk}_1, \text{pk}_1) \leftarrow \text{HE.Gen}(1^\lambda, 1^n)$.
 - Sample $\text{crs}_{\text{Com}} \leftarrow \text{Com.Setup}(1^\lambda, 1^{\ell_{\text{ct}}(\lambda, n)}, 2n - 1)$.

- For each $b \in \{0, 1\}$, sample $ct_{\text{zero}}^{(b)} \leftarrow \text{HE.Enc}(\text{pk}_b, 0)$. Then, for each $i \in [n]$ and $b \in \{0, 1\}$, if $i \in S_b$, sample $ct_i^{(b)} \leftarrow \text{HE.Enc}(\text{pk}_b, 1)$; otherwise, if $i \notin S_b$, sample $ct_i^{(b)} \leftarrow \text{HE.Enc}(\text{pk}_b, 0)$.
- For each $b \in \{0, 1\}$, let $(\text{com}_{\text{hk}}^{(b)}, \sigma_{\text{hk},1}^{(b)}, \dots, \sigma_{\text{hk},n}^{(b)}) \leftarrow \text{Com.Commit}(\text{crs}_{\text{Com}}, (ct_1^{(b)}, \dots, ct_n^{(b)}))$

4. Algorithm \mathcal{B} constructs hk and vk according to Eqs. (5.1) and (5.2):

$$\begin{aligned} \text{hk} &= \left(\text{crs}_{\text{Com}}, \text{crs}_{\text{BARG}}, \{ \text{pk}_b, ct_{\text{zero}}^{(b)}, ct_1^{(b)}, \dots, ct_n^{(b)}, \sigma_{\text{hk},1}^{(b)}, \dots, \sigma_{\text{hk},n}^{(b)} \}_{b \in \{0,1\}} \right) \\ \text{vk} &= (\text{crs}_{\text{Com}}, \text{vk}_{\text{BARG}}, \text{pk}_0, \text{pk}_1, ct_{\text{zero}}^{(0)}, ct_{\text{zero}}^{(1)}, \text{com}_{\text{hk}}^{(0)}, \text{com}_{\text{hk}}^{(1)}) \end{aligned}$$

Algorithm \mathcal{B} gives (hk, vk) to \mathcal{A} .

5. Algorithm \mathcal{A} outputs a digest $\text{dig} = (ct_{\text{root}}^{(0)}, ct_{\text{root}}^{(1)}, \text{com}_0, \text{com}_1, \pi_{\text{dig}})$ and a proof π .
6. Let $C_{i^*, x_{i^*}}$ be the circuit as defined in Definition 5.7. Algorithm \mathcal{B} outputs 1 if

$$\text{Verify}'(\text{vk}_{\text{BARG}}, C_{\text{id}_x}, 2n-1, \pi) = 1 \quad \text{and} \quad P(i^*, ct_{\text{root}}^{(0)}, ct_{\text{root}}^{(1)}, \text{sk}_0, \text{sk}_1, 2n-1) = 1$$

Otherwise, algorithm \mathcal{B} outputs 0.

We now consider the two possibilities:

- Suppose the challenger responds according to the specification of $\text{ExptSH}_{\mathcal{A}}(\lambda, 0)$. In this case, the challenger samples $(\text{crs}_{\text{BARG}}, \text{vk}_{\text{BARG}}) \leftarrow \text{Gen}'(1^\lambda, 1^{2n-1}, 1^s, 1^3)$. In this case, algorithm \mathcal{B} perfectly simulates an execution of Hyb_3 for \mathcal{A} . Moreover, algorithm \mathcal{B} computes the outputs according to the same specification of Hyb_3 , so we conclude that algorithm \mathcal{B} outputs 1 with $\Pr[\text{Hyb}_3(\mathcal{A}) = 1]$.
- Suppose the challenger responds according to the specification of $\text{ExptSH}_{\mathcal{A}}(\lambda, 1)$. In this case, the challenger samples $(\text{crs}_{\text{BARG}}, \text{vk}_{\text{BARG}}, \text{td}_{\text{BARG}}) \leftarrow \text{TrapGen}'(1^\lambda, 1^{2n-1}, 1^s, 1^3, \{j, j_L, j_R\})$. In this case, algorithm \mathcal{B} perfectly simulates an execution of Hyb_2 for \mathcal{A} , and correspondingly, algorithm \mathcal{B} outputs 1 with probability $\Pr[\text{Hyb}_2(\mathcal{A}) = 1]$.

We conclude that the distinguishing advantage of \mathcal{B} is exactly ε , which concludes the proof. \square

Combining Claims 5.16 to 5.18, we conclude that $|\Pr[\text{Hyb}_0(\mathcal{A}) = 1] - \Pr[\text{Hyb}_3(\mathcal{A}) = 1]| = \text{negl}(\lambda)$. By construction, $\text{Hyb}_0(\mathcal{A}) \equiv \text{Expt}_{2n-1}(\mathcal{A})$ and $\text{Hyb}_3(\mathcal{A}) \equiv \text{Expt}(\mathcal{A})$. From Eq. (5.10), we have that $\Pr[\text{Expt}_{2n-1}(\mathcal{A}) = 1] = \text{negl}(\lambda)$ and Theorem 5.9 follows. \square

5.1.2 Set Hiding

In this section, we show that Construction 5.2 satisfies set hiding. This follows immediately from CPA-security of the underlying encryption scheme. Recall that in Construction 5.2, the only difference between a hash key that binds to the empty set \emptyset versus the set S is that some of the ciphertexts in the hash key switch from encryptions of 0 (when binding to the empty set) to an encryption of 1 (when binding to the set S). We formalize this below:

Theorem 5.19 (Set Hiding). *If Π_{HE} is CPA-secure, then Construction 5.2 satisfies set hiding.*

Proof. Let \mathcal{A} be an adversary for the set hiding game. We define a hybrid experiment Hyb_β for each $\beta \in \{0, 1, 2\}$ as follows:

1. On input 1^λ , algorithm \mathcal{A} outputs the input length 1^n and set $S \subseteq [n]$.
2. The challenger now samples the following quantities:
 - Sample $(\text{sk}_0, \text{pk}_0) \leftarrow \text{HE.Gen}(1^\lambda, 1^n)$ and $(\text{sk}_1, \text{pk}_1) \leftarrow \text{HE.Gen}(1^\lambda, 1^n)$.
 - Sample $\text{crs}_{\text{Com}} \leftarrow \text{Com.Setup}(1^\lambda, 1^{\ell_{\text{ct}}(\lambda, n)}, 2n-1)$. and $(\text{crs}_{\text{BARG}}, \text{vk}_{\text{BARG}}) \leftarrow \text{Gen}'(1^\lambda, 1^{2n-1}, 1^s, 1^3)$.

- For each $b \in \{0, 1\}$, sample $\text{ct}_{\text{zero}}^{(b)} \leftarrow \text{HE.Enc}(\text{pk}_b, 0)$.
- For each $i \in [n]$ and $b \in \{0, 1\}$, if $i \in S$ and $b < \beta$, the challenger samples $\text{ct}_i^{(b)} \leftarrow \text{HE.Enc}(\text{pk}_b, 1)$. Otherwise, if $i \notin S$ or $b \geq \beta$, it samples $\text{ct}_i^{(b)} \leftarrow \text{HE.Enc}(\text{pk}_b, 0)$.
- For each $b \in \{0, 1\}$, let $(\text{com}_{\text{hk}}^{(b)}, \sigma_{\text{hk},1}^{(b)}, \dots, \sigma_{\text{hk},n}^{(b)}) \leftarrow \text{Com.Commit}(\text{crs}_{\text{Com}}, (\text{ct}_1^{(b)}, \dots, \text{ct}_n^{(b)}))$

3. Algorithm \mathcal{B} constructs hk and vk according to Eqs. (5.1) and (5.2):

$$\begin{aligned} \text{hk} &= \left(\text{crs}_{\text{Com}}, \text{crs}_{\text{BARG}}, \{ \text{pk}_b, \text{ct}_{\text{zero}}^{(b)}, \text{ct}_1^{(b)}, \dots, \text{ct}_n^{(b)}, \sigma_{\text{hk},1}^{(b)}, \dots, \sigma_{\text{hk},n}^{(b)} \}_{b \in \{0,1\}} \right) \\ \text{vk} &= (\text{crs}_{\text{Com}}, \text{vk}_{\text{BARG}}, \text{pk}_0, \text{pk}_1, \text{ct}_{\text{zero}}^{(0)}, \text{ct}_{\text{zero}}^{(1)}, \text{com}_{\text{hk}}^{(0)}, \text{com}_{\text{hk}}^{(1)}) \end{aligned}$$

Algorithm \mathcal{B} gives (hk, vk) to \mathcal{A} .

4. Algorithm \mathcal{A} outputs a bit b' which is the output of the experiment.

Let $\text{Hyb}_\beta(\mathcal{A})$ be the output of Hyb_β with adversary \mathcal{A} . By construction $\text{ExptSH}_{\mathcal{A}}(\lambda, 0) \equiv \text{Hyb}_0(\mathcal{A})$ and similarly, $\text{ExptSH}_{\mathcal{A}}(\lambda, 1) \equiv \text{Hyb}_2(\mathcal{A})$. We now argue that each adjacent pair of hybrid distributions are computationally indistinguishable.

Claim 5.20. *If Π_{HE} is CPA-secure, then there exists a negligible function $\text{negl}(\cdot)$ such that*

$$|\Pr[\text{Hyb}_1(\mathcal{A}) = 1] - \Pr[\text{Hyb}_0(\mathcal{A}) = 1]| = \text{negl}(\lambda).$$

Proof. Suppose that $|\Pr[\text{Hyb}_1(\mathcal{A}) = 1] - \Pr[\text{Hyb}_0(\mathcal{A}) = 1]| = \epsilon \geq \epsilon(\lambda)$ for some non-negligible ϵ . We construct a CPA-security adversary \mathcal{B} against Π_{HE} as follows:

1. On input 1^λ , algorithm \mathcal{B} runs \mathcal{A} to obtain the input length 1^n and a set $S \subseteq [n]$. Denote $S = \{i_1, \dots, i_t\} \subseteq [n]$, where $t = |S|$. Algorithm \mathcal{B} sends 1^n to the challenger as the input range. The challenger replies with a public key pk_0 .
2. Algorithm \mathcal{B} samples $(\text{sk}_1, \text{pk}_1) \leftarrow \text{HE.Gen}(1^\lambda, 1^n)$. It also samples the common reference strings $\text{crs}_{\text{Com}} \leftarrow \text{Com.Setup}(1^\lambda, 1^{\ell_{\text{ct}}(\lambda, n)}, 2n-1)$ and $(\text{crs}_{\text{BARG}}, \text{vk}_{\text{BARG}}) \leftarrow \text{Gen}'(1^\lambda, 1^{2n}, 1^s, 1^3)$. Finally, for $b \in \{0, 1\}$, it computes $\text{ct}_{\text{zero}}^{(b)} \leftarrow \text{HE.Enc}(\text{pk}_b, 0)$ for $b \in \{0, 1\}$.
3. Then, for each $i \in [n]$, algorithm \mathcal{B} does the following:
 - If $i \in S$, then it makes an encryption query on the pair $(0, 1)$ and receives the ciphertext ct_i^* . Algorithm \mathcal{B} sets $\text{ct}_i^{(0)} = \text{ct}_i^*$. If $i \notin S$, it sets $\text{ct}_i^{(0)} \leftarrow \text{HE.Enc}(\text{pk}_0, 0)$.
 - It computes $\text{ct}_i^{(1)} \leftarrow \text{HE.Enc}(\text{pk}_1, 0)$.
4. Finally, for $b \in \{0, 1\}$, compute $(\text{com}_{\text{hk}}^{(b)}, \sigma_{\text{hk},1}^{(b)}, \dots, \sigma_{\text{hk},n}^{(b)}) \leftarrow \text{Com.Commit}(\text{crs}_{\text{Com}}, (\text{ct}_1^{(b)}, \dots, \text{ct}_n^{(b)}))$.
5. Algorithm \mathcal{B} constructs hk and vk according to Eqs. (5.1) and (5.2) and gives (hk, vk) to \mathcal{A} .
6. Algorithm \mathcal{A} outputs a bit $b' \in \{0, 1\}$, which \mathcal{B} also outputs.

Observe that if the ct_i^* are encryptions of 0 then \mathcal{B} perfectly simulates Hyb_0 . If ct_i^* are encryptions of 1, then \mathcal{B} perfectly simulates Hyb_1 for \mathcal{A} . We conclude that the advantage of \mathcal{B} is ϵ . \square

Claim 5.21. *If Π_{HE} is CPA-secure, then there exists a negligible function $\text{negl}(\cdot)$ such that*

$$|\Pr[\text{Hyb}_2(\mathcal{A}) = 1] - \Pr[\text{Hyb}_1(\mathcal{A}) = 1]| = \text{negl}(\lambda).$$

Proof. Follows by an analogous argument as the proof of Claim 5.20. The only difference is the reduction algorithm \mathcal{B} sets pk_1 and the ciphertexts $\text{ct}_i^{(1)}$ for $i \in S$ to be the public key and challenge ciphertexts it receives for the CPA challenger, respectively. \square

Theorem 5.19 now follows by combining Claims 5.20 and 5.21. \square

5.1.3 Zero Fixing

In this section, we show that [Construction 5.2](#) satisfies zero-fixing security. In the zero-fixing game, the hash key in [Construction 5.2](#) is chosen to bind to a set S . This means that the ciphertext in the hash key associated with the set S are replaced by encryptions of 1. Suppose the adversary constructs an opening to 1 for an index $i \in S$. This means the adversary must have “used” the ciphertext associated with index i (which encrypts 1) when constructing the digest; formally, this will be enforced by the BARG. Since one of ciphertexts in the ciphertext tree is an encryption of 1, we can appeal to the predicate propagation property ([Theorem 5.9](#)) to argue that the root of the tree must encrypt a value that is non-zero. In this case, the extraction algorithm would declare the digest NotMatching and zero-fixing holds. We now give the formal argument:

Theorem 5.22 (Zero Fixing). *Suppose Π_{HE} satisfies perfect correctness and evaluation correctness, Π_{Com} is computationally binding, and Π_{BARG} satisfies set hiding with extraction, set hiding, and is somewhere extractable. Then [Construction 5.2](#) satisfies selective zero-fixing.*

Proof. We will leverage [Theorem 5.9](#). We start by defining a tree-based additive invariant P as follows. Let n be a power-of-two and take any index $i^* \in [n]$. We now inductively associate an interval $I_j^{(i^*)}$ with each node $j \in [2n - 1]$ of a complete binary tree with n leaves (indexed according to [Definition 5.1](#)) as follows:

- For $j \in [n]$, if $j = i^*$ then let $I_j^{(i^*)} = [1, 1]$. Otherwise, let $I_j^{(i^*)} = [0, 1]$.
- For an index $j \in [n + 1, 2n - 1]$, let j_L, j_R be the indices of the children of j according to [Definition 5.1](#). If $I_{j_L}^{(i^*)} = [a_L, b_L]$ and $I_{j_R}^{(i^*)} = [a_R, b_R]$, then define $I_j^{(i^*)} = [a_L + a_R, b_L + b_R] = I_{j_L}^{(i^*)} + I_{j_R}^{(i^*)}$, where we define interval addition to be component-wise addition: $[a_L, b_L] + [a_R, b_R] = [a_L + a_R, b_L + b_R]$.

We now define the interval-validity predicate P_{ValidInt} as follows: $P_{\text{ValidInt}}: \{0, 1\}^* \rightarrow \{0, 1\}$ as follows:

$$P_{\text{ValidInt}}(i^*, \text{ct}^{(0)}, \text{ct}^{(1)}, \text{sk}_0, \text{sk}_1, j) := \begin{cases} 1 & \text{HE.Dec}(\text{sk}_0, \text{ct}^{(0)}) \in I_j^{(i^*)} \\ 0 & \text{HE.Dec}(\text{sk}_0, \text{ct}^{(0)}) \notin I_j^{(i^*)}. \end{cases} \quad (5.11)$$

We now show that P_{ValidInt} is a tree-based additive invariant. We start by characterizing the intervals $I_j^{(i^*)}$ for all $j \in [2n - 1]$.

Claim 5.23. *For all $i^* \in [n]$ and any node $j \in [2n - 1]$ of height h in the binary tree, if i^* is in the subtree rooted at j then $I_j^{(i^*)} = [1, 2^h]$. Otherwise, $I_j^{(i^*)} = [0, 2^h]$.*

Proof. This follows by induction starting from the leaves. For every leaf node $j \in [n]$, the associated interval $I_j^{(i^*)}$ is $[0, 1]$ if $j \neq i^*$ and $[1, 1]$ if $j = i^*$. Thus, the claim holds for all of the leaf nodes. For the induction step, suppose j is a node of height $h + 1$. Let j_L, j_R be the indices of the children of j . If i^* is in the subtree of j , then it is either in the subtree of j_L or j_R but not both. By the induction hypothesis, $I_{j_L}^{(i^*)} + I_{j_R}^{(i^*)} = [0, 2^h] + [1, 2^h] = [1, 2^{h+1}]$. If i^* is not in the subtree of j , then by the induction hypothesis, $I_{j_L}^{(i^*)} + I_{j_R}^{(i^*)} = [0, 2^h] + [0, 2^h] = [0, 2^{h+1}]$. \square

Claim 5.24. *If Π_{HE} satisfies evaluation correctness, then P_{ValidInt} is a tree-based additive invariant.*

Proof. Take any $i^* \in [n]$, $(\text{sk}_0, \text{pk}_0)$, $(\text{sk}_1, \text{pk}_1)$ in the support of $\text{HE.Gen}(1^\lambda, 1^n)$, any triple of indices $j, j_L, j_R \in [2n - 1]$ where j_L, j_R are the indices of the children of j , and any set of ciphertexts $(\text{ct}_L^{(0)}, \text{ct}_L^{(1)})$, $(\text{ct}_R^{(0)}, \text{ct}_R^{(1)})$. Suppose j has height $h \leq \log n$. Let $\text{ct}_{\text{sum}}^{(0)} = \text{HE.Add}(\text{pk}_0, \text{ct}_L^{(0)}, \text{ct}_R^{(0)})$ and $\text{ct}_{\text{sum}}^{(1)} = \text{HE.Add}(\text{pk}_1, \text{ct}_L^{(1)}, \text{ct}_R^{(1)})$. Suppose

$$P_{\text{ValidInt}}(i^*, \text{ct}_L^{(0)}, \text{ct}_L^{(1)}, \text{sk}_0, \text{sk}_1, j_L) = 1 \quad \text{and} \quad P_{\text{ValidInt}}(i^*, \text{ct}_R^{(0)}, \text{ct}_R^{(1)}, \text{sk}_0, \text{sk}_1, j_R) = 1.$$

This means $\text{HE.Dec}(\text{sk}_0, \text{ct}_L^{(0)}) \in I_{j_L}^{(i^*)}$ and $\text{HE.Dec}(\text{sk}_0, \text{ct}_R^{(0)}) \in I_{j_R}^{(i^*)}$. By [Claim 5.23](#), this means

$$\text{HE.Dec}(\text{sk}_0, \text{ct}_L^{(0)}), \text{HE.Dec}(\text{sk}_0, \text{ct}_R^{(0)}) \in [0, 2^{h-1}].$$

Since $2^{h-1} \leq n/2$, we can appeal to evaluation correctness of Π_H and conclude that

$$\text{HE.Dec}(\text{sk}_0, \text{HE.Add}(\text{pk}_0, \text{ct}_L^{(0)}, \text{ct}_R^{(0)})) = \text{HE.Dec}(\text{sk}_0, \text{ct}_L^{(0)}) + \text{HE.Dec}(\text{sk}_0, \text{ct}_R^{(0)}) \in I_{j_L}^{(i^*)} + I_{j_R}^{(i^*)} = I_j^{(i^*)}.$$

We conclude that $P_{\text{ValidInt}}(i^*, \text{ct}_{\text{sum}}^{(0)}, \text{ct}_{\text{sum}}^{(1)}, \text{sk}_0, \text{sk}_1, j) = 1$. \square

Proof of Theorem 5.22. Returning now to the proof of Theorem 5.22, let \mathcal{A} be any efficient adversary for the zero-fixing security game. We start by defining a mapping `DeriveChal` as

$$\text{DeriveChal}(S, i) := (S, i) \mapsto (S, S, (i, 1)).$$

Let $\text{Expt} := \text{Expt}[P_{\text{ValidInt}}, \text{DeriveChal}]$ be the predicate propagation experiment from Definition 5.7. First, we claim that

$$\Pr[\text{ExptZF}_{\mathcal{A}}(\lambda) = 1] \leq \Pr[\text{Expt}(\mathcal{A}) = 1]. \quad (5.12)$$

By construction, the adversary's view in ExptZF and Expt is identical. It suffices to consider the outputs of the two experiments. Suppose $\text{ExptZF}_{\mathcal{A}}(\lambda) = 1$. This means the adversary \mathcal{A} outputs $\text{dig} = (\text{ct}_{\text{root}}^{(0)}, \text{ct}_{\text{root}}^{(1)}, \text{com}_0, \text{com}_1, \pi_{\text{dig}})$ and an opening π such that $\text{VerOpen}(\text{hk}, \text{dig}, i, 1, \pi) = 1$ and $\text{Extract}(\text{td}, \text{dig}) = \text{Matching}$. This means the following:

- By construction, $\text{Extract}(\text{td}, \text{dig})$ outputs `Matching` if $\text{HE.Dec}(\text{sk}_0, \text{ct}_{\text{root}}^{(0)}) = 0$. From Claim 5.24, we have that $I_{2n-1}^{(i^*)} = [1, n]$. Hence, in this case, $\text{HE.Dec}(\text{sk}_0, \text{ct}_{\text{root}}^{(0)}) \notin I_{2n-1}^{(i^*)}$, which implies that

$$P_{\text{ValidInt}}(i^*, \text{ct}_{\text{root}}^{(0)}, \text{ct}_{\text{root}}^{(1)}, \text{sk}_0, \text{sk}_1, 2n-1) = 0$$

in the predicate propagation experiment $\text{Expt}(\mathcal{A})$.

- By definition, VerOpen outputs 1 if $\text{Verify}'(\text{vk}_{\text{BARG}}, C_{i^*,1}, 2n-1, \pi) = 1$, where $i^* \in S$ is the index chosen by the adversary at the beginning of the (selective) zero-fixing game. By construction of `DeriveChal`, we have that $\text{idx} = (i^*, 1)$ in the execution of $\text{Expt}(\mathcal{A})$. This means that $\text{Verify}'(\text{vk}_{\text{BARG}}, C_{\text{idx}}, 2n-1, \pi) = 1$.

Since $P_{\text{ValidInt}}(i^*, \text{ct}_{\text{root}}^{(0)}, \text{ct}_{\text{root}}^{(1)}, \text{sk}_0, \text{sk}_1, 2n-1) = 0$ and $\text{Verify}'(\text{vk}_{\text{BARG}}, C_{\text{idx}}, 2n-1, \pi) = 1$, the predicate propagation experiment $\text{Expt}(\mathcal{A})$ also outputs 1. We now show using Theorem 5.9 that $\Pr[\text{Expt}(\mathcal{A}) = 1] \leq \text{negl}(\lambda)$. To leverage Theorem 5.9, we analyze the predicate propagation hybrid experiment $\text{Expt}_j := \text{Expt}_j[P_{\text{ValidInt}}, \text{DeriveChal}]$ from Definition 5.8.

Claim 5.25. *If Π_{HE} is perfectly correct and Π_{Com} satisfies computational binding, then there exists a negligible function $\text{negl}(\cdot)$ such that for all $j \in [n]$, it holds that $\Pr[\text{Expt}_j(\mathcal{A}) = 1] = \text{negl}(\lambda)$.*

Proof. Suppose there exists some $j \in [n]$ where $\Pr[\text{Expt}_j(\mathcal{A}) = 1] \geq \varepsilon(\lambda)$ for some non-negligible ε . We use \mathcal{A} to construct an adversary \mathcal{B} that breaks computational binding of Π_{Com} :

1. On input the security parameter 1^λ , algorithm \mathcal{B} runs algorithm \mathcal{A} to obtain the input length 1^n , the set $S \subseteq [n]$, and an index $i^* \in S$.
2. Algorithm \mathcal{B} outputs the block length $1^{\ell_{\text{ct}}(\lambda, n)}$ and the vector length $2n-1$ to the challenger. The challenger responds with crs_{Com} .
3. Algorithm \mathcal{B} computes $(S, S, (i^*, 1)) \leftarrow \text{DeriveChal}(S, i^*)$. It then samples the following components:
 - Sample $(\text{sk}_0, \text{pk}_0) \leftarrow \text{HE.Gen}(1^\lambda, 1^n)$ and $(\text{sk}_1, \text{pk}_1) \leftarrow \text{HE.Gen}(1^\lambda, 1^n)$.
 - Sample $(\text{crs}_{\text{BARG}}, \text{vk}_{\text{BARG}}, \text{td}_{\text{BARG}}) \leftarrow \text{TrapGen}'(1^\lambda, 1^{2n-1}, 1^s, 1^3, \{j\})$.
 - For each $b \in \{0, 1\}$, sample $\text{ct}_{\text{zero}}^{(b)} \leftarrow \text{HE.Enc}(\text{pk}_b, 0)$. Then, for each $i \in [n]$ and $b \in \{0, 1\}$, if $i \in S$, sample $\text{ct}_i^{(b)} \leftarrow \text{HE.Enc}(\text{pk}_b, 1)$; otherwise, if $i \notin S$, sample $\text{ct}_i^{(b)} \leftarrow \text{HE.Enc}(\text{pk}_b, 0)$.
 - For each $b \in \{0, 1\}$, let $(\text{com}_{\text{hk}}^{(b)}, \sigma_{\text{hk},1}^{(b)}, \dots, \sigma_{\text{hk},n}^{(b)}) \leftarrow \text{Com.Commit}(\text{crs}_{\text{Com}}, (\text{ct}_1^{(b)}, \dots, \text{ct}_n^{(b)}))$

4. Algorithm \mathcal{B} constructs hk and vk according to Eqs. (5.1) and (5.2):

$$\begin{aligned} \text{hk} &= \left(\text{crs}_{\text{Com}}, \text{crs}_{\text{BARG}}, \{ \text{pk}_b, \text{ct}_{\text{zero}}^{(b)}, \text{ct}_1^{(b)}, \dots, \text{ct}_n^{(b)}, \sigma_{\text{hk},1}^{(b)}, \dots, \sigma_{\text{hk},n}^{(b)} \}_{b \in \{0,1\}} \right) \\ \text{vk} &= (\text{crs}_{\text{Com}}, \text{vk}_{\text{BARG}}, \text{pk}_0, \text{pk}_1, \text{ct}_{\text{zero}}^{(0)}, \text{ct}_{\text{zero}}^{(1)}, \text{com}_{\text{hk}}^{(0)}, \text{com}_{\text{hk}}^{(1)}) \end{aligned}$$

Algorithm \mathcal{B} gives (hk, vk) to \mathcal{A} .

5. Algorithm \mathcal{A} outputs a digest $\text{dig} = (\text{ct}_{\text{root}}^{(0)}, \text{ct}_{\text{root}}^{(1)}, \text{com}_0, \text{com}_1, \pi_{\text{dig}})$ and a proof π .

6. Algorithm \mathcal{B} extracts $(\hat{v}_j^{(0)}, \hat{v}_j^{(1)}, \sigma_j^{(0)}, \sigma_j^{(1)}, \tilde{w}_j) \leftarrow \text{Extract}'(\text{td}_{\text{BARG}}, \pi, j)$ and parses the auxiliary witness $\tilde{w}_j = (\hat{\text{ct}}^{(0)}, \hat{\text{ct}}^{(1)}, \sigma_{\text{hk}}^{(0)}, \sigma_{\text{hk}}^{(1)})$.

7. Output the commitment $\text{com}_{\text{hk}}^{(0)}$, the index j , and the value-opening pairs $(\text{ct}_j^{(0)}, \sigma_{\text{hk},j}^{(0)})$ and $(\hat{\text{ct}}^{(0)}, \sigma_{\text{hk}}^{(0)})$.

By construction, the challenger samples $\text{crs}_{\text{Com}} \leftarrow \text{Com.Setup}(1^\lambda, 1^{\ell_{\text{ct}}(\lambda,n)}, 2n-1)$, which matches the specification in Expt_j . This, algorithm \mathcal{B} perfectly simulates an execution of Expt_j for \mathcal{A} . By assumption, with probability ε , algorithm \mathcal{A} outputs dig and π such that the experiment outputs 1. This means the following conditions hold:

$$C_{i^*,1}(j, (\hat{v}_j^{(0)}, \hat{v}_j^{(1)}, \sigma_j^{(0)}, \sigma_j^{(1)}, \tilde{w}_j)) = 1 \quad \text{and} \quad P_{\text{ValidInt}}(i^*, \hat{v}_j^{(0)}, \hat{v}_j^{(1)}, \text{sk}_0, \text{sk}_1, j) = 0.$$

By definition of $C_{i^*,1}$ and using the fact that $j \in [n]$, this means

$$\text{Com.Verify}(\text{crs}_{\text{com}}, \text{com}_{\text{hk}}^{(0)}, j, \hat{\text{ct}}^{(0)}, \sigma_{\text{hk}}^{(0)}) = 1 \quad \text{and} \quad \hat{v}_j^{(0)} \in \{\text{ct}_{\text{zero}}^{(0)}, \hat{\text{ct}}^{(0)}\}.$$

If $j = i^*$, we additionally have that $\hat{v}_j^{(0)} = \hat{\text{ct}}^{(0)}$. Next, by correctness of Π_{Com} ,

$$\text{Com.Verify}(\text{crs}_{\text{com}}, \text{com}_{\text{hk}}^{(0)}, j, \text{ct}_j^{(0)}, \sigma_{\text{hk},j}^{(0)}) = 1.$$

It suffices to argue that $\text{ct}_j^{(0)} \neq \hat{\text{ct}}^{(0)}$. We consider two cases:

- Suppose $j = i^*$. Recall that in this case, $\hat{v}_j^{(0)} = \hat{\text{ct}}^{(0)}$. By [Claim 5.24](#), we have that $I_j^{(i^*)} = [1, 1]$. Since $P_{\text{ValidInt}}(i^*, \hat{v}_j^{(0)}, \hat{v}_j^{(1)}, \text{sk}_0, \text{sk}_1, j) = 0$, this means that $\text{HE.Dec}(\text{sk}_0, \hat{v}_j^{(0)}) \notin I_j^{(i^*)}$. By definition of $I_j^{(i^*)}$, this implies $\text{HE.Dec}(\text{sk}_0, \hat{v}_j^{(0)}) \neq 1$. Next, algorithm \mathcal{B} constructs $\text{ct}_j^{(0)}$ to be an encryption of 1 (since $j = i^* \in S$). By perfect correctness of Π_{HE} , this means that $\hat{\text{ct}}^{(0)} = \hat{v}_j^{(0)} \neq \text{ct}_j^{(0)}$.
- Suppose $j \neq i^*$. By [Claim 5.24](#), we have that $I_j^{(i^*)} = [0, 1]$. By the same reasoning as in the previous case, this means that $\text{HE.Dec}(\text{sk}_0, \hat{v}_j^{(0)}) \notin I_j^{(i^*)}$. In other words, $\text{HE.Dec}(\text{sk}_0, \hat{v}_j^{(0)}) \notin \{0, 1\}$. By construction, $\text{ct}_{\text{zero}}^{(0)}$ is an encryption of 0, so by perfect correctness of Π_{HE} , we have that $\hat{v}_j^{(0)} \neq \text{ct}_{\text{zero}}^{(0)}$. Hence, it must be the case that $\hat{v}_j^{(0)} = \hat{\text{ct}}^{(0)}$. Next, the ciphertext $\text{ct}_j^{(0)}$ is an encryption of either 0 or 1, so we conclude that $\hat{\text{ct}}^{(0)} = \hat{v}_j^{(0)} \neq \text{ct}_j^{(0)}$.

In both cases, we conclude that $\hat{\text{ct}}^{(0)} \neq \text{ct}_j^{(0)}$. In this case, algorithm \mathcal{B} successfully opens $\text{com}_{\text{hk}}^{(0)}$ to two distinct values $\text{ct}_j^{(0)} \neq \hat{\text{ct}}^{(0)}$. Thus algorithm \mathcal{B} breaks binding with the same advantage ε . \square

Since for all $j \in [n]$, it holds that $\Pr[\text{Expt}_j(\mathcal{A}) = 1] \leq \text{negl}(\lambda)$, we can invoke [Theorem 5.9](#) to conclude that $\Pr[\text{Expt}(\mathcal{A}) = 1] \leq \text{negl}(\lambda)$. Zero-fixing security now follows via [Eq. \(5.12\)](#). \square

5.1.4 Extractor Validity

In this section, we show that [Construction 5.2](#) satisfies extractor validity. In the extractor validity game, the hash key is sampled to be zero-fixing on the empty set \emptyset , and the goal of the adversary is to produce a valid, but *non-matching* digest. In this setting, the ciphertexts in the hash key are all encryptions of 0. In order to break the extractor validity property, the adversary needs to produce a root ciphertext that encrypts a non-zero value, and yet, still argue that the root ciphertext was derived by summing a collection of ciphertexts that each encryption 0. The latter is ensured by security of the BARG, and specifically the predicate propagation theorem ([Theorem 5.9](#)). We give the formal theorem statement and proof below:

Theorem 5.26. *If Π_{HE} satisfies perfect correctness and evaluation correctness, Π_{Com} is computationally binding, and Π_{BARG} satisfies set hiding with extraction, and somewhere extractability, then [Construction 5.2](#) satisfies extractor validity.*

Proof. Similar to the proof of [Theorem 5.22](#), we leverage [Theorem 5.9](#). We start by defining a tree-based additive invariant P as follows. Define the “matching” predicate $P_{\text{Matching}}: \{0, 1\}^* \rightarrow \{0, 1\}$ as follows:

$$P_{\text{Matching}}(i^*, \text{ct}^{(0)}, \text{ct}^{(1)}, \text{sk}_0, \text{sk}_1, j) := \begin{cases} 1 & \text{HE.Dec}(\text{sk}_0, \text{ct}^{(0)}) = 0 \\ 0 & \text{HE.Dec}(\text{sk}_0, \text{ct}^{(0)}) \neq 0. \end{cases}$$

We first show that P_{Matching} is a tree-based additive invariant.

Claim 5.27. *If Π_{HE} satisfies evaluation correctness, then P_{Matching} is a tree-based additive invariant.*

Proof. Take any $i^* \in [n]$, any $(\text{sk}_0, \text{pk}_0), (\text{sk}_1, \text{pk}_1)$ in the support of $\text{HE.Gen}(1^\lambda, 1^n)$, any triple of indices $j, j_L, j_R \in [2n - 1]$ where j_L, j_R are the indices of the children of j , and any set of ciphertexts $(\text{ct}_L^{(0)}, \text{ct}_L^{(1)}), (\text{ct}_R^{(0)}, \text{ct}_R^{(1)})$. Suppose

$$P_{\text{Matching}}(i^*, \text{ct}_L^{(0)}, \text{ct}_L^{(1)}, \text{sk}_0, \text{sk}_1, j_L) = P_{\text{Matching}}(i^*, \text{ct}_R^{(0)}, \text{ct}_R^{(1)}, \text{sk}_0, \text{sk}_1, j_R) = 1.$$

This means $\text{HE.Dec}(\text{sk}_0, \text{ct}_L^{(0)}) = 0$ and $\text{HE.Dec}(\text{sk}_0, \text{ct}_R^{(0)}) = 0$. Let $\text{ct}_{\text{sum}}^{(b)} = \text{HE.Add}(\text{pk}_b, \text{ct}_L^{(b)}, \text{ct}_R^{(b)})$ for $b \in \{0, 1\}$. By evaluation correctness of Π_{HE} , we have $\text{HE.Dec}(\text{sk}_0, \text{ct}_{\text{sum}}^{(0)}) = 0$ and so by definition

$$P_{\text{Matching}}(i^*, \text{ct}_{\text{sum}}^{(0)}, \text{ct}_{\text{sum}}^{(1)}, \text{sk}_0, \text{sk}_1, j) = 1.$$

□

Let \mathcal{A} be an efficient adversary for the extractor-validity game. Define the mapping DeriveChal as

$$\text{DeriveChal}(S, i) := (S, i) \mapsto (\emptyset, \emptyset, \perp).$$

Let $\text{Expt} := \text{Expt}[P_{\text{Matching}}, \text{DeriveChal}]$ be the predicate propagation experiment from [Definition 5.7](#). First, we claim that we can use \mathcal{A} to construct an adversary \mathcal{A}' such that

$$\Pr[\text{ExptEV}_{\mathcal{A}}(\lambda) = 1] \leq \Pr[\text{Expt}(\mathcal{A}') = 1]. \quad (5.13)$$

Algorithm \mathcal{A}' works as follows:

1. On input the security parameter 1^λ , algorithm \mathcal{A}' runs \mathcal{A} on the same security parameter. Algorithm \mathcal{A} outputs an input length 1^n . Algorithm \mathcal{A}' outputs the input length 1^n , the set $S = \emptyset$, and the index $i^* = \perp$.
2. The challenger replies with (hk, vk) which \mathcal{A}' forwards to \mathcal{A} .
3. Algorithm \mathcal{A} outputs a digest $\text{dig} = (\text{ct}_{\text{root}}^{(0)}, \text{ct}_{\text{root}}^{(1)}, \text{com}_0, \text{com}_1, \pi_{\text{dig}})$. Algorithm \mathcal{A}' outputs the same digest dig and $\pi = \pi_{\text{dig}}$.

We now show that Eq. (5.13) holds. By construction, the pair (hk, vk) sampled by the challenger are distributed according to the real setup algorithm. Thus, algorithm \mathcal{A} perfectly simulates an execution of $\text{ExptEV}_{\mathcal{A}}(\lambda)$ for adversary \mathcal{A} . Thus, with probability $\Pr[\text{ExptEV}_{\mathcal{A}}(\lambda) = 1]$, algorithm \mathcal{A} outputs a digest dig where $\text{Extract}(\text{td}, \text{dig}) = \text{NotMatching}$ and $\text{ValidateDigest}(\text{hk}, \text{dig}) = 1$. This means the following:

- By construction, $\text{Extract}(\text{td}, \text{dig})$ outputs NotMatching if $\text{HE.Dec}(\text{sk}_0, \text{ct}_{\text{root}}^{(0)}) \neq 0$. By construction of P_{Matching} , this means $P_{\text{Matching}}(i^*, \text{ct}_{\text{root}}^{(0)}, \text{ct}_{\text{root}}^{(1)}, \text{sk}_0, \text{sk}_1, 2n - 1) = 0$.
- Next, ValidateDigest outputs 1 if $\text{Verify}'(\text{vk}_{\text{BARG}}, C_{\perp}, 2n - 1, \pi_{\text{dig}}) = 1$. By construction of DeriveChal , we have that $\text{idx} = \perp$ in the execution of $\text{Expt}(\mathcal{A})$, so this means that $\text{Verify}'(\text{vk}_{\text{BARG}}, C_{\text{idx}}, 2n - 1, \pi_{\text{dig}}) = 1$.

Since $P_{\text{Matching}}(i^*, \text{ct}_{\text{root}}^{(0)}, \text{ct}_{\text{root}}^{(1)}, \text{sk}_0, \text{sk}_1, 2n - 1) = 0$ and $\text{Verify}'(\text{vk}_{\text{BARG}}, C_{\text{idx}}, 2n - 1, \pi_{\text{dig}}) = 1$, the predicate propagation experiment $\text{Expt}(\mathcal{A}')$ also outputs 1. Hence, we conclude that $\Pr[\text{Expt}(\mathcal{A}') = 1] \geq \Pr[\text{ExptEV}_{\mathcal{A}}(\lambda) = 1]$. To complete the proof, we now show using Theorem 5.9 that $\Pr[\text{Expt}(\mathcal{A}') = 1] \leq \text{negl}(\lambda)$. To leverage Theorem 5.9, we analyze the predicate propagation hybrid experiment $\text{Expt}_j := \text{Expt}_j[P_{\text{Matching}}, \text{DeriveChal}]$ from Definition 5.8.

Claim 5.28. *If Π_{HE} is perfectly correct and Π_{Com} satisfies computational binding, then there exists a negligible function $\text{negl}(\cdot)$ such that for all $j \in [n]$, it holds that $\Pr[\text{Expt}_j(\mathcal{A}') = 1] = \text{negl}(\lambda)$.*

Proof. Suppose there exists some $j \in [n]$ where $\Pr[\text{Expt}_j(\mathcal{A}') = 1] \geq \varepsilon(\lambda)$ for some non-negligible ε . We use \mathcal{A}' to construct an adversary \mathcal{B} that breaks computational binding of Π_{Com} .

1. On input the security parameter 1^λ , algorithm \mathcal{B} runs algorithm \mathcal{A}' to obtain the input length 1^n , the set $S = \emptyset$, and the index $i^* = \perp$.
2. Algorithm \mathcal{B} outputs the block length $1^{\ell_{\text{ct}}(\lambda, n)}$ and the vector length $2n - 1$ to the challenger. The challenger responds with crs_{Com} .
3. Algorithm \mathcal{B} computes $(\emptyset, \emptyset, \perp) \leftarrow \text{DeriveChal}(S, i^*)$. It then samples the following components:
 - Sample $(\text{sk}_0, \text{pk}_0) \leftarrow \text{HE.Gen}(1^\lambda, 1^n)$ and $(\text{sk}_1, \text{pk}_1) \leftarrow \text{HE.Gen}(1^\lambda, 1^n)$.
 - Sample $(\text{crs}_{\text{BARG}}, \text{vk}_{\text{BARG}}, \text{td}_{\text{BARG}}) \leftarrow \text{TrapGen}'(1^\lambda, 1^{2n-1}, 1^s, 1^3, \{j\})$.
 - For each $b \in \{0, 1\}$, sample $\text{ct}_{\text{zero}}^{(b)} \leftarrow \text{HE.Enc}(\text{pk}_b, 0)$. Then, for each $i \in [n]$ and $b \in \{0, 1\}$, sample $\text{ct}_i^{(b)} \leftarrow \text{HE.Enc}(\text{pk}_b, 0)$.
 - For each $b \in \{0, 1\}$, let $(\text{com}_{\text{hk}}^{(b)}, \sigma_{\text{hk},1}^{(b)}, \dots, \sigma_{\text{hk},n}^{(b)}) \leftarrow \text{Com.Commit}(\text{crs}_{\text{Com}}, (\text{ct}_1^{(b)}, \dots, \text{ct}_n^{(b)}))$.
4. Algorithm \mathcal{B} constructs hk and vk according to Eqs. (5.1) and (5.2):

$$\begin{aligned} \text{hk} &= \left(\text{crs}_{\text{Com}}, \text{crs}_{\text{BARG}}, \left\{ \text{pk}_b, \text{ct}_{\text{zero}}^{(b)}, \text{ct}_1^{(b)}, \dots, \text{ct}_n^{(b)}, \sigma_{\text{hk},1}^{(b)}, \dots, \sigma_{\text{hk},n}^{(b)} \right\}_{b \in \{0,1\}} \right) \\ \text{vk} &= (\text{crs}_{\text{Com}}, \text{vk}_{\text{BARG}}, \text{pk}_0, \text{pk}_1, \text{ct}_{\text{zero}}^{(0)}, \text{ct}_{\text{zero}}^{(1)}, \text{com}_{\text{hk}}^{(0)}, \text{com}_{\text{hk}}^{(1)}) \end{aligned}$$

Algorithm \mathcal{B} gives (hk, vk) to \mathcal{A}' .

5. Algorithm \mathcal{A}' outputs a digest $\text{dig} = (\text{ct}_{\text{root}}^{(0)}, \text{ct}_{\text{root}}^{(1)}, \text{com}_0, \text{com}_1, \pi_{\text{dig}})$ and a proof $\pi = \pi_{\text{dig}}$.
6. Algorithm \mathcal{B} extracts $(\hat{\nu}_j^{(0)}, \hat{\nu}_j^{(1)}, \sigma_j^{(0)}, \sigma_j^{(1)}, \tilde{w}_j) \leftarrow \text{Extract}'(\text{td}_{\text{BARG}}, \pi, j)$ and parses the auxiliary witness $\tilde{w}_j = (\hat{\text{ct}}^{(0)}, \hat{\text{ct}}^{(1)}, \sigma_{\text{hk}}^{(0)}, \sigma_{\text{hk}}^{(1)})$.
7. Output the commitment $\text{com}_{\text{hk}}^{(0)}$, the index j , and the value-opening pairs $(\text{ct}_j^{(0)}, \sigma_{\text{hk},j}^{(0)})$ and $(\hat{\text{ct}}^{(0)}, \sigma_{\text{hk}}^{(0)})$.

By construction, the challenger samples $\text{crs}_{\text{Com}} \leftarrow \text{Com.Setup}(1^\lambda, 1^{\ell_{\text{ct}}(\lambda, n)}, 2n - 1)$, which matches the specification in Expt_j . Thus, algorithm \mathcal{B} perfectly simulates an execution of Expt_j for \mathcal{A}' . By assumption, with probability ϵ , algorithm \mathcal{A}' outputs dig and π such that the experiment outputs 1. This means the following conditions hold:

$$C_\perp(j, (\hat{\nu}_j^{(0)}, \hat{\nu}_j^{(1)}, \sigma_j^{(0)}, \sigma_j^{(1)}, \tilde{w}_j)) = 1 \quad \text{and} \quad P_{\text{Matching}}(i^*, \hat{\nu}_j^{(0)}, \hat{\nu}_j^{(1)}, \text{sk}_0, \text{sk}_1, j) = 0.$$

By definition of C_\perp and using the fact that $j \in [n]$, this means

$$\text{Com.Verify}(\text{crs}_{\text{Com}}, \text{com}_{\text{hk}}^{(0)}, j, \hat{\text{ct}}^{(0)}, \sigma_{\text{hk}}^{(0)}) = 1 \quad \text{and} \quad \hat{\nu}_j^{(0)} \in \{\text{ct}_{\text{zero}}^{(0)}, \hat{\text{ct}}^{(0)}\}.$$

Next, by correctness of Π_{Com} ,

$$\text{Com.Verify}(\text{crs}_{\text{Com}}, \text{com}_{\text{hk}}^{(0)}, j, \text{ct}_j^{(0)}, \sigma_{\text{hk}, j}^{(0)}) = 1.$$

It suffices to argue that $\text{ct}_j^{(0)} \neq \hat{\text{ct}}^{(0)}$. Since $P_{\text{Matching}}(i^*, \hat{\nu}_j^{(0)}, \hat{\nu}_j^{(1)}, \text{sk}_0, \text{sk}_1, j) = 0$, we have $\text{HE.Dec}(\text{sk}_0, \hat{\nu}_j^{(0)}) \neq 0$. Since $\text{ct}_{\text{zero}}^{(0)}$ is an encryption of 0, we can appeal to perfect correctness of Π_{HE} to conclude that $\hat{\nu}_j^{(0)} \neq \text{ct}_{\text{zero}}^{(0)}$. This means that $\hat{\nu}_j^{(0)} = \hat{\text{ct}}^{(0)}$. Moreover, $\text{ct}_j^{(0)}$ is also an encryption of 0, so again by perfect correctness of the encryption scheme, we can conclude that $\text{ct}_j^{(0)} \neq \hat{\nu}_j^{(0)} = \hat{\text{ct}}^{(0)}$. In this case, algorithm \mathcal{B} successfully opens $\text{com}_{\text{hk}}^{(0)}$ to two distinct values $\text{ct}_j^{(0)} \neq \hat{\text{ct}}^{(0)}$. Thus algorithm \mathcal{B} breaks binding with the same advantage ϵ . \square

Since for all $j \in [n]$, it holds that $\Pr[\text{Expt}_j(\mathcal{A}') = 1] \leq \text{negl}(\lambda)$, we can invoke [Theorem 5.9](#) to conclude that $\Pr[\text{Expt}(\mathcal{A}') = 1] \leq \text{negl}(\lambda)$. Extractor-validity security now follows via [Eq. \(5.13\)](#). \square

5.1.5 Index Hiding with Extracted Guess

In this section, we show that [Construction 5.2](#) satisfies the index hiding with extracted guess property. The challenge in this reduction is we need to switch from an encryption of 0 to an encryption of 1 (in the hash key) while retaining the ability to decide whether the digest is “Matching” or not (which in the real scheme, requires knowledge of the secret key for the underlying encryption scheme). As described in [Section 1.2](#), we solve this by adopting a Naor-Yung proof strategy.

Theorem 5.29. *If Π_{HE} satisfies perfect correctness, evaluation correctness, and CPA-security, Π_{Com} is computationally binding and Π_{BARG} satisfies set hiding with extraction, set hiding, and is somewhere extractable, then [Construction 5.2](#) satisfies index hiding with extracted guess.*

Proof. Let \mathcal{A} be an efficient adversary for the index hiding with extracted guess security game. We define a sequence of hybrid experiments:

- Hyb_0 : This is $\text{ExptIHE}_{\mathcal{A}}(\lambda, 0)$. Specifically, the game proceeds as follows:
 1. On input the security parameter 1^λ , algorithm \mathcal{A} outputs the input length 1^n , a set $S \subseteq [n]$, and an index $i^* \in S$.
 2. The challenger now samples the following quantities as in Setup:
 - Sample $(\text{sk}_0, \text{pk}_0) \leftarrow \text{HE.Gen}(1^\lambda, 1^n)$ and $(\text{sk}_1, \text{pk}_1) \leftarrow \text{HE.Gen}(1^\lambda, 1^n)$.
 - Sample $\text{crs}_{\text{Com}} \leftarrow \text{Com.Setup}(1^\lambda, 1^{\ell_{\text{ct}}(\lambda, n)}, 2n - 1)$.
 - Sample $(\text{crs}_{\text{BARG}}, \text{vk}_{\text{BARG}}) \leftarrow \text{Gen}'(1^\lambda, 1^{2n-1}, 1^s, 1^3)$, where s is a bound on the size of the circuit computing the index relation from [Fig. 2](#).
 - For each $b \in \{0, 1\}$, sample $\text{ct}_{\text{zero}}^{(b)} \leftarrow \text{HE.Enc}(\text{pk}_b, 0)$. Then, for each $i \in [n]$ and $b \in \{0, 1\}$, if $i \in S \setminus \{i^*\}$, sample $\text{ct}_i^{(b)} \leftarrow \text{HE.Enc}(\text{pk}_b, 1)$; otherwise sample $\text{ct}_i^{(b)} \leftarrow \text{HE.Enc}(\text{pk}_b, 0)$.
 - For each $b \in \{0, 1\}$, let $(\text{com}_{\text{hk}}^{(b)}, \sigma_{\text{hk}, 1}^{(b)}, \dots, \sigma_{\text{hk}, n}^{(b)}) \leftarrow \text{Com.Commit}(\text{crs}_{\text{Com}}, (\text{ct}_1^{(b)}, \dots, \text{ct}_n^{(b)}))$

3. The challenger constructs hk and vk according to Eqs. (5.1) and (5.2):

$$\begin{aligned} hk &= \left(crs_{Com}, crs_{BARG}, \{pk_b, ct_{zero}^{(b)}, ct_1^{(b)}, \dots, ct_n^{(b)}, \sigma_{hk,1}^{(b)}, \dots, \sigma_{hk,n}^{(b)}\}_{b \in \{0,1\}} \right) \\ vk &= (crs_{Com}, vk_{BARG}, pk_0, pk_1, ct_{zero}^{(0)}, ct_{zero}^{(1)}, com_{hk}^{(0)}, com_{hk}^{(1)}) \end{aligned}$$

The challenger gives (hk, vk) to \mathcal{A} .

4. Algorithm \mathcal{A} outputs a digest $dig = (ct_{root}^{(0)}, ct_{root}^{(1)}, com_0, com_1, \pi_{dig})$ and an opening π .
5. The output of the experiment is 1 if

$$\text{Verify}'(vk_{BARG}, C_{i^*,0}, 2n-1, \pi) = 1 \quad \text{and} \quad \text{HE.Dec}(sk_0, ct_{root}^{(0)}) = 0.$$

Otherwise, the output is 0.

- Hyb_1 : Same as Hyb_0 , except the challenger samples $ct_{i^*}^{(1)} \leftarrow \text{HE.Enc}(pk_1, 1)$.
- Hyb_2 : Same as Hyb_1 , except the output of the experiment is 1 if

$$\text{Verify}'(vk_{BARG}, C_{i^*,0}, 2n-1, \pi) = 1 \quad \text{and} \quad \text{HE.Dec}(sk_1, ct_{root}^{(1)}) = 0.$$

Notably, the challenger's behavior in this experiment does *not* depend on sk_0 .

- Hyb_3 : Same as Hyb_2 , except the output of the challenger samples $ct_{i^*}^{(0)} \leftarrow \text{HE.Enc}(pk_0, 1)$.
- Hyb_4 : Same as Hyb_3 , except the output of the experiment is 1 if

$$\text{Verify}'(vk_{BARG}, C_{i^*,0}, 2n-1, \pi) = 1 \quad \text{and} \quad \text{HE.Dec}(sk_0, ct_{root}^{(0)}) = 0.$$

This is experiment $\text{ExptIHE}_{\mathcal{A}}(\lambda, 1)$.

We write $\text{Hyb}_i(\mathcal{A})$ to denote the output of experiment of Hyb_i with adversary \mathcal{A} . We now analyze each pair of hybrid experiments.

Claim 5.30. *If Π_{HE} is CPA-secure, then there exists a negligible function $\text{negl}(\cdot)$ such that*

$$|\Pr[\text{Hyb}_1(\mathcal{A}) = 1] - \Pr[\text{Hyb}_0(\mathcal{A}) = 1]| = \text{negl}(\lambda).$$

Proof. Suppose $|\Pr[\text{Hyb}_1(\mathcal{A}) = 1] - \Pr[\text{Hyb}_0(\mathcal{A}) = 1]| \geq \varepsilon(\lambda)$ for some non-negligible ε . We use \mathcal{A} to construct an algorithm \mathcal{B} that breaks CPA security of Π_{HE} :

1. On input the security parameter 1^λ , algorithm \mathcal{B} runs algorithm \mathcal{A} on the same input to obtain the input length 1^n , the set $S \subseteq [n]$, and an index $i^* \in S$.
2. Algorithm \mathcal{B} sends 1^n as the input range. The challenger replies with a public key pk_1 .
3. Algorithm \mathcal{B} now samples the following components:
 - Sample $(sk_0, pk_0) \leftarrow \text{HE.Gen}(1^\lambda, 1^n)$.
 - Sample $(crs_{BARG}, vk_{BARG}, td_{BARG}) \leftarrow \text{TrapGen}'(1^\lambda, 1^{2n-1}, 1^S, 1^3, \{j\})$.
 - For each $b \in \{0, 1\}$, sample $ct_{zero}^{(b)} \leftarrow \text{HE.Enc}(pk_b, 0)$.
 - For each $i \in [n] \setminus \{i^*\}$ and $b \in \{0, 1\}$, if $i \in S$, sample $ct_i^{(b)} \leftarrow \text{HE.Enc}(pk_b, 1)$; otherwise, if $i \notin S$, sample $ct_i^{(b)} \leftarrow \text{HE.Enc}(pk_b, 0)$.
4. Algorithm \mathcal{B} makes an encryption query on the pair $(0, 1)$. The challenger replies with a ciphertext $ct_{i^*}^{(1)}$. Algorithm \mathcal{B} also computes $ct_{i^*}^{(0)} \leftarrow \text{HE.Enc}(pk_0, 0)$.

5. Finally, for each $b \in \{0, 1\}$, let $(\text{com}_{\text{hk}}^{(b)}, \sigma_{\text{hk},1}^{(b)}, \dots, \sigma_{\text{hk},n}^{(b)}) \leftarrow \text{Com.Commit}(\text{crs}_{\text{Com}}, (\text{ct}_1^{(b)}, \dots, \text{ct}_n^{(b)}))$. Next, algorithm \mathcal{B} constructs hk and vk according to Eqs. (5.1) and (5.2):

$$\begin{aligned} \text{hk} &= \left(\text{crs}_{\text{Com}}, \text{crs}_{\text{BARG}}, \{ \text{pk}_b, \text{ct}_{\text{zero}}^{(b)}, \text{ct}_1^{(b)}, \dots, \text{ct}_n^{(b)}, \sigma_{\text{hk},1}^{(b)}, \dots, \sigma_{\text{hk},n}^{(b)} \}_{b \in \{0,1\}} \right) \\ \text{vk} &= (\text{crs}_{\text{Com}}, \text{vk}_{\text{BARG}}, \text{pk}_0, \text{pk}_1, \text{ct}_{\text{zero}}^{(0)}, \text{ct}_{\text{zero}}^{(1)}, \text{com}_{\text{hk}}^{(0)}, \text{com}_{\text{hk}}^{(1)}) \end{aligned}$$

Algorithm \mathcal{B} gives (hk, vk) to \mathcal{A} .

6. Algorithm \mathcal{A} outputs a digest $\text{dig} = (\text{ct}_{\text{root}}^{(0)}, \text{ct}_{\text{root}}^{(1)}, \text{com}_0, \text{com}_1, \pi_{\text{dig}})$ and a proof π .
7. Algorithm \mathcal{B} outputs 1 if

$$\text{Verify}'(\text{vk}_{\text{BARG}}, C_{i^*,0}, 2n-1, \pi) = 1 \quad \text{and} \quad \text{HE.Dec}(\text{sk}_0, \text{ct}_{\text{root}}^{(0)}) = 0,$$

where $C_{i^*,0}$ is the circuit computing the relation from Fig. 2 (which is a function of the components from hk and dig).

Observe that if $\text{ct}_{i^*}^{(1)}$ is an encryption of 0 (under pk_1), then algorithm \mathcal{B} perfectly simulates Hyb_0 for \mathcal{A} . Alternatively, if $\text{ct}_{i^*}^{(1)}$ is an encryption of 1 (under pk_1), then algorithm \mathcal{B} perfectly simulates Hyb_1 for \mathcal{A} . We conclude that the advantage of \mathcal{B} is ε . \square

Claim 5.31. *If Π_{HE} is perfectly correct and satisfies evaluation correctness, Π_{Com} is computationally binding, Π_{BARG} satisfies set hiding with extraction, set hiding, and is somewhere extractable, then there exists a negligible function $\text{negl}(\cdot)$ such that $|\Pr[\text{Hyb}_2(\mathcal{A}) = 1] - \Pr[\text{Hyb}_1(\mathcal{A}) = 1]| = \text{negl}(\lambda)$.*

Proof. By construction, the only difference between the execution of Hyb_1 and Hyb_2 is the output condition. Let E be the following event in an execution of Hyb_1 and Hyb_2 :

$$\text{Verify}'(\text{vk}_{\text{BARG}}, C_{i^*,0}, 2n-1, \pi) = 1 \quad \text{and} \quad \text{HE.Dec}(\text{sk}_0, \text{ct}_{\text{root}}^{(0)}) \neq \text{HE.Dec}(\text{sk}_1, \text{ct}_{\text{root}}^{(1)}). \quad (5.14)$$

Observe that if E does not occur, then the output of Hyb_1 and Hyb_2 is identical. This means that

$$|\Pr[\text{Hyb}_2(\mathcal{A}) = 1] - \Pr[\text{Hyb}_1(\mathcal{A}) = 1]| \leq \Pr[\text{E}].$$

We now leverage Theorem 5.9 to argue that $\Pr[\text{E}] = \text{negl}(\lambda)$. To do so, we start by defining a suitable tree-based additive invariant. Similar to the proof of Theorem 5.22, we first associate a “validity interval” with each index $j \in [2n-1]$. For an index $i^* \in [n]$, we define the interval $I_j^{(i^*)}$ with each node j as follows:

- For $j \in [n]$, if $j = i^*$, let $I_j^{(i^*)} = [0, 0]$. Otherwise, let $I_j^{(i^*)} = [0, 1]$.
- For an index $j \in [n+1, 2n-1]$, let j_L, j_R be the indices of the children of j according to Definition 5.1. If $I_{j_L}^{(i^*)} = [a_L, b_L]$ and $I_{j_R}^{(i^*)} = [a_R, b_R]$, then define $I_j^{(i^*)} = [a_L + a_R, b_L + b_R] = I_{j_L}^{(i^*)} + I_{j_R}^{(i^*)}$, where we define interval addition to be component-wise addition: $[a_L, b_L] + [a_R, b_R] = [a_L + a_R, b_L + b_R]$.

By the same argument as in the proof of Claim 5.23, for all $j \in [2n-1]$ and all $i^* \in [n]$, we have that $I_j^{(i^*)} \subseteq [0, 2^h]$, where h is the height of node j . Now, we define the validity predicate $P_{\text{Valid}}: \{0, 1\}^* \rightarrow \{0, 1\}$ as follows:

- On input $(i^*, \text{ct}^{(0)}, \text{ct}^{(1)}, \text{sk}_0, \text{sk}_1, j)$, compute $x_b \leftarrow \text{HE.Dec}(\text{sk}_b, \text{ct}^{(b)})$ for each $b \in \{0, 1\}$.
- Output 1 if $x_0 = x_1 \in I_j^{(i^*)}$ and 0 otherwise.

In other words, the tuple $(i^*, \text{ct}^{(0)}, \text{ct}^{(1)}, \text{sk}_0, \text{sk}_1, j)$ is valid if the ciphertexts encrypt the same value, and moreover, they are within the valid range.⁸ We now show that P_{Valid} is a tree-based additive invariant.

⁸The range check is needed to ensure that all of the ciphertexts decrypt to values within the (bounded) support of the homomorphic encryption scheme. This is necessary to invoke evaluation correctness of Π_{HE} (see the proof of Lemma 5.32).

Lemma 5.32. *If Π_{HE} satisfies evaluation correctness, then P_{Valid} is a tree-based additive invariant.*

Proof. Take any $i^* \in [n]$, any $(\text{sk}_0, \text{pk}_0), (\text{sk}_1, \text{pk}_1)$ in the support of $\text{HE.Gen}(1^\lambda, 1^n)$, any triple of indices $j, j_L, j_R \in [2n-1]$ where j_L, j_R are the indices of the children of j , and any set of ciphertexts $(\text{ct}_L^{(0)}, \text{ct}_L^{(1)}), (\text{ct}_R^{(0)}, \text{ct}_R^{(1)})$. Suppose j has height $h \leq \log n$. For $b \in \{0, 1\}$, let $\text{ct}_{\text{sum}}^{(b)} = \text{HE.Add}(\text{pk}_b, \text{ct}_L^{(b)}, \text{ct}_R^{(b)})$. Suppose

$$P_{\text{Valid}}(i^*, \text{ct}_L^{(0)}, \text{ct}_L^{(1)}, \text{sk}_0, \text{sk}_1, j_L) = P_{\text{Valid}}(i^*, \text{ct}_R^{(0)}, \text{ct}_R^{(1)}, \text{sk}_0, \text{sk}_1, j_R) = 1.$$

This means $\text{HE.Dec}(\text{sk}_0, \text{ct}_L^{(0)}) = \text{HE.Dec}(\text{sk}_1, \text{ct}_L^{(1)}) \in I_{j_L}^{(i^*)}$ and $\text{HE.Dec}(\text{sk}_0, \text{ct}_R^{(0)}) = \text{HE.Dec}(\text{sk}_1, \text{ct}_R^{(1)}) \in I_{j_R}^{(i^*)}$. As argued above, since j_L, j_R have height $h-1$, we conclude that for $b \in \{0, 1\}$,

$$\text{HE.Dec}(\text{sk}_b, \text{ct}_L^{(b)}), \text{HE.Dec}(\text{sk}_b, \text{ct}_R^{(b)}) \in [0, 2^{h-1}].$$

Since $2^{h-1} \leq n/2$, we can appeal to evaluation correctness of Π_{H} and conclude that

$$\begin{aligned} \text{HE.Dec}(\text{sk}_0, \text{ct}_{\text{sum}}^{(0)}) &= \text{HE.Dec}(\text{sk}_0, \text{ct}_L^{(0)}) + \text{HE.Dec}(\text{sk}_0, \text{ct}_R^{(0)}) \\ &= \text{HE.Dec}(\text{sk}_1, \text{ct}_L^{(1)}) + \text{HE.Dec}(\text{sk}_1, \text{ct}_R^{(1)}) = \text{HE.Dec}(\text{sk}_1, \text{ct}_{\text{sum}}^{(1)}) \in I_{j_L} + I_{j_R} = I_j, \end{aligned}$$

We conclude that $P_{\text{Valid}}(i^*, \text{ct}_{\text{sum}}^{(0)}, \text{ct}_{\text{sum}}^{(1)}, \text{sk}_0, \text{sk}_1, j) = 1$. □

To leverage the predicate-propagation theorem ([Theorem 5.9](#)) to prove [Claim 5.31](#), we now define a mapping DeriveChal as

$$\text{DeriveChal}(S, i) := (S, i) \mapsto (S \setminus \{i\}, S, (i, 0)).$$

Let $\text{Expt} := \text{Expt}[P_{\text{Valid}}, \text{DeriveChal}]$ be the predicate propagation experiment from [Definition 5.7](#). First, we argue that

$$\Pr[E] \leq \Pr[\text{Expt}(\mathcal{A}) = 1], \tag{5.15}$$

where E is the event from [Eq. \(5.14\)](#). By construction, the adversary's view in Hyb_1 and Expt is identical. Suppose E occurs in an execution of Hyb_1 . Then the following hold:

- First $\text{Verify}'(\text{vk}_{\text{BARG}}, C_{i^*, 0}, 2n-1, \pi) = 1$. By construction of DeriveChal , we have that $\text{idx} = (i^*, 0)$ in the execution of $\text{Expt}(\mathcal{A})$. Hence, this means that $\text{Verify}'(\text{vk}_{\text{BARG}}, C_{\text{idx}}, 2n-1, \pi) = 1$.
- Next, $\text{HE.Dec}(\text{sk}_0, \text{ct}_{\text{root}}^{(0)}) \neq \text{HE.Dec}(\text{sk}_1, \text{ct}_{\text{root}}^{(1)})$. This means $P_{\text{Valid}}(\text{ct}_{\text{root}}^{(0)}, \text{ct}_{\text{root}}^{(1)}, \text{sk}_0, \text{sk}_1, 2n-1) = 0$.

Correspondingly, the output in Expt is also 1 in this case. Hence, we conclude that $\Pr[\text{Expt}(\mathcal{A}) = 1] \geq \Pr[E]$. To complete the proof, we analyze the predicate propagation hybrid experiment $\text{Expt}_j := \text{Expt}_j[P_{\text{Valid}}, \text{DeriveChal}]$.

Lemma 5.33. *If Π_{HE} is perfectly correct and Π_{Com} satisfies computational binding, then there exists a negligible function $\text{negl}(\cdot)$ such that for all $j \in [n]$, it holds that $\Pr[\text{Expt}_j(\mathcal{A}) = 1] = \text{negl}(\lambda)$.*

Proof. Suppose there exists some $j \in [n]$ where $\Pr[\text{Expt}_j(\mathcal{A}) = 1] \geq \varepsilon(\lambda)$ for some non-negligible ε . We use \mathcal{A} to construct an adversary \mathcal{B} that breaks computational binding of Π_{Com} .

1. On input the security parameter 1^λ , algorithm \mathcal{B} runs algorithm \mathcal{A} to obtain the input length 1^n , a set $S \subseteq [n]$, and an index $i^* \in S$.
2. Algorithm \mathcal{B} outputs the block length $1^{\ell_{\text{ct}}(\lambda, n)}$ and the vector length $2n-1$ to the challenger. The challenger responds with crs_{Com} .
3. Algorithm \mathcal{B} computes $(S \setminus \{i^*\}, S, (i^*, 0)) \leftarrow \text{DeriveChal}(S, i^*)$. It then samples the following components:
 - Sample $(\text{sk}_0, \text{pk}_0) \leftarrow \text{HE.Gen}(1^\lambda, 1^n)$ and $(\text{sk}_1, \text{pk}_1) \leftarrow \text{HE.Gen}(1^\lambda, 1^n)$.
 - Sample $(\text{crs}_{\text{BARG}}, \text{vk}_{\text{BARG}}, \text{td}_{\text{BARG}}) \leftarrow \text{TrapGen}'(1^\lambda, 1^{2n-1}, 1^s, 1^3, \{j\})$.
 - For each $b \in \{0, 1\}$, sample $\text{ct}_{\text{zero}}^{(b)} \leftarrow \text{HE.Enc}(\text{pk}_b, 0)$.

- For each $i \in [n] \setminus \{i^*\}$ and $b \in \{0, 1\}$, if $i \in S$, sample $ct_i^{(b)} \leftarrow \text{HE.Enc}(pk_b, 1)$. If $i \notin S$, sample $ct_i^{(b)} \leftarrow \text{HE.Enc}(pk_b, 0)$.
- Sample $ct_{i^*}^{(0)} \leftarrow \text{HE.Enc}(pk_0, 0)$ and $ct_{i^*}^{(1)} \leftarrow \text{HE.Enc}(pk_1, 1)$.
- For each $b \in \{0, 1\}$, let $(\text{com}_{\text{hk}}^{(b)}, \sigma_{\text{hk},1}^{(b)}, \dots, \sigma_{\text{hk},n}^{(b)}) \leftarrow \text{Com.Commit}(\text{crs}_{\text{Com}}, (ct_1^{(b)}, \dots, ct_n^{(b)}))$.

4. Algorithm \mathcal{B} constructs hk and vk according to Eqs. (5.1) and (5.2):

$$\begin{aligned} \text{hk} &= \left(\text{crs}_{\text{Com}}, \text{crs}_{\text{BARG}}, \{pk_b, ct_{\text{zero}}^{(b)}, ct_1^{(b)}, \dots, ct_n^{(b)}, \sigma_{\text{hk},1}^{(b)}, \dots, \sigma_{\text{hk},n}^{(b)}\}_{b \in \{0,1\}} \right) \\ \text{vk} &= (\text{crs}_{\text{Com}}, \text{vk}_{\text{BARG}}, pk_0, pk_1, ct_{\text{zero}}^{(0)}, ct_{\text{zero}}^{(1)}, \text{com}_{\text{hk}}^{(0)}, \text{com}_{\text{hk}}^{(1)}) \end{aligned}$$

Algorithm \mathcal{B} gives (hk, vk) to \mathcal{A} .

5. Algorithm \mathcal{A} outputs a digest $\text{dig} = (ct_{\text{root}}^{(0)}, ct_{\text{root}}^{(1)}, \text{com}_0, \text{com}_1, \pi_{\text{dig}})$ and a proof $\pi = \pi_{\text{dig}}$.
6. Algorithm \mathcal{B} extracts $(\hat{\vartheta}_j^{(0)}, \hat{\vartheta}_j^{(1)}, \sigma_j^{(0)}, \sigma_j^{(1)}, \tilde{w}_j) \leftarrow \text{Extract}'(\text{td}_{\text{BARG}}, \pi, j)$ and parses the auxiliary witness $\tilde{w}_j = (\hat{ct}^{(0)}, \hat{ct}^{(1)}, \sigma_{\text{hk}}^{(0)}, \sigma_{\text{hk}}^{(1)})$.
7. Algorithm \mathcal{B} checks if there exists $b \in \{0, 1\}$ where $\text{Com.Verify}(\text{crs}_{\text{Com}}, \text{com}_{\text{hk}}^{(b)}, j, \hat{ct}^{(b)}, \sigma_{\text{hk}}^{(b)}) = 1$ and $\hat{ct}^{(b)} \neq ct_j^{(b)}$. If so, it outputs the commitment $\text{com}_{\text{hk}}^{(b)}$, the index j , and the value-opening pairs $(ct_j^{(b)}, \sigma_{\text{hk},j}^{(b)})$ and $(\hat{ct}^{(b)}, \sigma_{\text{hk}}^{(b)})$.

By construction, the challenger samples $\text{crs}_{\text{Com}} \leftarrow \text{Com.Setup}(1^\lambda, 1^{\ell_{\text{ct}}(\lambda, n)}, 2n - 1)$, which matches the specification in Expt_j . This, algorithm \mathcal{B} perfectly simulates an execution of Expt_j for \mathcal{A} . By assumption, with probability ϵ , algorithm \mathcal{A} outputs dig and π such that the experiment outputs 1. This means the following conditions hold:

$$C_{i^*,0}(j, (\hat{\vartheta}_j^{(0)}, \hat{\vartheta}_j^{(1)}, \sigma_j^{(0)}, \sigma_j^{(1)}, \tilde{w}_j)) = 1 \quad \text{and} \quad P_{\text{Valid}}(i^*, \hat{\vartheta}_j^{(0)}, \hat{\vartheta}_j^{(1)}, \text{sk}_0, \text{sk}_1, j) = 0.$$

We consider two possibilities:

- Suppose $j = i^*$. By construction of $C_{i^*,0}$ (see Fig. 2), this means $\hat{\vartheta}_j^{(b)} = ct_{\text{zero}}^{(b)}$ for $b \in \{0, 1\}$. By construction, $ct_{\text{zero}}^{(b)}$ is an encryption of 0 under pk_b . In this case, $P_{\text{Valid}}(i^*, \hat{\vartheta}_j^{(0)}, \hat{\vartheta}_j^{(1)}, \text{sk}_0, \text{sk}_1, j) = 1$, which contradicts the premise.
- Suppose $j \neq i^*$. By construction of $C_{i^*,0}$, there are now two more possibilities:
 - Suppose for $b \in \{0, 1\}$, $\hat{\vartheta}_j^{(b)} = ct_{\text{zero}}^{(b)}$. As in the first case, this means $\hat{\vartheta}_j^{(0)}$ and $\hat{\vartheta}_j^{(1)}$ decrypt to 0 under sk_0 and sk_1 , respectively. In this case $P_{\text{Valid}}(i^*, \hat{\vartheta}_j^{(0)}, \hat{\vartheta}_j^{(1)}, \text{sk}_0, \text{sk}_1, j) = 1$, which again contradicts the premise.
 - Suppose for $b \in \{0, 1\}$, $\hat{\vartheta}_j^{(b)} = \hat{ct}^{(b)}$. In this case, we also have

$$\begin{aligned} \text{Com.Verify}(\text{crs}_{\text{Com}}, \text{com}_{\text{hk}}^{(0)}, j, \hat{ct}^{(0)}, \sigma_{\text{hk}}^{(0)}) &= 1 \\ \text{Com.Verify}(\text{crs}_{\text{Com}}, \text{com}_{\text{hk}}^{(1)}, j, \hat{ct}^{(1)}, \sigma_{\text{hk}}^{(1)}) &= 1. \end{aligned}$$

Suppose $\hat{ct}^{(b)} = ct_j^{(b)}$ for all $b \in \{0, 1\}$. In this case, since $j \neq i^*$, the ciphertexts $ct_j^{(0)}, ct_j^{(1)}$ are either both encryptions of 0 (if $j \notin S$) or both encryptions of 1 (if $j \in S$). In this case,

$$P_{\text{Valid}}(i^*, \hat{\vartheta}_j^{(0)}, \hat{\vartheta}_j^{(1)}, \text{sk}_0, \text{sk}_1, j) = P_{\text{Valid}}(i^*, ct_j^{(0)}, ct_j^{(1)}, \text{sk}_0, \text{sk}_1, j) = 1,$$

which contradicts the premise. Thus, if P_{Valid} is not satisfied, we conclude that there exists some $b \in \{0, 1\}$ such that $\hat{ct}^{(b)} \neq ct_j^{(b)}$.

Thus, there exists some $b \in \{0, 1\}$ such that the following holds:

$$\hat{\text{ct}}^{(b)} \neq \text{ct}_j^{(b)} \quad \text{and} \quad \text{Com.Verify}(\text{crs}_{\text{Com}}, \text{com}_{\text{hk}}^{(b)}, j, \hat{\text{ct}}^{(b)}, \sigma_{\text{hk}}^{(b)}) = 1.$$

Moreover, by correctness of Π_{Com} , we have that

$$\text{Com.Verify}(\text{crs}_{\text{Com}}, \text{com}_{\text{hk}}^{(b)}, j, \text{ct}_j^{(b)}, \sigma_{\text{hk},j}^{(b)}) = 1.$$

In this case, algorithm \mathcal{B} successfully breaks the binding property of the commitment scheme. \square

Since for all $j \in [n]$, it holds that $\Pr[\text{Expt}_j(\mathcal{A}) = 1] = \text{negl}(\lambda)$, we can invoke [Theorem 5.9](#) to conclude that $\Pr[\text{Expt}(\mathcal{A}) = 1] = \text{negl}(\lambda)$. [Claim 5.31](#) now follows via [Eqs. \(5.14\)](#) and [\(5.15\)](#). \square

Claim 5.34. *If Π_{HE} is CPA-secure, then there exists a negligible function $\text{negl}(\cdot)$ such that*

$$|\Pr[\text{Hyb}_3(\mathcal{A}) = 1] - \Pr[\text{Hyb}_2(\mathcal{A}) = 1]| = \text{negl}(\lambda).$$

Proof. Follows by an analogous argument as the proof of [Claim 5.30](#). In particular, the reduction obtains pk_0 and $\text{ct}_{i^*}^{(0)}$ from the challenger. It samples $(\text{pk}_1, \text{sk}_1)$ itself which it can use to compute the output (according to the specification in Hyb_2 and Hyb_3). \square

Claim 5.35. *If Π_{HE} is perfectly correct and satisfies evaluation correctness, Π_{Com} is computationally binding, and Π_{BARG} satisfies set hiding with extraction, set hiding, and is somewhere extractable, then there exists a negligible function $\text{negl}(\cdot)$ such that $|\Pr[\text{Hyb}_4(\mathcal{A}) = 1] - \Pr[\text{Hyb}_3(\mathcal{A}) = 1]| = \text{negl}(\lambda)$.*

Proof. Follows by an analogous argument as the proof of [Claim 5.31](#). The only difference is that we take the mapping DeriveChal to be

$$\text{DeriveChal}(S, i) := (S, i) \mapsto (S, S, (i, 0)).$$

The rest of the analysis proceeds exactly as before. \square

[Theorem 5.29](#) now follows by combining [Claims 5.30](#), [5.31](#), [5.34](#) and [5.35](#). \square

6 Zero-Fixing Hash Function from Bilinear Maps

In this section, we give a direct construction of a zero-fixing hash function from composite-order pairing groups. This construction does *not* require making non-black-box use of cryptography (in contrast to the construction from [Section 5](#)) and highlights an algebraic approach for building zero-fixing hash functions. We begin by recalling the concept of composite-order pairing groups [[BGN05](#)] and the generalized family of subgroup decision assumptions [[BWY11](#)].

Definition 6.1 (Composite-Order Bilinear Group). Let $k \in \mathbb{N}$ be a constant. A symmetric k -prime composite-order bilinear group generator is an efficient algorithm CompGroupGen that takes as input the security parameter λ and outputs a description $(\mathbb{G}, \mathbb{G}_T, \{p_i\}_{i \in [k]}, g, e)$ of a bilinear group where each p_i is a distinct prime where $p_i = 2^{\Omega(\lambda)}$, \mathbb{G} and \mathbb{G}_T are cyclic groups of order $N = \prod_{i \in [k]} p_i$, g is a generator of \mathbb{G} , and $e: \mathbb{G} \times \mathbb{G} \rightarrow \mathbb{G}_T$ is a non-degenerate bilinear map (called the “pairing”). We require that the group operation in \mathbb{G} and \mathbb{G}_T as well as the pairing operation be efficiently computable.

Notation. Let \mathbb{G} be a cyclic group with order $N = \prod_{i \in [k]} p_i$ and generator g . We write \mathbb{Z}_N to denote the ring of integers modulo N . In the following, for $i \in [k]$, we write $\mathbb{G}_i = \langle g^{N/p_i} \rangle$ to denote the subgroup of \mathbb{G} of order p_i . Throughout this section, we will write g_i to denote a *random* generator of \mathbb{G}_i . For a set $S \subseteq [k]$, we write $\mathbb{G}(S)$ to denote subgroup of \mathbb{G} of order $\prod_{i \in S} p_i$. By the Chinese Remainder Theorem, we can write \mathbb{G} as a direct product $\mathbb{G} \cong \mathbb{G}_{p_1} \times \cdots \times \mathbb{G}_{p_k}$. For a group element $h \in \mathbb{G}$, we can write $h = \prod_{i \in [k]} h_i$ where each $h_i \in \mathbb{G}_i$; we refer to h_i as the component of h in the subgroup \mathbb{G}_i . If two elements $h_1, h_2 \in \mathbb{G}$ are equal (i.e., $h_1 = h_2$), then for all $i \in [k]$, the component of h_1 and h_2 in \mathbb{G}_i are also equal. We extend this terminology to \mathbb{G}_T .

General subgroup decision assumption. We now recall the general subgroup decision assumption formalized in [BWY11]. The general subgroup decision assumption essentially states that for sets $S_0, S_1 \subseteq [k]$, no efficient adversary can distinguish between a random element of $\mathbb{G}(S_0)$ from $\mathbb{G}(S_1)$ even given random elements from $\mathbb{G}(S)$ for any $S \subseteq [k]$ where $S \cap S_0$ and $S \cap S_1$ are both empty or both non-empty. We give the formal definition below:

Definition 6.2 (General Subgroup Decision [BWY11, adapted]). Let $k \in \mathbb{N}$ be a constant and let CompGroupGen be a symmetric k -prime composite-order bilinear group generator. For an adversary \mathcal{A} and a bit $b \in \{0, 1\}$, we define the general subgroup decision game $\text{ExptSubgroup}_{\mathcal{A}}(\lambda, b)$ for CompGroupGen as follows:

1. At the beginning of the game, algorithm \mathcal{A} outputs two non-empty sets $S_0, S_1 \subseteq [k]$ and any number of sets $T_1, \dots, T_n \subseteq [k]$. We require that for all $i \in [n]$ either $S_0 \cap T_i = \emptyset = S_1 \cap T_i$ or $S_0 \cap T_i \neq \emptyset \neq S_1 \cap T_i$.
2. The challenger samples $(\mathbb{G}, \mathbb{G}_T, \{p_i\}_{i \in [k]}, g, e) \leftarrow \text{CompGroupGen}(1^\lambda)$. It compute $N = \prod_{i \in [k]} p_i$ and sets $\mathcal{G} = (\mathbb{G}, \mathbb{G}_T, N, g, e)$. For each $i \in [n]$, the challenger samples $X_i \xleftarrow{\mathbb{R}} \mathbb{G}(T_i)$. It also samples $Z \leftarrow \mathbb{G}(S_b)$, and gives the challenge $(\mathcal{G}, X_1, \dots, X_n, Z)$ to the adversary.
3. The adversary outputs a bit $b' \in \{0, 1\}$, which is the output of the experiment.

We say that the general subgroup decision assumption holds with respect to CompGroupGen if for all efficient adversaries \mathcal{A} $|\Pr[\text{ExptSubgroup}_{\mathcal{A}}(\lambda, 0) = 1] - \Pr[\text{ExptSubgroup}_{\mathcal{A}}(\lambda, 1) = 1]| \leq \text{negl}(\lambda)$.

Constructing zero-fixing hash functions. We now describe our construction of a zero-fixing hash function from composite-order bilinear groups. To simplify the main construction, we will describe our construction with a *long* verification key. The verification algorithm of our construction only requires *local* access to the long verification key, so it is straightforward to compile our construction into one with a short verification key using a collision-resistant hash function (see Remark 6.5).

Construction 6.3 (Zero-Fixing Hash Function from Composite-Order Bilinear Maps). Let CompGroupGen be a 6-prime composite-order pairing group. We construct a zero-fixing hash function $\Pi_H = (\text{Setup}, \text{Hash}, \text{ProveOpen}, \text{VerOpen}, \text{Extract}, \text{ValidateDigest})$ as follows:

- **Setup**($1^\lambda, 1^n, S$): On input a security parameter λ , an input length n , and a set $S \subseteq [n]$, the setup algorithm samples $(\mathbb{G}, \mathbb{G}_T, \{p_i\}_{i \in [6]}, g, e) \leftarrow \text{CompGroupGen}(1^\lambda)$. Let $N = \prod_{i \in [6]} p_i$. For each $i \in [6]$, let $g_i \xleftarrow{\mathbb{R}} \mathbb{G}_i$ be a random generator of \mathbb{G}_i . Let $\mathcal{G} = (\mathbb{G}, \mathbb{G}_T, N, g, e)$ be the group description. The setup algorithm now constructs the hash key components as follows:
 - **Main components:** For each $i \in [n]$, sample $\alpha_i, \beta_i \xleftarrow{\mathbb{R}} \mathbb{Z}_N$. Set

$$A_i = \begin{cases} (g_1 g_4)^{\alpha_i} & i \notin S \\ (g_1 g_2 g_3 g_4)^{\alpha_i} & i \in S. \end{cases}$$

For each $i \in [n]$, let $B_i = (g_1 g_5)^{\beta_i}$.

- **Cross-terms:** For each $i, j \in [n]$ where $i \neq j$, sample $r_{i,j} \xleftarrow{\mathbb{R}} \mathbb{Z}_N$ and let $C_{i,j} \leftarrow g_1^{\alpha_i \beta_j} (g_2 g_3 g_4 g_5)^{r_{i,j}}$.
- **Digest validation components:** Sample $\beta^* \xleftarrow{\mathbb{R}} \mathbb{Z}_N$ and let $B^* = (g_1 g_5)^{\beta^*}$. For each $i \in [n]$, sample $r_i^* \xleftarrow{\mathbb{R}} \mathbb{Z}_N$ and let $D_i = g_1^{\alpha_i \beta^*} (g_2 g_3 g_4 g_5)^{r_i^*}$.

Output the hash key hk and verification key vk where

$$\text{hk} = \text{vk} = (\mathcal{G}, g_1, \{A_i, B_i, D_i\}_{i \in [n]}, \{C_{i,j}\}_{i \neq j}, B^*)$$

along with the extraction trapdoor $\text{td} = g_2 g_3$.

- **Hash**(hk, x): On input a hash key $\text{hk} = (\mathcal{G}, g_1, \{A_i, B_i, D_i\}_{i \in [n]}, \{C_{i,j}\}_{i \neq j}, B^*)$ and an input $x \in \{0, 1\}^n$, the hash algorithm computes $h = \prod_{i \in [n]} A_i^{x_i}$ and $u = \prod_{i \in [n]} D_i^{x_i}$. It outputs the digest $\text{dig} = (h, u)$.

- **ValidateDigest**(vk, dig): On input the verification key $\text{vk} = (\mathcal{G}, \{A_i, B_i, D_i\}_{i \in [n]}, \{C_{i,j}\}_{i \neq j}, B^*)$ and a digest $\text{dig} = (h, u)$, the digest-validation algorithm outputs 1 if $e(h, B^*) = e(g_1, u)$ and 0 otherwise.
- **ProveOpen**(hk, x, i): On input a hash key $\text{hk} = (\mathcal{G}, g_1, \{A_i, B_i, D_i\}_{i \in [n]}, \{C_{i,j}\}_{i \neq j}, B^*)$, a string $x \in \{0, 1\}^n$, and an index $i \in [n]$, the prove algorithm outputs $\sigma = \prod_{j \neq i} C_{j,i}^{x_j}$.
- **VerOpen**(vk, dig, i, b, σ): On input a hash key $\text{vk} = (\mathcal{G}, g_1, \{(A_i, B_i)\}_{i \in [n]}, B^*)$, a digest $\text{dig} = (h, \pi)$, an index $i \in [n]$, a bit $b \in \{0, 1\}$, and an opening σ , the verification algorithm outputs 1 if $e(h, B_i) = e(A_i, B_i)^b \cdot e(g_1, \sigma)$ and 0 otherwise.
- **Extract**(td, dig): On input a trapdoor td and a digest $\text{dig} = (h, u)$, the extraction algorithm outputs Matching if $e(h, \text{td}) = 1$ and NotMatching otherwise.

Theorem 6.4 (Correctness). *Construction 6.3 is correct.*

Proof. Take any $\lambda, n \in \mathbb{N}$ and $x \in \{0, 1\}^n$. Let $i \in [n]$ be an index. Suppose $(\text{hk}, \text{vk}, \text{td}) \leftarrow \text{Setup}(1^\lambda, 1^n, \emptyset)$, $\text{dig} \leftarrow \text{Hash}(\text{hk}, x)$ and $\sigma \leftarrow \text{ProveOpen}(\text{hk}, x, i)$. By construction,

$$\text{hk} = \text{vk} = (\mathcal{G}, g_1, \{A_i, B_i, D_i\}_{i \in [n]}, \{C_{i,j}\}_{i \neq j}, B^*).$$

Next, $\text{dig} = (h, u)$ where $h = \prod_{i \in [n]} A_i^{x_i}$ and $u = \prod_{i \in [n]} D_i^{x_i}$. We consider the two properties:

- **Opening correctness:** By construction, $\sigma = \prod_{j \neq i} C_{j,i}^{x_j}$. By orthogonality, we have $e(A_j, B_i) = e(g_1, C_{j,i})$. Then,

$$e(h, B_i) = \prod_{j \in [n]} e(A_j, B_i)^{x_j} = e(A_i, B_i)^{x_i} \prod_{j \neq i} e(A_j, B_i)^{x_j} = e(A_i, B_i)^{x_i} \prod_{j \neq i} e(g_1, C_{j,i})^{x_j} = e(A_i, B_i)^{x_i} e(g_1, \sigma),$$

so $\text{VerOpen}(\text{vk}, \text{dig}, i, x_i, \sigma) = 1$, as required.

- **Digest correctness:** Again by orthogonality, we have $e(A_i, B^*) = e(g_1, D_i)$, so

$$e(h, B^*) = \prod_{j \in [n]} e(A_j, B^*)^{x_j} = \prod_{j \in [n]} e(g_1, D_j)^{x_j} = e(g_1, u),$$

and $\text{ValidateDigest}(\text{vk}, \text{dig}) = 1$. □

Remark 6.5 (Supporting Fast Verification). As described, the size of the verification key in [Construction 6.3](#) scales linearly with the input length n . This is incompatible with the succinctness requirements needed by our monotone BARG construction ([Construction 4.4](#)). However, it is straightforward to compress the verification key using a collision-resistant hash function. Observe that the verification algorithm VerOpen in [Construction 6.3](#) only requires *local* access to the verification key (i.e., it only needs to read elements A_i and B_i). The approach then is to only include a succinct *commitment* com to $(A_1, B_1), \dots, (A_n, B_n)$ in the verification key; the associated openings are included as part of the (long) hash key. Then, the opening for an index i would additionally contain the elements A_i, B_i as well as their openings with respect to com . The verifier would check that the correct elements A_i and B_i were provided and that they satisfy the verification relation. Security still holds as long as the scheme is computationally binding (since a computationally-bounded adversary would not be able to open com at i to any value other than (A_i, B_i)).

Security properties. We now show that each of the security requirements from [Definition 3.1](#) holds under the (general) subgroup decision assumption.

Theorem 6.6 (Set Hiding). *If the general subgroup decision holds with respect to CompGroupGen , then [Construction 6.3](#) satisfies set hiding.*

Proof. Let \mathcal{A} be an efficient adversary for the set hiding game. We begin by defining a sequence of hybrid experiments:

- Hyb_0 : This is experiment $\text{ExptSH}_{\mathcal{A}}(\lambda, 0)$. At the beginning of the game, the adversary outputs an input length n and a set $S \subseteq [n]$. Then the challenger samples $(\mathbb{G}, \mathbb{G}_T, \{p_i\}_{i \in [6]}, g, e) \leftarrow \text{CompGroupGen}(1^\lambda)$. It samples generators $g_i \xleftarrow{\mathbb{R}} \mathbb{G}_i$ and sets $\mathcal{G} = (\mathbb{G}, \mathbb{G}_T, N, g, e)$ where $N = \prod_{i \in [6]} p_i$. It constructs the hash key components as follows:

- **Main components:** For each $i \in [n]$, sample $\alpha_i, \beta_i \xleftarrow{\mathbb{R}} \mathbb{Z}_N$. Set $A_i = (g_1 g_4)^{\alpha_i}$ and $B_i = (g_1 g_5)^{\beta_i}$.
- **Cross-terms:** For each $i, j \in [n]$ where $i \neq j$, sample $r_{i,j} \xleftarrow{\mathbb{R}} \mathbb{Z}_N$ and let $C_{i,j} = g_1^{\alpha_i \beta_j} (g_2 g_3 g_4 g_5)^{r_{i,j}}$.
- **Digest validation components:** Sample $\beta^* \xleftarrow{\mathbb{R}} \mathbb{Z}_N$ and let $B^* = (g_1 g_5)^{\beta^*}$. For each $i \in [n]$, sample $r_i^* \xleftarrow{\mathbb{R}} \mathbb{Z}_N$ and let $D_i = g_1^{\alpha_i \beta^*} (g_2 g_3 g_4 g_5)^{r_i^*}$.

The challenger gives the hash key hk and verification key vk to \mathcal{A} where

$$\text{hk} = \text{vk} = (\mathcal{G}, g_1, \{A_i, B_i, D_i\}_{i \in [n]}, \{C_{i,j}\}_{i \neq j}, B^*).$$

Algorithm \mathcal{A} outputs a bit $b' \in \{0, 1\}$, which is the output of the experiment.

- Hyb_1 : Same as Hyb_0 , except the challenger now sets $C_{i,j} = A_i^{\beta_j} (g_2 g_3 g_4 g_5)^{r_{i,j}}$ and $D_i = A_i^{\beta^*} (g_2 g_3 g_4 g_5)^{r_i^*}$. In particular, in this experiment, the exponents α_i only show up in the definition of A_i .
- Hyb_2 : Same as Hyb_1 , except for $i \in S$, the challenger now sets $A_i = (g_1 g_2 g_3 g_4)^{\alpha_i}$.
- Hyb_3 : Same as Hyb_2 , except the challenger now sets $C_{i,j} = g_1^{\alpha_i \beta_j} (g_2 g_3 g_4 g_5)^{r_{i,j}}$ and $D_i = g_1^{\alpha_i \beta^*} (g_2 g_3 g_4 g_5)^{r_i^*}$. This is experiment $\text{ExptSH}_{\mathcal{A}}(\lambda, 1)$.

We write $\text{Hyb}_i(\mathcal{A})$ to denote the output of an execution of Hyb_i with adversary \mathcal{A} . We now analyze each pair of adjacent hybrid experiments.

Lemma 6.7. $\Pr[\text{Hyb}_0(\mathcal{A}) = 1] = \Pr[\text{Hyb}_1(\mathcal{A}) = 1]$

Proof. The outputs of $\text{Hyb}_0(\mathcal{A})$ and $\text{Hyb}_1(\mathcal{A})$ are identically distributed. The only difference between these two distributions is the distribution of the cross-terms $C_{i,j}$ and D_i . According to the specification of Hyb_1 ,

$$C_{i,j} = A_i^{\beta_j} (g_2 g_3 g_4 g_5)^{r_{i,j}} = (g_1 g_4)^{\alpha_i \beta_j} (g_2 g_3 g_4 g_5)^{r_{i,j}} = g_1^{\alpha_i \beta_j} (g_2 g_3 g_5)^{r_{i,j}} g_4^{r_{i,j} + \alpha_i \beta_j}.$$

Since $r_{i,j} \xleftarrow{\mathbb{R}} \mathbb{Z}_N$ (and independent of all other quantities in hk, vk), the distribution of $r_{i,j} + \alpha_i \beta_j \pmod{p_4}$ is uniform over \mathbb{Z}_{p_4} . We conclude that the distribution of $C_{i,j}$ in Hyb_1 is distributed exactly as in Hyb_0 . A similar analysis applies to D_i , and the claim holds. \square

Lemma 6.8. *If the general subgroup decision assumption holds with respect to CompGroupGen , then it holds that $|\Pr[\text{Hyb}_1(\mathcal{A}) = 1] - \Pr[\text{Hyb}_2(\mathcal{A}) = 1]| = \text{negl}(\lambda)$.*

Proof. Suppose there exists an adversary \mathcal{A} where $|\Pr[\text{Hyb}_1(\mathcal{A}) = 1] - \Pr[\text{Hyb}_2(\mathcal{A}) = 1]| \geq \epsilon(\lambda)$ for some non-negligible ϵ . We use \mathcal{A} to construct an adversary \mathcal{B} that breaks the general subgroup decision assumption:

1. At the beginning of the game, algorithm \mathcal{B} submits $S_0 = \{1, 4\}$ and $S_1 = \{1, 2, 3, 4\}$ and the sets $\{1\}$, $\{4\}$, $\{2, 3, 4\}$, and $\{5\}$ to the subgroup decision challenger.
2. The challenger replies with the challenge $(\mathcal{G}, g_1, g_4, X_{234}, g_5, Z)$, where g_i is used to denote the random generator of \mathbb{G}_i and $X_{234} \xleftarrow{\mathbb{R}} \mathbb{G}(\{2, 3, 4\})$.
3. Algorithm \mathcal{B} starts running algorithm \mathcal{A} who outputs the input length 1^n and a set $S \subseteq [n]$. The challenger samples $\alpha_i \xleftarrow{\mathbb{R}} \mathbb{Z}_N, \beta_i \xleftarrow{\mathbb{R}} \mathbb{Z}_N, r_{i,j} \xleftarrow{\mathbb{R}} \mathbb{Z}_N$ for $i, j \in [n]$. It constructs the components of the hash key as follows:

$$A_i = \begin{cases} (g_1 g_4)^{\alpha_i} & i \notin S \\ Z^{\alpha_i} & i \in S \end{cases} \quad \text{and} \quad B_i = (g_1 g_5)^{\beta_i} \quad \text{and} \quad C_{i,j} = A_i^{\beta_j} (X_{234} g_5)^{r_{i,j}}.$$

Next, it samples $\beta^* \xleftarrow{R} \mathbb{Z}_N$ and $r_i^* \xleftarrow{R} \mathbb{Z}_N$ for each $i \in [n]$. It sets

$$B^* = (g_1 g_5)^{\beta^*} \quad \text{and} \quad D_i = A_i^{\beta^*} (X_{234} g_5)^{r_i^*}.$$

It sets $\text{hk} = \text{vk} = (\mathcal{G}, g_1, \{A_i, B_i, D_i\}_{i \in [n]}, \{C_{i,j}\}_{i \neq j}, B^*)$ and gives (hk, vk) to \mathcal{A} .

4. Algorithm \mathcal{A} outputs a bit $b' \in \{0, 1\}$, which algorithm \mathcal{B} also outputs.

The subgroup decision challenger samples the generators $g_i \xleftarrow{R} \mathbb{G}_i$ exactly as in Hyb_1 and Hyb_2 . Moreover, $X_{234} = (g_2 g_3 g_4)^\gamma$ for $\gamma \xleftarrow{R} \mathbb{Z}_N$. Since the only element that depends on $r_{i,j}$ is $C_{i,j}$ and $r_{i,j} \xleftarrow{R} \mathbb{Z}_N$, the distribution of $X_{234}^{r_{i,j}}$ is identical to the distribution of $(g_2 g_3 g_4)^{r_{i,j}}$. Similarly, the distribution of $X_{234}^{r_i^*}$ is identically distributed to $(g_2 g_3 g_4)^{r_i^*}$. We now consider the two possibilities:

- Suppose $Z = (g_1 g_4)^t$ where $t \xleftarrow{R} \mathbb{Z}_N$. This corresponds to an execution of Hyb_1 with α_i replaced by $\alpha_i t$ when $i \in S$. As long as t is non-zero modulo p_1 and p_4 (which happens with negligible probability), the distribution of $\alpha_i t$ is uniform over $\mathbb{Z}_{p_1 p_4}$. In this case, algorithm \mathcal{B} outputs 1 with probability $\Pr[\text{Hyb}_1(\mathcal{A}) = 1]$.
- Suppose $Z = (g_1 g_2 g_3 g_4)^t$ where $t \xleftarrow{R} \mathbb{Z}_N$. This corresponds to an execution of Hyb_2 with α_i replaced by $\alpha_i t$ whenever $i \in S$. As long as t is non-zero modulo $p_1, p_2, p_3,$ and p_4 (which happens with negligible probability), the distribution of $\alpha_i t$ is uniform over $\mathbb{Z}_{p_1 p_2 p_3 p_4}$. In this case, algorithm \mathcal{B} outputs 1 with probability $\Pr[\text{Hyb}_2(\mathcal{A}) = 1]$.

We conclude that algorithm \mathcal{B} succeeds with probability that is negligibly close to ε and the claim holds. \square

Lemma 6.9. $\Pr[\text{Hyb}_2(\mathcal{A}) = 1] = \Pr[\text{Hyb}_3(\mathcal{A}) = 1]$.

Proof. The outputs of $\text{Hyb}_2(\mathcal{A})$ and $\text{Hyb}_3(\mathcal{A})$ are identically distributed by an analogous argument as the proof of [Lemma 6.7](#). In particular, in Hyb_2 , for $i \in S$,

$$C_{i,j} = A_i^{\beta_j} (g_2 g_3 g_4 g_5)^{r_{i,j}} = (g_1 g_2 g_3 g_4)^{\alpha_i \beta_j} (g_2 g_3 g_4 g_5)^{r_{i,j}} = g_1^{\alpha_i \beta_j} (g_2 g_3 g_4)^{r_{i,j} + \alpha_i \beta_j} g_5^{r_{i,j}}.$$

Again since $r_{i,j} \xleftarrow{R} \mathbb{Z}_N$ and independent of all other quantities in hk, vk , the distribution of $r_{i,j} + \alpha_i \beta_j \pmod{p_2 p_3 p_4}$ is identical to the distribution of $r_{i,j} \xleftarrow{R} \mathbb{Z}_{p_2 p_3 p_4}$. A similar argument applies to $C_{i,j}$ for $i \notin S$ and the D_i terms. \square

Set hiding now follows by combining [Lemmas 6.7](#) to [6.9](#). \square

Theorem 6.10 (Index Hiding with Extracted Guess). *Assume the general subgroup decision holds with respect to CompGroupGen , then [Construction 6.3](#) satisfies one-sided index hiding with extracted guess ([Definition 3.3](#)).*

Proof. Let \mathcal{A} be an efficient adversary for the index hiding with extracted guess game. We define a sequence of hybrid experiments:

- Hyb_0 : This is the index hiding with extracted guess experiment $\text{ExptIHE}_{\mathcal{A}}(\lambda, 0)$. Namely, the adversary starts by outputting the input length 1^n , a set $S \subseteq [n]$ and an index $i^* \in S$. The challenger samples $(\mathbb{G}, \mathbb{G}_T, \{p_i\}_{i \in [6]}, g, e) \leftarrow \text{CompGroupGen}(1^\lambda)$. It samples generators $g_i \xleftarrow{R} \mathbb{G}_i$ and sets $\mathcal{G} = (\mathbb{G}, \mathbb{G}_T, N, g, e)$ where $N = \prod_{i \in [6]} p_i$. It constructs the hash key components as follows:
 - **Main components:** For each $i \in [n]$, sample $\alpha_i, \beta_i \xleftarrow{R} \mathbb{Z}_N$. Set $A_i = (g_1 g_2 g_3 g_4)^{\alpha_i}$ if $i \in S \setminus \{i^*\}$ and $A_i = (g_1 g_4)^{\alpha_i}$ if $i \notin S$. Set $A_{i^*} = (g_1 g_4)^{\alpha_{i^*}}$. Then, set $B_i = (g_1 g_5)^{\beta_i}$.
 - **Cross-terms:** For each $i, j \in [n]$ where $i \neq j$, sample $r_{i,j} \xleftarrow{R} \mathbb{Z}_N$ and let $C_{i,j} = g_1^{\alpha_i \beta_j} (g_2 g_3 g_4 g_5)^{r_{i,j}}$.
 - **Digest validation components:** Sample $\beta^* \xleftarrow{R} \mathbb{Z}_N$ and let $B^* = (g_1 g_5)^{\beta^*}$. For each $i \in [n]$, sample $r_i^* \xleftarrow{R} \mathbb{Z}_N$ and let $D_i = g_1^{\alpha_i \beta^*} (g_2 g_3 g_4 g_5)^{r_i^*}$.

The challenger gives the hash key hk and verification key vk to \mathcal{A} where

$$hk = vk = (\mathcal{G}, g_1, \{A_i, B_i, D_i\}_{i \in [n]}, \{C_{i,j}\}_{i \neq j}, B^*).$$

Algorithm \mathcal{A} then outputs a digest $dig = (h, u)$ and an opening σ . The output of the experiment is 1 if

$$e(h, B_{i^*}) = e(g_1, \sigma) \quad \text{and} \quad e(h, g_2 g_3) = 1.$$

- Hyb_1 : Same as Hyb_0 except the challenger sets $C_{i^*,j} = A_{i^*}^{\beta_j} (g_2 g_3 g_4 g_5)^{r_{i^*,j}}$ for all $j \neq i^*$ and $C_{i,i^*} = B_{i^*}^{\alpha_{i^*}} (g_2 g_3 g_4 g_5)^{r_{i,i^*}}$ for all $i \neq i^*$. Similarly, the challenger sets $D_{i^*} = A_{i^*}^{\beta_{i^*}} (g_2 g_3 g_4 g_5)^{r_{i^*,i^*}}$. In particular, the exponents α_{i^*} and β_{i^*} only shows up in A_{i^*} and B_{i^*} , respectively.
- Hyb_2 : Same as Hyb_1 except the challenger sets $B_{i^*} = (g_1 g_5 g_6)^{\beta_{i^*}}$.
- Hyb_3 : Same as Hyb_2 except the challenger sets $A_{i^*} = (g_1 g_4 g_6)^{\alpha_{i^*}}$.
- Hyb_4 : Same as Hyb_3 except the challenger sets $A_{i^*} = (g_1 g_2 g_4 g_6)^{\alpha_{i^*}}$.
- Hyb_5 : Same as Hyb_4 except the challenger sets $A_{i^*} = (g_1 g_2 g_3 g_4 g_6)^{\alpha_{i^*}}$.
- Hyb_6 : Same as Hyb_5 except the challenger sets $A_{i^*} = (g_1 g_2 g_3 g_4)^{\alpha_{i^*}}$. **Namely, there is no longer a \mathbb{G}_6 component in A_{i^*} .**
- Hyb_7 : Same as Hyb_6 except the challenger sets $B_{i^*} = (g_1 g_5)^{\beta_{i^*}}$. **Namely, there is no longer a \mathbb{G}_6 component in B_{i^*} .**
- Hyb_8 : Same as Hyb_7 except the challenger sets $C_{i^*,j} = g_1^{\alpha_{i^*} \beta_j} (g_2 g_3 g_4 g_5)^{r_{i^*,j}}$ and $C_{i,i^*} = g_1^{\alpha_i \beta_{i^*}} (g_2 g_3 g_4 g_5)^{r_{i,i^*}}$. Similarly, the challenger sets $D_{i^*} = g_1^{\alpha_{i^*} \beta_{i^*}} (g_2 g_3 g_4 g_5)^{r_{i^*,i^*}}$. This is experiment $\text{ExptIHE}_{\mathcal{A}}(\lambda, 1)$.

We write $\text{Hyb}_i(\mathcal{A})$ to denote the output of an execution of Hyb_i with adversary \mathcal{A} . We now analyze each pair of adjacent hybrid experiments. Our goal is to show that $\Pr[\text{Hyb}_8(\mathcal{A}) = 1] \geq \Pr[\text{Hyb}_0(\mathcal{A}) = 1] - \text{negl}(\lambda)$.

Lemma 6.11. $\Pr[\text{Hyb}_0(\mathcal{A}) = 1] = \Pr[\text{Hyb}_1(\mathcal{A}) = 1]$

Proof. The outputs of $\text{Hyb}_0(\mathcal{A})$ and $\text{Hyb}_1(\mathcal{A})$ are identically distributed. The only difference between these two distributions is the distribution of $C_{i^*,j}$, C_{i,i^*} , and D_{i^*} . According to the specification of Hyb_1 ,

$$\begin{aligned} C_{i^*,j} &= A_{i^*}^{\beta_j} (g_2 g_3 g_4 g_5)^{r_{i^*,j}} = (g_1 g_4)^{\alpha_{i^*} \beta_j} (g_2 g_3 g_4 g_5)^{r_{i^*,j}} = g_1^{\alpha_{i^*} \beta_j} g_4^{r_{i^*,j} + \alpha_{i^*} \beta_j} (g_2 g_3 g_5)^{r_{i^*,j}} \\ C_{i,i^*} &= B_{i^*}^{\alpha_{i^*}} (g_2 g_3 g_4 g_5)^{r_{i,i^*}} = (g_1 g_5)^{\alpha_{i^*} \beta_{i^*}} (g_2 g_3 g_4 g_5)^{r_{i,i^*}} = g_1^{\alpha_i \beta_{i^*}} (g_2 g_3 g_4)^{r_{i,i^*}} g_5^{r_{i,i^*} + \alpha_i \beta_{i^*}} \\ D_{i^*} &= A_{i^*}^{\beta_{i^*}} (g_2 g_3 g_4 g_5)^{r_{i^*,i^*}} = (g_1 g_4)^{\alpha_{i^*} \beta_{i^*}} (g_2 g_3 g_4 g_5)^{r_{i^*,i^*}} = g_1^{\alpha_{i^*} \beta_{i^*}} g_4^{r_{i^*,i^*} + \alpha_{i^*} \beta_{i^*}} (g_2 g_3 g_5)^{r_{i^*,i^*}} \end{aligned}$$

Since $r_{i^*,j}, r_{i,i^*}, r_{i^*,i^*} \xleftarrow{\mathbb{R}} \mathbb{Z}_N$ (and independent of all other quantities in hk, vk), the elements $C_{i^*,j}, C_{i,i^*}, D_{i^*}$ are distributed exactly as in Hyb_0 . \square

Lemma 6.12. *If the general subgroup decision assumption holds with respect to CompGroupGen , then it holds that $|\Pr[\text{Hyb}_1(\mathcal{A}) = 1] - \Pr[\text{Hyb}_2(\mathcal{A}) = 1]| = \text{negl}(\lambda)$.*

Proof. Suppose there exists an adversary \mathcal{A} where $|\Pr[\text{Hyb}_1(\mathcal{A}) = 1] - \Pr[\text{Hyb}_2(\mathcal{A}) = 1]| \geq \varepsilon(\lambda)$ for some non-negligible ε . We use \mathcal{A} to construct an adversary \mathcal{B} that breaks the general subgroup decision assumption:

1. At the beginning of the game, algorithm \mathcal{B} submits $S_0 = \{1, 5\}$ and $S_1 = \{1, 5, 6\}$ and the sets $\{1\}, \{2\}, \{3\}, \{4\}, \{5\}$ to the subgroup decision challenger.
2. The challenger replies with the challenge $(\mathcal{G}, g_1, g_2, g_3, g_4, g_5, Z)$, where g_i is used to denote the random generator of \mathbb{G}_i .

3. Algorithm \mathcal{B} starts running algorithm \mathcal{A} who outputs the input length 1^n , a set $S \subseteq [n]$ and an index $i^* \in S$. Algorithm \mathcal{B} samples $\alpha_i \xleftarrow{\mathbb{R}} \mathbb{Z}_N$, $\beta_i \xleftarrow{\mathbb{R}} \mathbb{Z}_N$, $r_{i,j} \xleftarrow{\mathbb{R}} \mathbb{Z}_N$ for $i, j \in [n]$. It sets

$$A_i = \begin{cases} (g_1 g_4)^{\alpha_i} & i \notin S \setminus \{i^*\} \\ (g_1 g_2 g_3 g_4)^{\alpha_i} & i \in S \setminus \{i^*\} \end{cases} \quad \text{and} \quad B_i = \begin{cases} Z & i = i^* \\ (g_1 g_5)^{\beta_i} & i \neq i^* \end{cases} \quad \text{and} \quad C_{i,j} = \begin{cases} g_1^{\alpha_i \beta_j} (g_2 g_3 g_4 g_5)^{r_{i,j}} & i, j \neq i^* \\ A_{i^*}^{\beta_j} (g_2 g_3 g_4 g_5)^{r_{i^*,j}} & i = i^* \\ B_{i^*}^{\alpha_i} (g_2 g_3 g_4 g_5)^{r_{i,i^*}} & j = i^*. \end{cases}$$

Next, it samples $\beta^* \xleftarrow{\mathbb{R}} \mathbb{Z}_N$ and $r_i^* \xleftarrow{\mathbb{R}} \mathbb{Z}_N$ for each $i \in [n]$. It sets

$$B^* = (g_1 g_5)^{\beta^*} \quad \text{and} \quad D_i = \begin{cases} g_1^{\alpha_i \beta^*} (g_2 g_3 g_4 g_5)^{r_i^*} & i \neq i^* \\ A_{i^*}^{\beta^*} (g_2 g_3 g_4 g_5)^{r_{i^*,i^*}} & i = i^*. \end{cases}$$

It sets $\text{hk} = \text{vk} = (\mathcal{G}, g_1, \{A_i, B_i, D_i\}_{i \in [n]}, \{C_{i,j}\}_{i \neq j}, B^*)$ and gives (hk, vk) to \mathcal{A} .

4. Algorithm \mathcal{A} outputs a digest $\text{dig} = (h, u)$ and an opening σ . Algorithm \mathcal{B} outputs 1 if

$$e(h, B_{i^*}) = e(A_{i^*}, B_{i^*}) \cdot e(g_1, \sigma) \quad \text{and} \quad e(h, g_2 g_3) = 1.$$

The subgroup decision challenger samples the generators $g_i \xleftarrow{\mathbb{R}} \mathbb{G}_i$ exactly as in Hyb_1 and Hyb_2 . Moreover, all of the components other than B_{i^*} is constructed exactly as described in Hyb_1 and Hyb_2 . Thus, it suffices to consider the distribution of B^* . We consider the two possibilities:

- Suppose $Z = (g_1 g_5)^t$ where $t \xleftarrow{\mathbb{R}} \mathbb{Z}_N$. This corresponds to an execution of Hyb_1 with $\beta_{i^*} = t$. In this case, algorithm \mathcal{B} outputs 1 with probability $\Pr[\text{Hyb}_1(\mathcal{A}) = 1]$.
- Suppose $Z = (g_1 g_5 g_6)^t$ where $t \xleftarrow{\mathbb{R}} \mathbb{Z}_N$. This corresponds to an execution of Hyb_2 with $\beta_{i^*} = t$. In this case, algorithm \mathcal{B} outputs 1 with probability $\Pr[\text{Hyb}_2(\mathcal{A}) = 1]$.

We conclude that algorithm \mathcal{B} succeeds with probability ε and the claim holds. \square

Lemma 6.13. *If the general subgroup decision assumption holds with respect to CompGroupGen , then it holds that $|\Pr[\text{Hyb}_2(\mathcal{A}) = 1] - \Pr[\text{Hyb}_3(\mathcal{A}) = 1]| = \text{negl}(\lambda)$.*

Proof. Suppose there exists an adversary \mathcal{A} where $|\Pr[\text{Hyb}_2(\mathcal{A}) = 1] - \Pr[\text{Hyb}_3(\mathcal{A}) = 1]| \geq \varepsilon(\lambda)$ for some non-negligible ε . We use \mathcal{A} to construct an adversary \mathcal{B} that breaks the general subgroup decision assumption:

1. At the beginning of the game, algorithm \mathcal{B} submits $S_0 = \{1, 4\}$ and $S_1 = \{1, 4, 6\}$ and the sets $\{1\}, \{2\}, \{3\}, \{4\}, \{5\}, \{1, 6\}$ to the subgroup decision challenger.
2. The challenger replies with the challenge $(\mathcal{G}, g_1, g_2, g_3, g_4, g_5, X_{16}, Z)$, where g_i is used to denote the random generator of \mathbb{G}_i and $X_{16} \xleftarrow{\mathbb{R}} \mathbb{G}(\{1, 6\})$.
3. Algorithm \mathcal{B} starts running algorithm \mathcal{A} who outputs the input length 1^n . Algorithm \mathcal{B} samples $\alpha_i \xleftarrow{\mathbb{R}} \mathbb{Z}_N$, $\beta_i \xleftarrow{\mathbb{R}} \mathbb{Z}_N$, $r_{i,j} \xleftarrow{\mathbb{R}} \mathbb{Z}_N$ for $i, j \in [n]$. It sets

$$A_i = \begin{cases} Z & i = i^* \\ (g_1 g_4)^{\alpha_i} & i \notin S \\ (g_1 g_2 g_3 g_4)^{\alpha_i} & i \in S \setminus \{i^*\} \end{cases} \quad \text{and} \quad B_i = \begin{cases} X_{16} g_5^{\beta_{i^*}} & i = i^* \\ (g_1 g_5)^{\beta_i} & i \neq i^* \end{cases} \quad \text{and} \quad C_{i,j} = \begin{cases} g_1^{\alpha_i \beta_j} (g_2 g_3 g_4 g_5)^{r_{i,j}} & i, j \neq i^* \\ A_{i^*}^{\beta_j} (g_2 g_3 g_4 g_5)^{r_{i^*,j}} & i = i^* \\ B_{i^*}^{\alpha_i} (g_2 g_3 g_4 g_5)^{r_{i,i^*}} & j = i^*. \end{cases}$$

Next, it samples $\beta^* \xleftarrow{\mathbb{R}} \mathbb{Z}_N$ and $r_i^* \xleftarrow{\mathbb{R}} \mathbb{Z}_N$ for each $i \in [n]$. It sets

$$B^* = (g_1 g_5)^{\beta^*} \quad \text{and} \quad D_i = \begin{cases} g_1^{\alpha_i \beta^*} (g_2 g_3 g_4 g_5)^{r_i^*} & i \neq i^* \\ A_{i^*}^{\beta^*} (g_2 g_3 g_4 g_5)^{r_{i^*,i^*}} & i = i^*. \end{cases}$$

It sets $\text{hk} = \text{vk} = (\mathcal{G}, g_1, \{A_i, B_i, D_i\}_{i \in [n]}, \{C_{i,j}\}_{i \neq j}, B^*)$ and gives (hk, vk) to \mathcal{A} .

4. Algorithm \mathcal{A} outputs a digest $\text{dig} = (h, u)$ and an opening σ . Algorithm \mathcal{B} outputs 1 if

$$e(h, B_{i^*}) = e(A_{i^*}, B_{i^*}) \cdot e(g_1, \sigma) \quad \text{and} \quad e(h, g_2 g_3) = 1.$$

The subgroup decision challenger samples the generators $g_i \xleftarrow{\mathbb{R}} \mathbb{G}_i$ exactly as in Hyb₂ and Hyb₃. We can write $X_{16} = (g_1 g_6)^{\gamma_{16}}$ where $\gamma_{16} \xleftarrow{\mathbb{R}} \mathbb{Z}_N$. The value $\gamma_{16} \bmod p_1 p_6$ corresponds to the value of $\beta_{i^*} \bmod p_1 p_6$ in Hyb₂ and Hyb₃. The remaining components other than A_{i^*} are sampled exactly as required in Hyb₂ and Hyb₃, so it suffices to consider A_{i^*} . We consider the two possibilities:

- Suppose $Z = (g_1 g_4)^t$ where $t \xleftarrow{\mathbb{R}} \mathbb{Z}_N$. This corresponds to an execution of Hyb₁ with $\alpha_{i^*} = t$. In this case, algorithm \mathcal{B} outputs 1 with probability $\Pr[\text{Hyb}_2(\mathcal{A}) = 1]$.
- Suppose $Z = (g_1 g_4 g_6)^t$ where $t \xleftarrow{\mathbb{R}} \mathbb{Z}_N$. This corresponds to an execution of Hyb₃ with $\alpha_{i^*} = t$. In this case, algorithm \mathcal{B} outputs 1 with probability $\Pr[\text{Hyb}_3(\mathcal{A}) = 1]$.

We conclude that algorithm \mathcal{B} succeeds with probability ε and the claim holds. \square

Lemma 6.14. *If the general subgroup decision assumption holds with respect to CompGroupGen, then it holds that $\Pr[\text{Hyb}_4(\mathcal{A}) = 1] \geq \Pr[\text{Hyb}_3(\mathcal{A}) = 1] - \text{negl}(\lambda)$.*

Proof. Let $\text{dig} = (h, u)$ be the digest output by \mathcal{A} in an execution of Hyb₃ and Hyb₄. For an index $i \in \{3, 4\}$, we define events $E_{i,1}$ and $E_{i,2}$:

- $E_{i,1}$: This is the event $e(h, B_{i^*}) = e(g_1, \sigma)$ and $e(h, g_3) = 1$ occurring in Hyb_i.
- $E_{i,2}$: This is the event that $e(h, g_2) = 1$ occurring in Hyb_i.

By definition, the output in Hyb_i is 1 if and only if both events $E_{i,1}$ and $E_{i,2}$ occur. To complete the proof, we start by showing the following two properties: (1) $|\Pr[E_{3,1}] - \Pr[E_{4,1}]| = \text{negl}(\lambda)$; and (2) $\Pr[E_{4,1} \wedge \neg E_{4,2}] = \text{negl}(\lambda)$.

Claim 6.15. *If the subgroup decision assumption holds with respect to CompGroupGen, then $|\Pr[E_{3,1}] - \Pr[E_{4,1}]| = \text{negl}(\lambda)$.*

Proof. Suppose $|\Pr[E_{3,1}] - \Pr[E_{4,1}]| \geq \varepsilon(\lambda)$ for some non-negligible ε . We use \mathcal{A} to construct an adversary \mathcal{B} that breaks the general subgroup decision assumption:

1. At the beginning of the game, algorithm \mathcal{B} submits $S_0 = \{1, 4, 6\}$ and $S_1 = \{1, 2, 4, 6\}$ and the sets $\{1\}, \{3\}, \{4\}, \{5\}, \{6\}, \{2, 3, 4\}$.
2. The challenger replies with the challenge $(\mathcal{G}, g_1, g_3, g_4, g_5, g_6, X_{234}, Z)$ where g_i is a random generator of \mathbb{G}_i and $X_{234} \xleftarrow{\mathbb{R}} \mathbb{G}(\{2, 3, 4\})$.
3. Algorithm \mathcal{B} starts running algorithm \mathcal{A} who outputs the input length 1^n . Algorithm \mathcal{B} samples $\alpha_i \xleftarrow{\mathbb{R}} \mathbb{Z}_N$, $\beta_i \xleftarrow{\mathbb{R}} \mathbb{Z}_N$, $r_{i,j} \xleftarrow{\mathbb{R}} \mathbb{Z}_N$ for $i, j \in [n]$. It sets

$$A_i = \begin{cases} Z & i = i^* \\ (g_1 g_4)^{\alpha_i} & i \notin S \\ (g_1 X_{234})^{\alpha_i} & i \in S \setminus \{i^*\} \end{cases} \quad \text{and} \quad B_i = \begin{cases} (g_1 g_5 g_6)^{\beta_{i^*}} & i = i^* \\ (g_1 g_5)^{\beta_i} & i \neq i^* \end{cases} \quad \text{and} \quad C_{i,j} = \begin{cases} g_1^{\alpha_i \beta_j} (X_{234} g_5)^{r_{i,j}} & i, j \neq i^* \\ A_{i^*}^{\beta_j} (X_{234} g_5)^{r_{i^*,j}} & i = i^* \\ B_{i^*}^{\alpha_i} (X_{234} g_5)^{r_{i,i^*}} & j = i^*. \end{cases}$$

It samples $\beta^* \xleftarrow{\mathbb{R}} \mathbb{Z}_N$ and $r_i^* \xleftarrow{\mathbb{R}} \mathbb{Z}_N$ for each $i \in [n]$, and sets

$$B^* = (g_1 g_5)^{\beta^*} \quad \text{and} \quad D_i = \begin{cases} g_1^{\alpha_i \beta^*} (X_{234} g_5)^{r_i^*} & i \neq i^* \\ A_{i^*}^{\beta^*} (X_{234} g_5)^{r_{i^*,i^*}} & i = i^*. \end{cases}$$

It sets $\text{hk} = \text{vk} = (\mathcal{G}, g_1, \{A_i, B_i, D_i\}_{i \in [n]}, \{C_{i,j}\}_{i \neq j}, B^*)$ and gives (hk, vk) to \mathcal{A} .

4. After \mathcal{A} outputs the digest $\text{dig} = (h, u)$, algorithm \mathcal{B} outputs 1 if

$$e(h, B_{i^*}) = e(g_1, \sigma) \quad \text{and} \quad e(h, g_3) = 1.$$

We first argue that \mathcal{B} correctly simulates the hash key according to the specification of Hyb_3 and Hyb_4 . First, we can write $X_{234} = (g_2 g_3 g_4)^{\gamma_{234}}$ where $\gamma_{234} \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_N$. Since the challenger samples $r_{i,j}$ and $r_i^* \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_N$ and each of these values is used exactly once in the construction of hk , the distributions of $C_{i,j}$ and D_i are distributed exactly as they are in Hyb_3 and Hyb_4 unless γ_{234} is zero in the p_2, p_3 , or p_4 components. This happens with negligible probability over the choice of γ_{234} . Similarly, A_i for $i \in S \setminus \{i^*\}$ is distributed identically; the distributions of $\alpha_i \bmod p_2 p_3 p_4$ and that of $\alpha_i \gamma_{234} \bmod p_2 p_3 p_4$ when $\alpha_i \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_N$ are identical as long as γ_{234} is non-zero in the p_2, p_3 , and p_4 subgroups. It suffices to consider the distribution of A_{i^*} :

- Suppose $Z = (g_1 g_4 g_6)^t$ for $t \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_N$. This corresponds to an execution of Hyb_3 with $\alpha_{i^*} = t \bmod p_1 p_4 p_6$. Thus, algorithm \mathcal{B} outputs 1 with probability $\Pr[E_{3,1}]$.
- Suppose $Z = (g_1 g_2 g_4 g_6)^t$ for $t \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_N$. This corresponds to an execution of Hyb_2 with $\alpha_{i^*} = t \bmod p_1 p_4 p_6$. In this case, algorithm \mathcal{B} outputs 1 with probability $\Pr[E_{4,1}]$.

We conclude that \mathcal{B} succeeds with probability ε and the claim holds. \square

Claim 6.16. *If the subgroup decision assumption holds with respect to CompGroupGen , then $\Pr[E_{4,1} \wedge \neg E_{4,2}] = \text{negl}(\lambda)$.*

Proof. Suppose $\Pr[E_{4,1} \wedge \neg E_{4,2}] \geq \varepsilon(\lambda)$ for some non-negligible ε . We use \mathcal{A} to construct an adversary \mathcal{B} that breaks the general subgroup decision assumption:

1. At the beginning of the game, algorithm \mathcal{B} submits $S_0 = \{3, 6\}$ and $S_1 = \{2, 3, 6\}$ and the sets $\{1\}, \{3\}, \{4\}, \{5\}, \{6\}, \{2, 3\}, \{2, 6\}$.
2. The challenger replies with the challenge $(\mathcal{G}, g_1, g_3, g_4, g_5, g_6, X_{23}, X_{26}, Z)$ where g_i is a random generator of \mathbb{G}_i , $X_{23} \stackrel{\mathbb{R}}{\leftarrow} \mathbb{G}(\{2, 3\})$, and $X_{26} \stackrel{\mathbb{R}}{\leftarrow} \mathbb{G}(\{2, 6\})$.
3. Algorithm \mathcal{B} starts running algorithm \mathcal{A} who outputs the input length 1^n . Algorithm \mathcal{B} samples $\alpha_i \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_N$, $\beta_i \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_N$, $r_{i,j} \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_N$ for $i, j \in [n]$. It sets

$$A_i = \begin{cases} (g_1 g_4)^{\alpha_{i^*}} X_{26} & i = i^* \\ (g_1 g_4)^{\alpha_i} & i \notin S \\ (g_1 g_4 X_{23})^{\alpha_i} & i \in S \setminus \{i^*\} \end{cases} \quad \text{and} \quad B_i = \begin{cases} (g_1 g_5 g_6)^{\beta_{i^*}} & i = i^* \\ (g_1 g_5)^{\beta_i} & i \neq i^* \end{cases} \quad \text{and} \quad C_{i,j} = \begin{cases} g_1^{\alpha_i \beta_j} (X_{23} g_4 g_5)^{r_{i,j}} & i, j \neq i^* \\ A_{i^*}^{\beta_j} (X_{23} g_4 g_5)^{r_{i^*,j}} & i = i^* \\ B_{i^*}^{\alpha_i} (X_{23} g_4 g_5)^{r_{i,i^*}} & j = i^*. \end{cases}$$

It samples $\beta^* \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_N$ and $r_i^* \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_N$ for each $i \in [n]$, and sets

$$B^* = (g_1 g_5)^{\beta^*} \quad \text{and} \quad D_i = \begin{cases} g_1^{\alpha_i \beta^*} (X_{23} g_4 g_5)^{r_i^*} & i \neq i^* \\ A_{i^*}^{\beta^*} (X_{23} g_4 g_5)^{r_{i^*}^*} & i = i^*. \end{cases}$$

It sets $\text{hk} = \text{vk} = (\mathcal{G}, g_1, \{A_i, B_i, D_i\}_{i \in [n]}, \{C_{i,j}\}_{i \neq j}, B^*)$ and gives (hk, vk) to \mathcal{A} .

4. After \mathcal{A} outputs the digest $\text{dig} = (h, u)$, algorithm \mathcal{B} outputs 1 if

$$e(h, B_{i^*}) = e(g_1, \sigma) \quad \text{and} \quad e(h, g_3) = 1 \quad \text{and} \quad e(h, Z) = 1.$$

We first argue that algorithm \mathcal{B} correctly simulates an execution of Hyb_4 for \mathcal{A} . First, we can write $X_{23} = (g_2 g_3)^{\gamma_{23}}$ and $X_{26} = (g_2 g_6)^{\gamma_{26}}$, where $\gamma_{23}, \gamma_{26} \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_N$. With overwhelming probability, γ_{23} is non-zero in the p_2 and p_3 subgroups. In the following, we will assume this is the case. Since the challenger samples $r_{i,j}, r_i^* \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_N$ and each of these values is used exactly once in the construction of hk , the distributions of $C_{i,j}$ and D_i are distributed exactly as required in Hyb_4 . Similarly, the distribution of A_{i^*} coincides with setting $\alpha_{i^*} \bmod p_2 p_6$ as $\gamma_{26} \bmod p_2 p_6$ and $\alpha_i \bmod p_2 p_3$ as

$\alpha_i \gamma_{23} \bmod p_2 p_3$. Since each $\alpha_i \stackrel{R}{\leftarrow} \mathbb{Z}_N$ and $\gamma_{26} \stackrel{R}{\leftarrow} \mathbb{Z}_N$, this matches the distribution in Hyb_4 . Thus, with probability at least $\varepsilon - \text{negl}(\lambda)$, algorithm \mathcal{B} outputs $\text{dig} = (h, u)$ such that

$$e(h, B_{i^*}) = e(g_1, \sigma) \quad \text{and} \quad e(h, g_3) = 1 \quad \text{and} \quad e(h, g_2) \neq 1. \quad (6.1)$$

Suppose Eq. (6.1) holds. We first claim that with overwhelming probability, $e(h, g_6) = 1$. Suppose otherwise. If $\beta_{i^*} \neq 0 \bmod p_6$ (which happens with overwhelming probability), then B_{i^*} is non-zero in the order p_6 subgroup. If $e(h, g_6) \neq 1$, then $e(h, B_{i^*})$ is non-zero in the p_6 subgroup. By construction $e(g_1, \sigma)$ is necessarily 0 in the p_6 subgroup, so it can no longer be the case that $e(h, B_{i^*}) = e(g_1, \sigma)$. Now consider the probability that \mathcal{B} outputs 1:

- Suppose $Z = (g_3 g_6)^t$ for some $t \stackrel{R}{\leftarrow} \mathbb{Z}_N$. As argued previously, with overwhelming probability, if Eq. (6.1) holds, then $e(h, g_6) = 1$. Since $e(h, g_3) = 1$, this means $e(h, Z) = 1$ and algorithm \mathcal{B} outputs 1 with overwhelming probability.
- Suppose $Z = (g_2 g_3 g_6)^t$ for some $t \stackrel{R}{\leftarrow} \mathbb{Z}_N$. Since $e(h, g_2) \neq 1$, then $e(h, Z) \neq 1$ so long as $t \neq 0 \bmod p_2$, which holds with overwhelming probability. Thus, in this case, algorithm \mathcal{B} outputs 1 with negligible probability.

We now compute the advantage of \mathcal{B} . We consider three possibilities:

- Suppose \mathcal{A} outputs (h, u) such that $e(h, B_{i^*}) \neq e(g_1, \sigma)$ or $e(h, g_3) \neq 1$. Then, the output of \mathcal{B} is always 0.
- Suppose \mathcal{A} outputs (h, u) such that $e(h, B_{i^*}) = e(g_1, \sigma)$ and $e(h, g_3) = 1$ and $e(h, g_2) = 1$. By the previous analysis, if $e(h, B_{i^*}) = e(g_1, \sigma)$, then $e(h, g_6) = 1$ with overwhelming probability. Since $Z = (g_3 g_6)^t$ or $Z = (g_2 g_3 g_6)^t$, this means that \mathcal{B} outputs 1 with overwhelming probability regardless for both possible values of Z .
- Suppose \mathcal{A} outputs (h, u) such that $e(h, B_{i^*}) = e(g_1, \sigma)$ and $e(h, g_3) = 1$ and $e(h, g_2) \neq 1$. By the earlier analysis, this case occurs with probability at least $\varepsilon - \text{negl}(\lambda)$, and in this case, algorithm \mathcal{B} outputs 1 with probability $1 - \text{negl}(\lambda)$ if $Z = (g_3 g_6)^t$ and with probability $\text{negl}(\lambda)$ if $Z = (g_2 g_3 g_6)^t$.

Let ρ_1, ρ_2, ρ_3 be the probabilities of each of these cases. Then,

$$\begin{aligned} \Pr[\text{ExptSubgroup}_{\mathcal{B}}(\lambda, 0) = 1] &= \rho_2(1 - \text{negl}(\lambda)) + \rho_3 \cdot (1 - \text{negl}(\lambda)) \\ \Pr[\text{ExptSubgroup}_{\mathcal{B}}(\lambda, 1) = 1] &= \rho_2(1 - \text{negl}(\lambda)) + \rho_3 \cdot \text{negl}(\lambda) \end{aligned}$$

The advantage of \mathcal{B} is thus $\rho_3 - \text{negl}(\lambda) \geq \varepsilon - \text{negl}(\lambda)$, and the claim holds. \square

To complete the proof we have that

$$\begin{aligned} \Pr[\text{Hyb}_4(\mathcal{A}) = 1] &= \Pr[E_{4,1} \wedge E_{4,2}] \\ &= \Pr[E_{4,1}] - \Pr[E_{4,1} \wedge \neg E_{4,2}] \\ &\geq \Pr[E_{4,1}] - \text{negl}(\lambda) && \text{by Claim 6.16} \\ &\geq \Pr[E_{3,1}] - \text{negl}(\lambda) && \text{by Claim 6.15,} \end{aligned}$$

and Lemma 6.14 follows. \square

Lemma 6.17. *If the general subgroup decision assumption holds with respect to CompGroupGen , then it holds that $\Pr[\text{Hyb}_5(\mathcal{A}) = 1] \geq \Pr[\text{Hyb}_4(\mathcal{A}) = 1] - \text{negl}(\lambda)$.*

Proof. The proof follows by a similar argument as that of Lemma 6.14. Let $\text{dig} = (h, u)$ be the digest output by \mathcal{A} in an execution of Hyb_4 and Hyb_5 . For an index $i \in \{4, 5\}$, we define events $E_{i,1}$ and $E_{i,2}$ as in the proof of Lemma 6.14 (changes marked in green):

- $E_{i,1}$: This is the event $e(h, B_{i^*}) = e(g_1, \sigma)$ and $e(h, g_2) = 1$ occurring in Hyb_i .
- $E_{i,2}$: This is the event that $e(h, g_3) = 1$ occurring in Hyb_i .

By definition, the output in Hyb_i is 1 if and only if both events $E_{i,1}$ and $E_{i,2}$ occur. To complete the proof, we start by showing the following two properties: (1) $|\Pr[E_{4,1}] - \Pr[E_{5,1}]| = \text{negl}(\lambda)$; and (2) $\Pr[E_{5,1} \wedge \neg E_{5,2}] = \text{negl}(\lambda)$.

Claim 6.18. *If the subgroup decision assumption holds with respect to CompGroupGen , then $|\Pr[E_{4,1}] - \Pr[E_{5,1}]| = \text{negl}(\lambda)$.*

Proof. Suppose $|\Pr[E_{4,1}] - \Pr[E_{5,1}]| \geq \epsilon(\lambda)$ for some non-negligible ϵ . We use \mathcal{A} to construct an adversary \mathcal{B} that breaks the general subgroup decision assumption:

1. At the beginning of the game, algorithm \mathcal{B} submits $S_0 = \{1, 2, 4, 6\}$ and $S_1 = \{1, 2, 3, 4, 6\}$ and the sets $\{1\}, \{2\}, \{4\}, \{5\}, \{6\}, \{2, 3, 4\}$.
2. The challenger replies with the challenge $(\mathcal{G}, g_1, g_2, g_4, g_5, g_6, X_{234}, Z)$ where g_i is a random generator of \mathbb{G}_i and $X_{234} \xleftarrow{\mathbb{R}} \mathbb{G}(\{2, 3, 4\})$.
3. Algorithm \mathcal{B} starts running algorithm \mathcal{A} who outputs the input length 1^n . Algorithm \mathcal{B} samples $\alpha_i \xleftarrow{\mathbb{R}} \mathbb{Z}_N$, $\beta_i \xleftarrow{\mathbb{R}} \mathbb{Z}_N$, $r_{i,j} \xleftarrow{\mathbb{R}} \mathbb{Z}_N$ for $i, j \in [n]$. It sets

$$A_i = \begin{cases} Z & i = i^* \\ (g_1 g_4)^{\alpha_i} & i \notin S \\ (g_1 X_{234})^{\alpha_i} & i \in S \setminus \{i^*\} \end{cases} \quad \text{and} \quad B_i = \begin{cases} (g_1 g_5 g_6)^{\beta_{i^*}} & i = i^* \\ (g_1 g_5)^{\beta_i} & i \neq i^* \end{cases} \quad \text{and} \quad C_{i,j} = \begin{cases} g_1^{\alpha_i \beta_j} (X_{234} g_5)^{r_{i,j}} & i, j \neq i^* \\ A_{i^*}^{\beta_j} (X_{234} g_5)^{r_{i^*,j}} & i = i^* \\ B_{i^*}^{\alpha_i} (X_{234} g_5)^{r_{i,i^*}} & j = i^*. \end{cases}$$

It samples $\beta^* \xleftarrow{\mathbb{R}} \mathbb{Z}_N$ and $r_i^* \xleftarrow{\mathbb{R}} \mathbb{Z}_N$ for each $i \in [n]$, and sets

$$B^* = (g_1 g_5)^{\beta^*} \quad \text{and} \quad D_i = \begin{cases} g_1^{\alpha_i \beta^*} (X_{234} g_5)^{r_i^*} & i \neq i^* \\ A_{i^*}^{\beta^*} (X_{234} g_5)^{r_{i^*,i^*}} & i = i^*. \end{cases}$$

It sets $\text{hk} = \text{vk} = (\mathcal{G}, g_1, \{A_i, B_i, D_i\}_{i \in [n]}, \{C_{i,j}\}_{i \neq j}, B^*)$ and gives (hk, vk) to \mathcal{A} .

4. After \mathcal{A} outputs the digest $\text{dig} = (h, u)$, algorithm \mathcal{B} outputs 1 if

$$e(h, B_{i^*}) = e(g_1, \sigma) \quad \text{and} \quad e(h, g_2) = 1.$$

We first argue that \mathcal{B} correctly simulates the hash key according to the specification of Hyb_4 and Hyb_5 . First, we can write $X_{234} = (g_2 g_3 g_4)^{\gamma_{234}}$ where $\gamma_{234} \xleftarrow{\mathbb{R}} \mathbb{Z}_N$. Since the challenger samples $r_{i,j}$ and $r_i^* \xleftarrow{\mathbb{R}} \mathbb{Z}_N$ and each of these values is used exactly once in the construction of hk , the distributions of $C_{i,j}$ and D_i are distributed exactly as they are in Hyb_4 and Hyb_5 unless γ_{234} is zero in the p_2, p_3 , or p_4 components. This happens with negligible probability over the choice of γ_{234} . Similarly, A_i for $i \in S \setminus \{i^*\}$ is distributed identically; the reduction algorithm effectively samples $\alpha_i \bmod p_2 p_3 p_4$, $\alpha_i \gamma_{234} \bmod p_2 p_3 p_4$ which are identically distributed when $\alpha_i \xleftarrow{\mathbb{R}} \mathbb{Z}_N$ and γ_{234} is non-zero in the p_2, p_3 , and p_4 subgroups. It suffices to consider the distribution of A_{i^*} :

- Suppose $Z = (g_1 g_2 g_4 g_6)^t$ for $t \xleftarrow{\mathbb{R}} \mathbb{Z}_N$. This corresponds to an execution of Hyb_4 with $\alpha_{i^*} = t \bmod p_1 p_2 p_4 p_6$. Thus, algorithm \mathcal{B} outputs 1 with probability $\Pr[E_{4,1}]$.
- Suppose $Z = (g_1 g_2 g_3 g_4 g_6)^t$ for $t \xleftarrow{\mathbb{R}} \mathbb{Z}_N$. This corresponds to an execution of Hyb_5 with $\alpha_{i^*} = t \bmod p_1 p_2 p_3 p_4 p_6$. In this case, algorithm \mathcal{B} outputs 1 with probability $\Pr[E_{5,1}]$.

We conclude that \mathcal{B} succeeds with probability ϵ and the claim holds. \square

Claim 6.19. *If the subgroup decision assumption holds with respect to CompGroupGen , then $\Pr[E_{5,1} \wedge \neg E_{5,2}] = \text{negl}(\lambda)$.*

Proof. Suppose $\Pr[E_{5,1} \wedge \neg E_{5,2}] \geq \epsilon(\lambda)$ for some non-negligible ϵ . We use \mathcal{A} to construct an adversary \mathcal{B} that breaks the general subgroup decision assumption:

1. At the beginning of the game, algorithm \mathcal{B} submits $S_0 = \{2, 6\}$ and $S_1 = \{2, 3, 6\}$ and the sets $\{1\}, \{2\}, \{4\}, \{5\}, \{6\}, \{2, 3\}, \{3, 6\}$.

2. The challenger replies with the challenge $(\mathcal{G}, g_1, g_2, g_4, g_5, g_6, X_{23}, X_{36}, Z)$ where g_i is a random generator of \mathbb{G}_i , $X_{23} \xleftarrow{\mathbb{R}} \mathbb{G}(\{2, 3\})$, and $X_{36} \xleftarrow{\mathbb{R}} \mathbb{G}(\{3, 6\})$.
3. Algorithm \mathcal{B} starts running algorithm \mathcal{A} who outputs the input length 1^n . Algorithm \mathcal{B} samples $\alpha_i \xleftarrow{\mathbb{R}} \mathbb{Z}_N$, $\beta_i \xleftarrow{\mathbb{R}} \mathbb{Z}_N$, $r_{i,j} \xleftarrow{\mathbb{R}} \mathbb{Z}_N$ for $i, j \in [n]$. It sets

$$A_i = \begin{cases} (g_1 g_2 g_4)^{\alpha_{i^*}} X_{36} & i = i^* \\ (g_1 g_4)^{\alpha_i} & i \notin S \\ (g_1 g_4 X_{23})^{\alpha_i} & i \in S \setminus \{i^*\} \end{cases} \quad \text{and} \quad B_i = \begin{cases} (g_1 g_5 g_6)^{\beta_{i^*}} & i = i^* \\ (g_1 g_5)^{\beta_i} & i \neq i^* \end{cases} \quad \text{and} \quad C_{i,j} = \begin{cases} g_1^{\alpha_i \beta_j} (X_{23} g_4 g_5)^{r_{i,j}} & i, j \neq i^* \\ A_{i^*}^{\beta_j} (X_{23} g_4 g_5)^{r_{i^*,j}} & i = i^* \\ B_{i^*}^{\alpha_i} (X_{23} g_4 g_5)^{r_{i,i^*}} & j = i^*. \end{cases}$$

It samples $\beta^* \xleftarrow{\mathbb{R}} \mathbb{Z}_N$ and $r_i^* \xleftarrow{\mathbb{R}} \mathbb{Z}_N$ for each $i \in [n]$, and sets

$$B^* = (g_1 g_5)^{\beta^*} \quad \text{and} \quad D_i = \begin{cases} g_1^{\alpha_i \beta^*} (X_{23} g_4 g_5)^{r_i^*} & i \neq i^* \\ A_{i^*}^{\beta^*} (X_{23} g_4 g_5)^{r_{i^*,i^*}} & i = i^*. \end{cases}$$

It sets $\text{hk} = \text{vk} = (\mathcal{G}, g_1, \{A_i, B_i, D_i\}_{i \in [n]}, \{C_{i,j}\}_{i \neq j}, B^*)$ and gives (hk, vk) to \mathcal{A} .

4. After \mathcal{A} outputs the digest $\text{dig} = (h, u)$, algorithm \mathcal{B} outputs 1 if

$$e(h, B_{i^*}) = e(g_1, \sigma) \quad \text{and} \quad e(h, g_2) = 1 \quad \text{and} \quad e(h, Z) = 1.$$

We first argue that algorithm \mathcal{B} correctly simulates an execution of Hyb_4 for \mathcal{A} . First, we can write $X_{23} = (g_2 g_3)^{\gamma_{23}}$ and $X_{36} = (g_3 g_6)^{\gamma_{36}}$, where $\gamma_{23}, \gamma_{36} \xleftarrow{\mathbb{R}} \mathbb{Z}_N$. With overwhelming probability γ_{23} is non-zero in the p_2 and p_3 subgroups. In the following, we will assume this is the case. Since the challenger samples $r_{i,j}, r_i^* \xleftarrow{\mathbb{R}} \mathbb{Z}_N$ and each of these values is used exactly once in the construction of hk , the distributions of $C_{i,j}$ and D_i are distributed exactly as required in Hyb_5 . Similarly, the distribution of A_{i^*} coincides with setting $\alpha_{i^*} \bmod p_2 p_3 p_6$ as $\gamma_{236} \bmod p_2 p_3 p_6$ and $\alpha_i \bmod p_2 p_3$ as $\alpha_i \gamma_{23} \bmod p_2 p_3$. Since each $\alpha_i \xleftarrow{\mathbb{R}} \mathbb{Z}_N$ and $\gamma_{23}, \gamma_{36} \xleftarrow{\mathbb{R}} \mathbb{Z}_N$, this matches the distribution in Hyb_5 . Thus, with probability at least $\varepsilon - \text{negl}(\lambda)$, algorithm \mathcal{B} outputs $\text{dig} = (h, u)$ such that

$$e(h, B_{i^*}) = e(g_1, \sigma) \quad \text{and} \quad e(h, g_2) = 1 \quad \text{and} \quad e(h, g_3) \neq 1. \quad (6.2)$$

Suppose Eq. (6.2) holds. We first claim that with overwhelming probability, $e(h, g_6) = 1$. Suppose otherwise. If $\beta_{i^*} \neq 0 \bmod p_6$ (which happens with overwhelming probability), then B_{i^*} is non-zero in the order p_6 subgroup. If $e(h, g_6) \neq 1$, then $e(h, B_{i^*})$ is non-zero in the p_6 subgroup. By construction $e(g_1, \sigma)$ is necessarily 0 in the p_6 subgroup, so it can no longer be the case that $e(h, B_{i^*}) = e(g_1, \sigma)$. Now consider the probability that \mathcal{B} outputs 1:

- Suppose $Z = (g_2 g_6)^t$ for some $t \xleftarrow{\mathbb{R}} \mathbb{Z}_N$. As argued previously, with overwhelming probability, if Eq. (6.2) holds, then $e(h, g_6) = 1$. Since $e(h, g_2) = 1$, this means $e(h, Z) = 1$ and algorithm \mathcal{B} outputs 1 with overwhelming probability.
- Suppose $Z = (g_2 g_3 g_6)^t$ for some $t \xleftarrow{\mathbb{R}} \mathbb{Z}_N$. Since $e(h, g_3) \neq 1$, then $e(h, Z) \neq 1$ so long as $t \neq 0 \bmod p_3$, which holds with overwhelming probability. Thus, in this case, algorithm \mathcal{B} outputs 1 with negligible probability.

We now compute the advantage of \mathcal{B} . We consider three possibilities:

- Suppose \mathcal{A} outputs (h, u) such that $e(h, B_{i^*}) \neq e(g_1, \sigma)$ or $e(h, g_2) \neq 1$. Then, the output of \mathcal{B} is always 0.
- Suppose \mathcal{A} outputs (h, u) such that $e(h, B_{i^*}) = e(g_1, \sigma)$ and $e(h, g_2) = 1$ and $e(h, g_3) = 1$. By the previous analysis, if $e(h, B_{i^*}) = e(g_1, \sigma)$, then $e(h, g_6) = 1$ with overwhelming probability. Since $Z = (g_2 g_6)^t$ or $Z = (g_2 g_3 g_6)^t$, this means that \mathcal{B} outputs 1 with overwhelming probability regardless for both possible values of Z .
- Suppose \mathcal{A} outputs (h, u) such that $e(h, B_{i^*}) = e(g_1, \sigma)$ and $e(h, g_2) = 1$ and $e(h, g_3) \neq 1$. By the earlier analysis, this case occurs with probability at least $\varepsilon - \text{negl}(\lambda)$, and in this case, algorithm \mathcal{B} outputs 1 with probability $1 - \text{negl}(\lambda)$ if $Z = (g_2 g_6)^t$ and with probability $\text{negl}(\lambda)$ if $Z = (g_2 g_3 g_6)^t$.

Let ρ_1, ρ_2, ρ_3 be the probabilities of each of these cases. Then,

$$\begin{aligned}\Pr[\text{ExptSubgroup}_{\mathcal{B}}(\lambda, 0) = 1] &= \rho_2(1 - \text{negl}(\lambda)) + \rho_3 \cdot (1 - \text{negl}(\lambda)) \\ \Pr[\text{ExptSubgroup}_{\mathcal{B}}(\lambda, 1) = 1] &= \rho_2(1 - \text{negl}(\lambda)) + \rho_3 \cdot \text{negl}(\lambda)\end{aligned}$$

The advantage of \mathcal{B} is thus $\rho_3 - \text{negl}(\lambda) \geq \varepsilon - \text{negl}(\lambda)$, and the claim holds. \square

To complete the proof we have that

$$\begin{aligned}\Pr[\text{Hyb}_5(\mathcal{A}) = 1] &= \Pr[E_{5,1} \wedge E_{5,2}] \\ &= \Pr[E_{5,1}] - \Pr[E_{5,1} \wedge \neg E_{5,2}] \\ &\geq \Pr[E_{5,1}] - \text{negl}(\lambda) && \text{by Claim 6.19} \\ &\geq \Pr[E_{4,1}] - \text{negl}(\lambda) && \text{by Claim 6.18,}\end{aligned}$$

and Lemma 6.17 follows. \square

Lemma 6.20. *If the general subgroup decision assumption holds with respect to CompGroupGen, then it holds that $|\Pr[\text{Hyb}_5(\mathcal{A}) = 1] - \Pr[\text{Hyb}_6(\mathcal{A}) = 1]| = \text{negl}(\lambda)$.*

Proof. The proof is analogous to the proof of Lemma 6.13, except that the challenge subgroups are $S_0 = \{1, 2, 3, 4\}$ and $S_1 = \{1, 2, 3, 4, 6\}$. \square

Lemma 6.21. *If the general subgroup decision assumption holds with respect to CompGroupGen, then it holds that $|\Pr[\text{Hyb}_6(\mathcal{A}) = 1] - \Pr[\text{Hyb}_7(\mathcal{A}) = 1]| = \text{negl}(\lambda)$.*

Proof. The proof is analogous to the proof of Lemma 6.12, except that A_{i^*} is now in $\mathbb{G}(\{1, 2, 3, 4\})$. \square

Lemma 6.22. $\Pr[\text{Hyb}_7(\mathcal{A}) = 1] = \Pr[\text{Hyb}_8(\mathcal{A}) = 1]$.

Proof. This follows by an analogous argument as the proof of Lemma 6.11. \square

Combining Lemmas 6.11 to 6.14, 6.17 and 6.20 to 6.22, the index hiding with extracted guess property holds. \square

Theorem 6.23 (Zero Fixing). *If the general subgroup decision holds with respect to CompGroupGen, then Construction 6.3 satisfies selective zero-fixing.*

Proof. Let \mathcal{A} be an efficient adversary for the zero fixing game. We begin by defining a sequence of hybrid experiments:

- **Hyb₀:** This is the selective version of the experiment $\text{ExptZF}_{\mathcal{A}}(\lambda)$. Namely, the adversary starts by outputting an input length 1^n , a set $S \subseteq [n]$, and an index $i^* \in S$. The challenger samples $(\mathbb{G}, \mathbb{G}_T, \{p_i\}_{i \in [6]}, g, e) \leftarrow \text{CompGroupGen}(1^\lambda)$. It samples generators $g_i \xleftarrow{\mathbb{R}} \mathbb{G}_i$ and sets $\mathcal{G} = (\mathbb{G}, \mathbb{G}_T, N, g, e)$ where $N = \prod_{i \in [6]} p_i$. It constructs the hash key components as follows:
 - **Main components:** For each $i \in [n]$, sample $\alpha_i, \beta_i \xleftarrow{\mathbb{R}} \mathbb{Z}_N$ Set $A_i = (g_1 g_4)^{\alpha_i}$ if $i \notin S$ and $A_i = (g_1 g_2 g_3 g_4)^{\alpha_i}$ if $i \in S$. Then, set $B_i = (g_1 g_5)^{\beta_i}$.
 - **Cross-terms:** For each $i, j \in [n]$ where $i \neq j$, sample $r_{i,j} \xleftarrow{\mathbb{R}} \mathbb{Z}_N$ and let $C_{i,j} = g_1^{\alpha_i \beta_j} (g_2 g_3 g_4 g_5)^{r_{i,j}}$.
 - **Digest validation components:** Sample $\beta^* \xleftarrow{\mathbb{R}} \mathbb{Z}_N$ and let $B^* = (g_1 g_5)^{\beta^*}$. For each $i \in [n]$, sample $r_i^* \xleftarrow{\mathbb{R}} \mathbb{Z}_N$ and let $D_i = g_1^{\alpha_i \beta^*} (g_2 g_3 g_4 g_5)^{r_i^*}$.

The challenger gives the hash key hk and verification key vk to \mathcal{A} where

$$\text{hk} = \text{vk} = (\mathcal{G}, g_1, \{A_i, B_i, D_i\}_{i \in [n]}, \{C_{i,j}\}_{i \neq j}, B^*).$$

Algorithm \mathcal{A} then outputs a digest $\text{dig} = (h, u)$ and an opening σ . The output of the experiment is 1 if

$$e(h, B_{i^*}) = e(A_{i^*}, B_{i^*}) \cdot e(g_1, \sigma) \quad \text{and} \quad e(h, g_2 g_3) = 1.$$

- Hyb₁: Same as Hyb₀ except the challenger now sets $C_{i^*,j} = A_{i^*}^{\beta_j} (g_2 g_3 g_4 g_5)^{r_{i^*,j}}$ for all $j \neq i^*$ and $C_{i,i^*} = B_{i^*}^{\alpha_i} (g_2 g_3 g_4 g_5)^{r_{i,i^*}}$ for all $i \neq i^*$. Similarly, the challenger sets $D_{i^*} = A_{i^*}^{\beta_{i^*}} (g_2 g_3 g_4 g_5)^{r_{i^*,i^*}}$. In particular, the exponents α_{i^*} and β_{i^*} only shows up in A_{i^*} and B_{i^*} , respectively.
- Hyb₂: Same as Hyb₁ except the challenger now sets $B_{i^*} = (g_1 g_5 g_6)^{\beta_{i^*}}$.
- Hyb₃: Same as Hyb₂ except the challenger now sets $A_{i^*} = (g_1 g_2 g_3 g_4 g_6)^{\alpha_{i^*}}$.
- Hyb₄: Same as Hyb₃ except the experiment **outputs 0 if $e(h, g_6) \neq 1$** .

We write $\text{Hyb}_i(\mathcal{A})$ to denote the output of an execution of Hyb_i with adversary \mathcal{A} . We now analyze each pair of adjacent hybrid experiments.

Lemma 6.24. $\Pr[\text{Hyb}_0(\mathcal{A}) = 1] = \Pr[\text{Hyb}_1(\mathcal{A}) = 1]$

Proof. This follows by the same argument as in the proof of Lemma 6.11. \square

Lemma 6.25. *If the general subgroup decision assumption holds with respect to CompGroupGen , then it holds that $|\Pr[\text{Hyb}_1(\mathcal{A}) = 1] - \Pr[\text{Hyb}_2(\mathcal{A}) = 1]| = \text{negl}(\lambda)$.*

Proof. Suppose there exists an adversary \mathcal{A} where $|\Pr[\text{Hyb}_1(\mathcal{A}) = 1] - \Pr[\text{Hyb}_2(\mathcal{A}) = 1]| \geq \varepsilon(\lambda)$ for some non-negligible ε . We use \mathcal{A} to construct an adversary \mathcal{B} that breaks the general subgroup decision assumption:

1. At the beginning of the game, algorithm \mathcal{B} submits $S_0 = \{1, 5\}$ and $S_1 = \{1, 5, 6\}$ and the sets $\{1\}, \{2\}, \{3\}, \{4\}, \{5\}$ to the subgroup decision challenger.
2. The challenger replies with the challenge $(\mathcal{G}, g_1, g_2, g_3, g_4, g_5, Z)$, where g_i is used to denote the random generator of \mathbb{G}_i .
3. Algorithm \mathcal{B} starts running algorithm \mathcal{A} who outputs the input length 1^n , a set $S \subseteq [n]$ and an index $i^* \in S$. Algorithm \mathcal{B} samples $\alpha_i \xleftarrow{\mathbb{R}} \mathbb{Z}_N, \beta_i \xleftarrow{\mathbb{R}} \mathbb{Z}_N, r_{i,j} \xleftarrow{\mathbb{R}} \mathbb{Z}_N$ for $i, j \in [n]$. It sets

$$A_i = \begin{cases} (g_1 g_4)^{\alpha_i} & i \notin S \\ (g_1 g_2 g_3 g_4)^{\alpha_i} & i \in S \end{cases} \quad \text{and} \quad B_i = \begin{cases} Z & i = i^* \\ (g_1 g_5)^{\beta_i} & i \neq i^* \end{cases} \quad \text{and} \quad C_{i,j} = \begin{cases} g_1^{\alpha_i \beta_j} (g_2 g_3 g_4 g_5)^{r_{i,j}} & i, j \neq i^* \\ A_{i^*}^{\beta_j} (g_2 g_3 g_4 g_5)^{r_{i^*,j}} & i = i^* \\ B_{i^*}^{\alpha_i} (g_2 g_3 g_4 g_5)^{r_{i,i^*}} & j = i^*. \end{cases}$$

Next, it samples $\beta^* \xleftarrow{\mathbb{R}} \mathbb{Z}_N$ and $r_i^* \xleftarrow{\mathbb{R}} \mathbb{Z}_N$ for each $i \in [n]$. It sets

$$B^* = (g_1 g_5)^{\beta^*} \quad \text{and} \quad D_i = \begin{cases} g_1^{\alpha_i \beta^*} (g_2 g_3 g_4 g_5)^{r_i^*} & i \neq i^* \\ A_{i^*}^{\beta^*} (g_2 g_3 g_4 g_5)^{r_{i^*,i^*}} & i = i^*. \end{cases}$$

It sets $\text{hk} = \text{vk} = (\mathcal{G}, g_1, \{A_i, B_i, D_i\}_{i \in [n]}, \{C_{i,j}\}_{i \neq j}, B^*)$ and gives (hk, vk) to \mathcal{A} .

4. Algorithm \mathcal{A} outputs a digest $\text{dig} = (h, u)$ and an opening σ . Algorithm \mathcal{B} outputs 1 if

$$e(h, B_{i^*}) = e(A_{i^*}, B_{i^*}) \cdot e(g_1, \sigma) \quad \text{and} \quad e(h, g_2 g_3) = 1.$$

The subgroup decision challenger samples the generators $g_i \xleftarrow{\mathbb{R}} \mathbb{G}_i$ exactly as in Hyb₁ and Hyb₂. Moreover, all of the components other than B_{i^*} is constructed exactly as described in Hyb₁ and Hyb₂. Thus, it suffices to consider the distribution of B^* . We consider the two possibilities:

- Suppose $Z = (g_1 g_5)^t$ where $t \xleftarrow{\mathbb{R}} \mathbb{Z}_N$. This corresponds to an execution of Hyb₁ with $\beta_{i^*} = t$. In this case, algorithm \mathcal{B} outputs 1 with probability $\Pr[\text{Hyb}_1(\mathcal{A}) = 1]$.
- Suppose $Z = (g_1 g_5 g_6)^t$ where $t \xleftarrow{\mathbb{R}} \mathbb{Z}_N$. This corresponds to an execution of Hyb₂ with $\beta_{i^*} = t$. In this case, algorithm \mathcal{B} outputs 1 with probability $\Pr[\text{Hyb}_2(\mathcal{A}) = 1]$.

We conclude that algorithm \mathcal{B} succeeds with probability ε and the claim holds. \square

Lemma 6.26. *If the general subgroup decision assumption holds with respect to CompGroupGen, then it holds that $|\Pr[\text{Hyb}_2(\mathcal{A}) = 1] - \Pr[\text{Hyb}_3(\mathcal{A}) = 1]| = \text{negl}(\lambda)$.*

Proof. Suppose there exists an adversary \mathcal{A} where $|\Pr[\text{Hyb}_2(\mathcal{A}) = 1] - \Pr[\text{Hyb}_3(\mathcal{A}) = 1]| \geq \varepsilon(\lambda)$ for some non-negligible ε . We use \mathcal{A} to construct an adversary \mathcal{B} that breaks the general subgroup decision assumption:

1. At the beginning of the game, algorithm \mathcal{B} submits $S_0 = \{1, 2, 3, 4\}$ and $S_1 = \{1, 2, 3, 4, 6\}$ and the sets $\{1\}$, $\{2\}$, $\{3\}$, $\{4\}$, $\{5\}$, $\{1, 6\}$ to the subgroup decision challenger.
2. The challenger replies with the challenge $(\mathcal{G}, g_1, g_2, g_3, g_4, g_5, X_{16}, Z)$, where g_i is used to denote the random generator of \mathbb{G}_i and $X_{16} \stackrel{\mathbb{R}}{\leftarrow} \mathbb{G}(\{1, 6\})$.
3. Algorithm \mathcal{B} starts running algorithm \mathcal{A} who outputs the input length 1^n . Algorithm \mathcal{B} samples $\alpha_i \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_N$, $\beta_i \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_N$, $r_{i,j} \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_N$ for $i, j \in [n]$. It sets

$$A_i = \begin{cases} Z & i = i^* \\ (g_1 g_4)^{\alpha_i} & i \notin S \\ (g_1 g_2 g_3 g_4)^{\alpha_i} & i \in S \setminus \{i^*\} \end{cases} \quad \text{and} \quad B_i = \begin{cases} X_{16} g_5^{\beta_{i^*}} & i = i^* \\ (g_1 g_5)^{\beta_i} & i \neq i^* \end{cases} \quad \text{and} \quad C_{i,j} = \begin{cases} g_1^{\alpha_i \beta_j} (g_2 g_3 g_4 g_5)^{r_{i,j}} & i, j \neq i^* \\ A_{i^*}^{\beta_j} (g_2 g_3 g_4 g_5)^{r_{i^*,j}} & i = i^* \\ B_{i^*}^{\alpha_i} (g_2 g_3 g_4 g_5)^{r_{i,i^*}} & j = i^* \end{cases}$$

Next, it samples $\beta^* \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_N$ and $r_i^* \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_N$ for each $i \in [n]$. It sets

$$B^* = (g_1 g_5)^{\beta^*} \quad \text{and} \quad D_i = \begin{cases} g_1^{\alpha_i \beta^*} (g_2 g_3 g_4 g_5)^{r_i^*} & i \neq i^* \\ A_{i^*}^{\beta^*} (g_2 g_3 g_4 g_5)^{r_{i^*,i^*}} & i = i^* \end{cases}$$

It sets $\text{hk} = \text{vk} = (\mathcal{G}, g_1, \{A_i, B_i, D_i\}_{i \in [n]}, \{C_{i,j}\}_{i \neq j}, B^*)$ and gives (hk, vk) to \mathcal{A} .

4. Algorithm \mathcal{A} outputs a digest $\text{dig} = (h, u)$ and an opening σ . Algorithm \mathcal{B} outputs 1 if

$$e(h, B_{i^*}) = e(A_{i^*}, B_{i^*}) \cdot e(g_1, \sigma) \quad \text{and} \quad e(h, g_2 g_3) = 1.$$

The subgroup decision challenger samples the generators $g_i \stackrel{\mathbb{R}}{\leftarrow} \mathbb{G}_i$ exactly as in Hyb_2 and Hyb_3 . We can write $X_{16} = (g_1 g_6)^{\gamma_{16}}$ where $\gamma_{16} \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_N$. The value $\gamma_{16} \bmod p_1 p_6$ corresponds to the value of $\beta_{i^*} \bmod p_1 p_6$ in Hyb_2 and Hyb_3 . The remaining components other than A_{i^*} are sampled exactly as required in Hyb_2 and Hyb_3 , so it suffices to consider A_{i^*} . We consider the two possibilities:

- Suppose $Z = (g_1 g_2 g_3 g_4)^t$ where $t \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_N$. This corresponds to an execution of Hyb_1 with $\alpha_{i^*} = t$. In this case, algorithm \mathcal{B} outputs 1 with probability $\Pr[\text{Hyb}_2(\mathcal{A}) = 1]$.
- Suppose $Z = (g_1 g_2 g_3 g_4 g_6)^t$ where $t \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_N$. This corresponds to an execution of Hyb_3 with $\alpha_{i^*} = t$. In this case, algorithm \mathcal{B} outputs 1 with probability $\Pr[\text{Hyb}_3(\mathcal{A}) = 1]$.

We conclude that algorithm \mathcal{B} succeeds with probability ε and the claim holds. \square

Lemma 6.27. *If the general subgroup decision assumption holds with respect to CompGroupGen, then it holds that $|\Pr[\text{Hyb}_3(\mathcal{A}) = 1] - \Pr[\text{Hyb}_4(\mathcal{A}) = 1]| = \text{negl}(\lambda)$.*

Proof. Suppose there exists an adversary \mathcal{A} where $|\Pr[\text{Hyb}_3(\mathcal{A}) = 1] - \Pr[\text{Hyb}_4(\mathcal{A}) = 1]| \geq \varepsilon(\lambda)$ for some non-negligible ε . Since the only difference between Hyb_3 and Hyb_4 is the extra condition, it must be the case that with probability ε , algorithm \mathcal{A} outputs (h, u, σ) such that

$$e(h, B_{i^*}) = e(A_{i^*}, B_{i^*}) \cdot e(g_1, \sigma) \quad \text{and} \quad e(h, g_2 g_3) = 1 \quad \text{and} \quad e(h, g_6) \neq 1. \quad (6.3)$$

In all other cases, the output in Hyb_3 and Hyb_4 are identical. We use \mathcal{A} to construct an adversary \mathcal{B} for the general subgroup decision assumption:

1. At the beginning of the game, algorithm \mathcal{B} submits $S_0 = \{2, 3, 5\}$ and $S_1 = \{2, 3, 5, 6\}$ and the sets $\{1\}, \{2\}, \{3\}, \{4\}, \{5\}, \{2, 6\}, \{5, 6\}$ to the subgroup decision challenger.
2. The challenger replies with the challenge $(\mathcal{G}, g_1, g_2, g_3, g_4, g_5, X_{26}, X_{56}, Z)$, where g_i is used to denote the random generator of \mathbb{G}_i , $X_{26} \stackrel{\mathbb{R}}{\leftarrow} \mathbb{G}(\{2, 6\})$, and $X_{56} \stackrel{\mathbb{R}}{\leftarrow} \mathbb{G}(\{5, 6\})$.
3. Algorithm \mathcal{B} starts running algorithm \mathcal{A} who outputs the input length 1^n . Algorithm \mathcal{B} samples $\alpha_i \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_N$, $\beta_i \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_N$, $r_{i,j} \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_N$ for $i, j \in [n]$. It sets

$$A_i = \begin{cases} (g_1 g_3 g_4)^{\alpha_{i^*}} X_{26} & i = i^* \\ (g_1 g_4)^{\alpha_i} & i \notin S \\ (g_1 g_2 g_3 g_4)^{\alpha_i} & i \in S \setminus \{i^*\} \end{cases} \quad \text{and} \quad B_i = \begin{cases} g_1^{\beta_{i^*}} X_{56} & i = i^* \\ (g_1 g_5)^{\beta_i} & i \neq i^* \end{cases} \quad \text{and} \quad C_{i,j} = \begin{cases} g_1^{\alpha_i \beta_j} (g_2 g_3 g_4 g_5)^{r_{i,j}} & i, j \neq i^* \\ A_{i^*}^{\beta_j} (g_2 g_3 g_4 g_5)^{r_{i^*,j}} & i = i^* \\ B_{i^*}^{\alpha_i} (g_2 g_3 g_4 g_5)^{r_{i,i^*}} & j = i^*. \end{cases}$$

Next, it samples $\beta^* \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_N$ and $r_i^* \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_N$ for each $i \in [n]$. It sets

$$B^* = (g_1 g_5)^{\beta^*} \quad \text{and} \quad D_i = \begin{cases} g_1^{\alpha_i \beta^*} (g_2 g_3 g_4 g_5)^{r_i^*} & i \neq i^* \\ A_{i^*}^{\beta^*} (g_2 g_3 g_4 g_5)^{r_{i^*,i^*}} & i = i^*. \end{cases}$$

It sets $hk = vk = (\mathcal{G}, g_1, \{A_i, B_i, D_i\}_{i \in [n]}, \{C_{i,j}\}_{i \neq j}, B^*)$ and gives (hk, vk) to \mathcal{A} .

4. Algorithm \mathcal{A} outputs a digest $dig = (h, u)$ and an opening σ . Algorithm \mathcal{B} outputs 1 if

$$e(h, B_{i^*}) = e(A_{i^*}, B_{i^*}) \cdot e(g_1, \sigma) \quad \text{and} \quad e(h, g_2 g_3) = 1 \quad \text{and} \quad e(h, Z) = 1.$$

The subgroup decision challenger samples the generators $g_i \stackrel{\mathbb{R}}{\leftarrow} \mathbb{G}_i$ exactly as in Hyb_3 and Hyb_4 . We can write $X_{26} = (g_2 g_6)^{\gamma_{26}}$ and $X_{56} = (g_5 g_6)^{\gamma_{56}}$ where $\gamma_{26}, \gamma_{56} \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_N$. The value $\gamma_{26} \bmod p_1 p_6$ corresponds to the value of $\alpha_{i^*} \bmod p_2 p_6$ while the value γ_{56} corresponds to the value of $\beta_{i^*} \bmod p_5 p_6$. Thus, algorithm \mathcal{B} perfectly simulates the hash key for algorithm \mathcal{A} . Thus, with probability at least $\varepsilon - \text{negl}(\lambda)$, algorithm \mathcal{A} outputs (h, u, σ) that satisfies Eq. (6.3). Then, we have the following:

- It must be the case that $e(h, g_5) = 1$. Suppose otherwise. This means that h is non-zero in the G_5 subgroup. Consider the first verification condition $e(h, B_{i^*}) = e(A_{i^*}, B_{i^*}) \cdot e(g_1, \sigma)$. If h is non-zero in the \mathbb{G}_5 component, then the left-hand side $e(h, B_{i^*})$ is non-zero in the order- p_5 subgroup unless $\gamma_{56} = 0 \bmod p_5$, which happens with negligible probability. However, the right-hand side is guaranteed to be zero in the order p_5 subgroup (since neither A_{i^*} nor g_1 have non-zero components in \mathbb{G}_5).
- Suppose $Z = (g_2 g_3 g_5)^t$ for some $t \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_N$. Since $e(h, g_2 g_3) = 1$ and $e(h, g_5) = 1$, this means that $e(g, Z) = 1$ and algorithm \mathcal{B} always outputs 1.
- Suppose $Z = (g_2 g_3 g_5 g_6)^t$ for some $t \stackrel{\mathbb{R}}{\leftarrow} \mathbb{Z}_N$. From Eq. (6.3), we have that $e(h, g_6) \neq 1$, so h has a non-zero component in the G_6 subgroup. As long as $t \bmod p_6$ is non-zero (which happens with overwhelming probability), then $e(h, Z) \neq 1$. In this case, algorithm \mathcal{B} outputs 1 with negligible probability.

We have established that when Eq. (6.3) holds, algorithm \mathcal{B} is able to successfully distinguish the subgroup decision challenge. To complete the proof, we show that when Eq. (6.3) does not hold,⁹ then algorithm \mathcal{B} 's behavior is *independent* of the challenge Z .

1. Suppose \mathcal{A} outputs (h, u, σ) where either $e(h, B_{i^*}) \neq e(A_{i^*}, B_{i^*}) \cdot e(g_1, \sigma)$ or $e(h, g_2 g_3) \neq 1$. In this case, algorithm \mathcal{B} always outputs 0.

⁹Note that algorithm \mathcal{B} cannot check for itself whether Eq. (6.3) occurs or not since it does not know g_6 (and indeed, knowledge of g_6 would trivially break the subgroup decision assumption). Thus, our proof strategy is simply to argue that when Eq. (6.3) does not happen, then the behavior of algorithm \mathcal{B} is independent of the challenge Z .

2. Suppose \mathcal{A} outputs (h, u, σ) where $e(h, B_{i^*}) = e(A_{i^*}, B_{i^*}) \cdot e(g_1, \sigma)$, $e(h, g_2g_3) = 1$, and $e(h, g_6) = 1$. Since $e(h, B_{i^*}) = e(A_{i^*}, B_{i^*}) \cdot e(g_1, \sigma)$, our earlier analysis implies that with overwhelming probability over the choice of γ_{56} , $e(h, g_5) = 1$. Thus, in this case, $e(h, g_2g_3g_5g_6) = 1$, so $e(h, Z) = 1$ for both possible choice of Z . As such, algorithm \mathcal{B} always outputs 1 in this case.
3. Finally, suppose \mathcal{A} outputs (h, u, σ) that satisfies Eq. (6.3). By our analysis above, algorithm \mathcal{B} outputs 1 with probability 1 when $Z = (g_2g_3g_5)^t$ and probability $\text{negl}(\lambda)$ when $Z = (g_2g_3g_5g_6)^t$.

If we let ρ_1, ρ_2, ρ_3 be the probabilities of each of these possible cases, then we have

$$\begin{aligned} \Pr[\text{ExptSubgroup}_{\mathcal{B}}(\lambda, 0) = 1] &= \rho_2 + \rho_3 \\ \Pr[\text{ExptSubgroup}_{\mathcal{B}}(\lambda, 1) = 1] &= \rho_2 + \rho_3 \cdot \text{negl}(\lambda) \end{aligned}$$

The advantage of \mathcal{B} is thus $\rho_3(1 - \text{negl}(\lambda))$. By our above analysis, we have that $\rho_3 \geq \varepsilon - \text{negl}(\lambda)$ and so algorithm \mathcal{B} breaks the general subgroup decision assumption with advantage at least $\varepsilon - \text{negl}(\lambda)$. \square

Lemma 6.28. $\Pr[\text{Hyb}_4(\mathcal{A}) = 1] = \text{negl}(\lambda)$.

Proof. In order for the output of Hyb_4 to be 1, the adversary \mathcal{A} must output (h, u, σ) such that

$$e(h, B_{i^*}) = e(A_{i^*}, B_{i^*}) \cdot e(g_1, \sigma) \quad \text{and} \quad e(h, g_2g_3) = 1 \quad \text{and} \quad e(h, g_6) = 1.$$

We claim that this can only happen with negligible probability over the choice of α_{i^*} and β_{i^*} . By construction in Hyb_4 , as long as $\alpha_{i^*}, \beta_{i^*} \neq 0 \pmod{p_6}$ (which holds with overwhelming probability), $e(A_{i^*}, B_{i^*})$ will have a non-zero component in the order p_6 -subgroup. However, if $e(h, g_6) = 1$, then h is zero in the order p_6 -subgroup. Likewise, $e(g_1, \sigma)$ does not have an order p_6 subgroup. This means the left-hand side of the verification relation is zero in the order- p_6 subgroup while the right-hand side is non-zero. As such, the verification relation is unsatisfiable as long as $\alpha_{i^*}, \beta_{i^*} \neq 0 \pmod{p_6}$. \square

Combining Lemmas 6.24 to 6.28, we have that $\Pr[\text{Hyb}_0(\mathcal{A}) = 1] = \text{negl}(\lambda)$, and zero fixing security holds. \square

Theorem 6.29 (Extractor Validity). *If the general subgroup decision holds with respect to CompGroupGen , then Construction 6.3 satisfies extractor validity.*

Proof. Let \mathcal{A} be an efficient adversary for the extractor validity game. We begin by defining a sequence of hybrid experiments:

- Hyb_0 : This is experiment $\text{ExptEV}_{\mathcal{A}}(\lambda)$. Namely, the adversary starts by outputting an input length 1^n . The challenger samples $(\mathbb{G}, \mathbb{G}_T, \{p_i\}_{i \in [6]}, g, e) \leftarrow \text{CompGroupGen}(1^\lambda)$. It samples generators $g_i \xleftarrow{\mathbb{R}} \mathbb{G}_i$ and sets $\mathcal{G} = (\mathbb{G}, \mathbb{G}_T, N, g, e)$, where $N = \prod_{i \in [6]} p_i$. It constructs the hash key components as follows:
 - **Main components:** For each $i \in [n]$, sample $\alpha_i, \beta_i \xleftarrow{\mathbb{R}} \mathbb{Z}_N$. Set $A_i = (g_1g_4)^{\alpha_i}$ and $B_i = (g_1g_5)^{\beta_i}$.
 - **Cross-terms:** For each $i, j \in [n]$ where $i \neq j$, sample $r_{i,j} \xleftarrow{\mathbb{R}} \mathbb{Z}_N$ and let $C_{i,j} = g_1^{\alpha_i\beta_j} (g_2g_3g_4g_5)^{r_{i,j}}$.
 - **Digest validation components:** Sample $\beta^* \xleftarrow{\mathbb{R}} \mathbb{Z}_N$ and let $B^* = (g_1g_5)^{\beta^*}$. For each $i \in [n]$, sample $r_i^* \xleftarrow{\mathbb{R}} \mathbb{Z}_N$ and let $D_i = g_1^{\alpha_i\beta^*} (g_2g_3g_4g_5)^{r_i^*}$.

The challenger gives the hash key hk and verification key vk to \mathcal{A} where

$$\text{hk} = \text{vk} = (\mathcal{G}, g_1, \{A_i, B_i, D_i\}_{i \in [n]}, \{C_{i,j}\}_{i \neq j}, B^*).$$

Algorithm \mathcal{A} then outputs a digest $\text{dig} = (h, u)$ and the output of the experiment is 1 if

$$e(h, B^*) = e(g_1, u) \quad \text{and} \quad e(h, g_2g_3) \neq 1.$$

- Hyb_1 : Same as Hyb_0 , except the challenger now sets $D_i = (B^*)^{\alpha_i} (g_2g_3g_4g_5)^{r_i^*}$ for all $i \in [n]$. In particular, the exponent β^* only shows up in the definition of B^* .

- Hyb₂: Same as Hyb₁, except the challenger now sets $B^* = (g_1 g_2 g_5)^{\beta^*}$.
- Hyb₃: Same as Hyb₂, except the challenger now sets $B^* = (g_1 g_2 g_3 g_5)^{\beta^*}$.

We write $\text{Hyb}_i(\mathcal{A})$ to denote the output of an execution of Hyb_i with adversary \mathcal{A} . We now analyze each pair of adjacent hybrid experiments.

Lemma 6.30. $\Pr[\text{Hyb}_0(\mathcal{A}) = 1] = \Pr[\text{Hyb}_1(\mathcal{A}) = 1]$

Proof. The outputs of Hyb₀(\mathcal{A}) and Hyb₁(\mathcal{A}) are identically distributed. The only difference between these two distributions is the distribution of D_i . According to the specification of Hyb₁,

$$D_i = (B^*)^{\alpha_i} (g_2 g_3 g_4 g_5)^{r_i^*} = (g_1 g_5)^{\alpha_i \beta^*} (g_2 g_3 g_4 g_5)^{r_i^*} = g_1^{\alpha_i \beta^*} (g_2 g_3 g_4)^{r_i^*} g_5^{r_i^* + \alpha_i \beta^*}.$$

Since $r_i^* \xleftarrow{\mathbb{R}} \mathbb{Z}_N$ (and independent of all other quantities in hk, vk), the distribution of $r_i^* + \alpha_i \beta^* \pmod{p_5}$ is uniform over \mathbb{Z}_{p_5} . We conclude that the distribution of D_i in Hyb₁ is distributed exactly as in Hyb₀. \square

Lemma 6.31. *If the general subgroup decision assumption holds with respect to CompGroupGen, then it holds that $\Pr[\text{Hyb}_1(\mathcal{A}) = 1] \leq \Pr[\text{Hyb}_2(\mathcal{A}) = 1] + \text{negl}(\lambda)$.*

Proof. Let $\text{dig} = (h, u)$ be the digest output by \mathcal{A} in an execution of Hyb₁ and Hyb₂. For an index $i \in \{1, 2\}$, we define events $E_{i,1}$ and $E_{i,2}$:

- $E_{i,1}$: This is the event $e(h, B^*) = e(g_1, u)$ and $e(h, g_3) \neq 1$ occurring in Hyb_i.
- $E_{i,2}$: This is the event $e(h, B^*) = e(g_1, u)$ and $e(h, g_3) = 1$ and $e(h, g_2) \neq 1$ occurring in Hyb_i.

If the output of Hyb_i is 1, exactly one of $E_{i,1}$ or $E_{i,2}$ *must* happen (note that these events are mutually exclusive). Thus, for $i \in \{1, 2\}$, we can write

$$\Pr[\text{Hyb}_i(\mathcal{A}) = 1] = \Pr[E_{i,1}] + \Pr[E_{i,2}]. \quad (6.4)$$

We now analyze the probabilities of these events:

Claim 6.32. *If the general subgroup decision assumption holds with respect to CompGroupGen, then it holds that $|\Pr[E_{1,1}] - \Pr[E_{2,1}]| = \text{negl}(\lambda)$.*

Proof. Suppose $|\Pr[E_{1,1}] - \Pr[E_{2,1}]| \geq \varepsilon(\lambda)$ for some non-negligible ε . We use \mathcal{A} to construct an adversary \mathcal{B} that breaks the general subgroup decision assumption:

1. At the beginning of the game, algorithm \mathcal{B} submits $S_0 = \{5\}$ and $S_1 = \{2, 5\}$ and the sets $\{1\}, \{3\}, \{4\}, \{5\}, \{2, 5\}$.
2. The challenger replies with the challenge $(\mathcal{G}, g_1, g_3, g_4, g_5, X_{25}, Z)$ where g_i is a random generator of \mathbb{G}_i , $X_{25} \xleftarrow{\mathbb{R}} \mathbb{G}(\{2, 5\})$.
3. Algorithm \mathcal{B} starts running algorithm \mathcal{A} who outputs the input length 1^n . Algorithm \mathcal{B} samples $\alpha_i \xleftarrow{\mathbb{R}} \mathbb{Z}_N$, $\beta_i \xleftarrow{\mathbb{R}} \mathbb{Z}_N$, $r_{i,j} \xleftarrow{\mathbb{R}} \mathbb{Z}_N$ for $i, j \in [n]$. It sets

$$A_i = (g_1 g_4)^{\alpha_i} \quad \text{and} \quad B_i = (g_1 g_5)^{\beta_i} \quad \text{and} \quad C_{i,j} = g_1^{\alpha_i \beta_j} (X_{25} g_3 g_4)^{r_{i,j}}.$$

It samples $\beta^* \xleftarrow{\mathbb{R}} \mathbb{Z}_N$ and $r_i^* \xleftarrow{\mathbb{R}} \mathbb{Z}_N$ for each $i \in [n]$, and sets

$$B^* = g_1^{\beta^*} Z \quad \text{and} \quad D_i = (B^*)^{\alpha_i} (X_{25} g_3 g_4)^{r_i^*}.$$

It sets $\text{hk} = \text{vk} = (\mathcal{G}, g_1, \{A_i, B_i, D_i\}_{i \in [n]}, \{C_{i,j}\}_{i \neq j}, B^*)$ and gives (hk, vk) to \mathcal{A} .

4. After \mathcal{A} outputs the digest $\text{dig} = (h, u)$, algorithm \mathcal{B} outputs 1 if

$$e(h, B^*) = e(g_1, u) \quad \text{and} \quad e(h, g_3) \neq 1.$$

We first argue that \mathcal{B} correctly simulates the hash key according to the specification of Hyb_1 and Hyb_2 . First, we can write $X_{25} = (g_2 g_5)^{\gamma_{25}}$ where $\gamma_{25} \xleftarrow{\mathbb{R}} \mathbb{Z}_N$. Since the challenger samples $r_{i,j}$ and $r_i^* \xleftarrow{\mathbb{R}} \mathbb{Z}_N$ and each of these values is used exactly once in the construction of hk , the distributions of $C_{i,j}$ and D_i are distributed exactly as they are in Hyb_1 and Hyb_2 unless γ_{25} is zero in the p_2 or p_5 components. This happens with negligible probability over the choice of γ_{25} . It suffices to consider the distribution of B^* :

- Suppose $Z = g_5^t$ for $t \xleftarrow{\mathbb{R}} \mathbb{Z}_N$. This corresponds to an execution of Hyb_1 with $\beta^* = t \bmod p_5$. Thus, algorithm \mathcal{B} outputs 1 with probability $\Pr[E_{1,1}]$.
- Suppose $Z = (g_2 g_5)^t$ for $t \xleftarrow{\mathbb{R}} \mathbb{Z}_N$. This corresponds to an execution of Hyb_2 with $\beta^* = t \bmod p_2 p_5$. In this case, algorithm \mathcal{B} outputs 1 with probability $\Pr[E_{2,1}]$.

We conclude that algorithm \mathcal{B} succeeds with advantage $\varepsilon - \text{negl}(\lambda)$ and the claim follows. \square

Claim 6.33. *If the general subgroup decision assumption holds with respect to CompGroupGen , then $\Pr[E_{1,2}] = \text{negl}(\lambda)$.*

Proof. Suppose $\Pr[E_{1,2}] \geq \varepsilon(\lambda)$ for some non-negligible ε . We use \mathcal{A} to construct an adversary \mathcal{B} that breaks the general subgroup decision assumption:

1. At the beginning of the game, algorithm \mathcal{B} submits $S_0 = \{2, 3\}$ and $S_1 = \{3\}$ and the sets $\{1\}, \{3\}, \{4\}, \{5\}, \{2, 3\}$ to the subgroup decision challenger.
2. The challenger replies with the challenge $(\mathcal{G}, g_1, g_3, g_4, g_5, X_{23}, Z)$ where g_i is a random generator of \mathbb{G}_i and $X_{23} \xleftarrow{\mathbb{R}} \mathbb{G}(\{2, 3\})$.
3. Algorithm \mathcal{B} starts running algorithm \mathcal{A} who outputs the input length 1^n . Algorithm \mathcal{B} samples $\alpha_i \xleftarrow{\mathbb{R}} \mathbb{Z}_N$, $\beta_i \xleftarrow{\mathbb{R}} \mathbb{Z}_N$, $r_{i,j} \xleftarrow{\mathbb{R}} \mathbb{Z}_N$ for $i, j \in [n]$. It sets

$$A_i = (g_1 g_4)^{\alpha_i} \quad \text{and} \quad B_i = (g_1 g_5)^{\beta_i} \quad \text{and} \quad C_{i,j} = g_1^{\alpha_i \beta_j} (X_{23} g_4 g_5)^{r_{i,j}}.$$

It samples $\beta^* \xleftarrow{\mathbb{R}} \mathbb{Z}_N$ and $r_i^* \xleftarrow{\mathbb{R}} \mathbb{Z}_N$ for each $i \in [n]$, and sets

$$B^* = (g_1 g_5)^{\beta^*} \quad \text{and} \quad D_i = (B^*)^{\alpha_i} (X_{23} g_4 g_5)^{r_i^*}.$$

It sets $\text{hk} = \text{vk} = (\mathcal{G}, g_1, \{A_i, B_i, D_i\}_{i \in [n]}, \{C_{i,j}\}_{i \neq j}, B^*)$ and gives (hk, vk) to \mathcal{A} .

4. After \mathcal{A} outputs the digest $\text{dig} = (h, u)$, algorithm \mathcal{B} output 1 if the following hold:

$$e(h, B^*) = e(g_1, u) \quad \text{and} \quad e(h, g_3) = 1 \quad \text{and} \quad e(h, Z) = 1.$$

We first argue that algorithm \mathcal{B} correctly simulates an execution of Hyb_1 for \mathcal{A} . First, we can write $X_{23} = (g_2 g_3)^{\gamma_{23}}$ where $\gamma_{23} \xleftarrow{\mathbb{R}} \mathbb{Z}_N$. Since the challenger samples $r_{i,j}, r_i^* \xleftarrow{\mathbb{R}} \mathbb{Z}_N$ and each of these values is used exactly once in the construction of hk , the distributions of $C_{i,j}$ and D_i are statistically close to that in Hyb_1 unless γ_{23} is zero in the p_2 or p_3 components. This happens with negligible probability over the choice of γ_{23} . In the following analysis, we assume that γ_{23} is non-zero in both its p_2 and p_3 components. Thus, with probability at least $\varepsilon - \text{negl}(\lambda)$, algorithm \mathcal{B} outputs $\text{dig} = (h, u)$ such that event $E_{1,2}$ occurs. This means

$$e(h, B^*) = e(g_1, u) \quad \text{and} \quad e(h, g_3) = 1 \quad \text{and} \quad e(h, g_2) \neq 1. \tag{6.5}$$

Suppose Eq. (6.5) holds. We consider the probability that \mathcal{B} outputs 1:

- If $Z = (g_2 g_3)^t$ for some $t \xleftarrow{\mathbb{R}} \mathbb{Z}_N$ and $e(h, g_2) \neq 1$, then as long as $t \neq 0 \bmod p_2$, it will be the case that $e(h, Z) \neq 1$, so algorithm \mathcal{B} outputs 0.
- If $Z = g_3^t$ for some $t \xleftarrow{\mathbb{R}} \mathbb{Z}_N$ and $e(h, g_3) = 1$, then algorithm \mathcal{B} always outputs 1.

We now compute the advantage of \mathcal{B} . We consider three possibilities:

1. Suppose \mathcal{A} outputs (h, u) where either $e(h, B^*) \neq e(g_1, u)$ or $e(h, g_3) \neq 1$. Then, the output of \mathcal{B} is always 0.
2. Suppose \mathcal{A} outputs (h, u) where $e(h, B^*) = e(g_1, u)$, $e(h, g_3) = 1$, and $e(h, g_2) = 1$. Since $Z = (g_2 g_3)^t$ or $Z = g_3^t$, in both cases, $e(h, Z) = 1$ and algorithm \mathcal{B} always outputs 1.
3. Suppose \mathcal{A} outputs (h, u) such that Eq. (6.5) holds. By the above analysis, this case happens with probability at least $\varepsilon - \text{negl}(\lambda)$. Then algorithm \mathcal{B} outputs 1 with negligible probability if $Z = (g_2 g_3)^t$ and with probability 1 if $Z = g_3^t$.

Let ρ_1, ρ_2, ρ_3 be the probabilities of each of these cases. Then,

$$\begin{aligned}\Pr[\text{ExptSubgroup}_{\mathcal{B}}(\lambda, 0) = 1] &= \rho_2 + \rho_3 \cdot \text{negl}(\lambda) \\ \Pr[\text{ExptSubgroup}_{\mathcal{B}}(\lambda, 1) = 1] &= \rho_2 + \rho_3\end{aligned}$$

The advantage of \mathcal{B} is thus $\rho_3(1 - \text{negl}(\lambda)) \geq \varepsilon - \text{negl}(\lambda)$, and the claim holds. \square

Returning to the proof of Lemma 6.31, we appeal to Eq. (6.4) to write

$$\begin{aligned}\Pr[\text{Hyb}_1(\mathcal{A}) = 1] &= \Pr[E_{1,1}] + \Pr[E_{1,2}] && \text{by Eq. (6.4)} \\ &\leq \Pr[E_{1,1}] + \text{negl}(\lambda) && \text{by Claim 6.33} \\ &\leq \Pr[E_{2,1}] + \text{negl}(\lambda) && \text{by Claim 6.32} \\ &\leq \Pr[E_{2,1}] + \Pr[E_{2,2}] + \text{negl}(\lambda) \\ &= \Pr[\text{Hyb}_2(\mathcal{A}) = 1] + \text{negl}(\lambda) && \text{by Eq. (6.4),}\end{aligned}$$

which proves the lemma. \square

Lemma 6.34. *If the general subgroup decision assumption holds with respect to CompGroupGen, then it holds that $\Pr[\text{Hyb}_2(\mathcal{A}) = 1] \leq \Pr[\text{Hyb}_3(\mathcal{A}) = 1] + \text{negl}(\lambda)$.*

Proof. This follows by a similar argument as the proof of Lemma 6.31. Let $\text{dig} = (h, u)$ be the digest output by \mathcal{A} in an execution of Hyb_2 and Hyb_3 . For an index $i \in \{2, 3\}$, we define an analogous set of events $E_{i,1}$ and $E_{i,2}$ as in the proof of Lemma 6.31 (changes marked in green):

- $E_{i,1}$: This is the event $e(h, B^*) = e(g_1, u)$ and $e(h, g_2) \neq 1$ occurring in Hyb_i .
- $E_{i,2}$: This is the event $e(h, B^*) = e(g_1, u)$ and $e(h, g_2) = 1$ and $e(h, g_3) \neq 1$ occurring in Hyb_i .

Once again, we can write $\Pr[\text{Hyb}_i(\mathcal{A}) = 1] = \Pr[E_{i,1}] + \Pr[E_{i,2}]$. We analyze the probabilities of each of these events:

Claim 6.35. *If the general subgroup decision assumption holds with respect to CompGroupGen, then it holds that $|\Pr[E_{2,1}] - \Pr[E_{3,1}]| = \text{negl}(\lambda)$.*

Proof. Suppose $|\Pr[E_{2,1}] - \Pr[E_{3,1}]| \geq \varepsilon(\lambda)$ for some non-negligible ε . We use \mathcal{A} to construct an adversary \mathcal{B} that breaks the general subgroup decision assumption:

1. At the beginning of the game, algorithm \mathcal{B} submits $S_0 = \{5\}$ and $S_1 = \{3, 5\}$ and the sets $\{1\}, \{2\}, \{4\}, \{5\}, \{3, 5\}$.
2. The challenger replies with the challenge $(\mathcal{G}, g_1, g_2, g_4, g_5, X_{35}, Z)$ where g_i is a random generator of \mathbb{G}_i , $X_{35} \xleftarrow{\mathbb{R}} \mathbb{G}(\{3, 5\})$.
3. Algorithm \mathcal{B} starts running algorithm \mathcal{A} who outputs the input length 1^n . Algorithm \mathcal{B} samples $\alpha_i \xleftarrow{\mathbb{R}} \mathbb{Z}_N$, $\beta_i \xleftarrow{\mathbb{R}} \mathbb{Z}_N$, $r_{i,j} \xleftarrow{\mathbb{R}} \mathbb{Z}_N$ for $i, j \in [n]$. It sets

$$A_i = (g_1 g_4)^{\alpha_i} \quad \text{and} \quad B_i = (g_1 g_5)^{\beta_i} \quad \text{and} \quad C_{i,j} = g_1^{\alpha_i \beta_j} (X_{35} g_2 g_4)^{r_{i,j}}.$$

It samples $\beta^* \xleftarrow{\mathbb{R}} \mathbb{Z}_N$ and $r_i^* \xleftarrow{\mathbb{R}} \mathbb{Z}_N$ for each $i \in [n]$, and sets

$$B^* = (g_1 g_2)^{\beta^*} Z \quad \text{and} \quad D_i = (B^*)^{\alpha_i} (X_{35} g_2 g_4)^{r_i^*}.$$

It sets $\text{hk} = \text{vk} = (\mathcal{G}, g_1, \{A_i, B_i, D_i\}_{i \in [n]}, \{C_{i,j}\}_{i \neq j}, B^*)$ and gives (hk, vk) to \mathcal{A} .

4. After \mathcal{A} outputs the digest $\text{dig} = (h, u)$, algorithm \mathcal{B} outputs 1 if

$$e(h, B^*) = e(g_1, u) \quad \text{and} \quad e(h, g_2) \neq 1.$$

We first argue that \mathcal{B} correctly simulates the hash key according to the specification of Hyb_2 and Hyb_3 . First, we can write $X_{35} = (g_3 g_5)^{\gamma_{35}}$ where $\gamma_{35} \xleftarrow{\mathbb{R}} \mathbb{Z}_N$. Since the challenger samples $r_{i,j}$ and $r_i^* \xleftarrow{\mathbb{R}} \mathbb{Z}_N$ and each of these values is used exactly once in the construction of hk , the distributions of $C_{i,j}$ and D_i are distributed exactly as they are in Hyb_2 and Hyb_3 unless γ_{35} is zero in the p_3 or p_5 components. This happens with negligible probability over the choice of γ_{35} . It suffices to consider the distribution of B^* :

- Suppose $Z = g_5^t$ for $t \xleftarrow{\mathbb{R}} \mathbb{Z}_N$. This corresponds to an execution of Hyb_2 with $\beta^* = t \bmod p_5$. Thus, algorithm \mathcal{B} outputs 1 with probability $\Pr[E_{2,1}]$.
- Suppose $Z = (g_2 g_5)^t$ for $t \xleftarrow{\mathbb{R}} \mathbb{Z}_N$. This corresponds to an execution of Hyb_3 with $\beta^* = t \bmod p_2 p_5$. In this case, algorithm \mathcal{B} outputs 1 with probability $\Pr[E_{3,1}]$.

We conclude that algorithm \mathcal{B} succeeds with advantage $\varepsilon - \text{negl}(\lambda)$ and the claim follows. \square

Claim 6.36. *If the general subgroup decision assumption holds with respect to CompGroupGen , then $\Pr[E_{2,2}] = \text{negl}(\lambda)$.*

Proof. Suppose $\Pr[E_{2,2}] \geq \varepsilon(\lambda)$ for some non-negligible ε . We use \mathcal{A} to construct an adversary \mathcal{B} that breaks the general subgroup decision assumption:

1. At the beginning of the game, algorithm \mathcal{B} submits $S_0 = \{2, 3\}$ and $S_1 = \{2\}$ and the sets $\{1\}, \{2\}, \{4\}, \{5\}, \{2, 3\}$ to the subgroup decision challenger.
2. The challenger replies with the challenge $(\mathcal{G}, g_1, g_2, g_4, g_5, X_{23}, Z)$ where g_i is a random generator of \mathbb{G}_i and $X_{23} \xleftarrow{\mathbb{R}} \mathbb{G}(\{2, 3\})$.
3. Algorithm \mathcal{B} starts running algorithm \mathcal{A} who outputs the input length 1^n . Algorithm \mathcal{B} samples $\alpha_i \xleftarrow{\mathbb{R}} \mathbb{Z}_N$, $\beta_i \xleftarrow{\mathbb{R}} \mathbb{Z}_N$, $r_{i,j} \xleftarrow{\mathbb{R}} \mathbb{Z}_N$ for $i, j \in [n]$. It sets

$$A_i = (g_1 g_4)^{\alpha_i} \quad \text{and} \quad B_i = (g_1 g_5)^{\beta_i} \quad \text{and} \quad C_{i,j} = g_1^{\alpha_i \beta_j} (X_{23} g_4 g_5)^{r_{i,j}}.$$

It samples $\beta^* \xleftarrow{\mathbb{R}} \mathbb{Z}_N$ and $r_i^* \xleftarrow{\mathbb{R}} \mathbb{Z}_N$ for each $i \in [n]$, and sets

$$B^* = (g_1 g_2 g_5)^{\beta^*} \quad \text{and} \quad D_i = (B^*)^{\alpha_i} (X_{23} g_4 g_5)^{r_i^*}.$$

It sets $\text{hk} = \text{vk} = (\mathcal{G}, g_1, \{A_i, B_i, D_i\}_{i \in [n]}, \{C_{i,j}\}_{i \neq j}, B^*)$ and gives (hk, vk) to \mathcal{A} .

4. After \mathcal{A} outputs the digest $\text{dig} = (h, u)$, algorithm \mathcal{B} output 1 if the following hold:

$$e(h, B^*) = e(g_1, u) \quad \text{and} \quad e(h, g_2) = 1 \quad \text{and} \quad e(h, Z) = 1.$$

We first argue that algorithm \mathcal{B} correctly simulates an execution of Hyb_2 for \mathcal{A} . First, we can write $X_{23} = (g_2 g_3)^{\gamma_{23}}$ where $\gamma_{23} \xleftarrow{\mathbb{R}} \mathbb{Z}_N$. Since the challenger samples $r_{i,j}, r_i^* \xleftarrow{\mathbb{R}} \mathbb{Z}_N$ and each of these values is used exactly once in the construction of hk , the distributions of $C_{i,j}$ and D_i are statistically close to that in Hyb_2 unless γ_{23} is zero in the p_2 or p_3 component. This happens with negligible probability over the choice of γ_{23} , so in the following analysis, we assume that this is not the case. Thus, with probability at least $\varepsilon - \text{negl}(\lambda)$, algorithm \mathcal{B} outputs $\text{dig} = (h, u)$ such that

$$e(h, B^*) = e(g_1, u) \quad \text{and} \quad e(h, g_2) = 1 \quad \text{and} \quad e(h, g_3) \neq 1. \quad (6.6)$$

Suppose Eq. (6.6) holds. We consider the probability that \mathcal{B} outputs 1:

- If $Z = (g_2 g_3)^t$ for some $t \xleftarrow{\mathbb{R}} \mathbb{Z}_N$ and $e(h, g_3) \neq 1$, then with overwhelming probability over the choice of t , $e(h, Z) \neq 1$, and algorithm \mathcal{B} outputs 0.

- If $Z = g_2^t$ for some $t \stackrel{R}{\leftarrow} \mathbb{Z}_N$ and $e(h, g_2) = 1$, so algorithm \mathcal{B} outputs 1.

We now compute the advantage of \mathcal{B} . We consider three possibilities:

1. Suppose \mathcal{A} outputs (h, u) where either $e(h, B^*) \neq e(g_1, u)$ or $e(h, g_2) \neq 1$. Then, the output of \mathcal{B} is always 0.
2. Suppose \mathcal{A} outputs (h, u) where $e(h, B^*) = e(g_1, u)$, $e(h, g_2) = 1$, and $e(h, g_3) = 1$. Since either $Z = (g_2 g_3)^t$ or $Z = g_2^t$, in both cases, $e(h, Z) = 1$ and algorithm \mathcal{B} outputs 1.
3. Suppose \mathcal{A} outputs (h, u) such that Eq. (6.6) holds. By the above analysis, this case happens with probability at least $\varepsilon - \text{negl}(\lambda)$. Then algorithm \mathcal{B} outputs 1 with negligible probability if $Z = (g_2 g_3)^t$ and with probability 1 if $Z = g_2^t$.

Let ρ_1, ρ_2, ρ_3 be the probabilities of each of these cases. Then,

$$\begin{aligned} \Pr[\text{ExptSubgroup}_{\mathcal{B}}(\lambda, 0) = 1] &= \rho_2 + \rho_3 \cdot \text{negl}(\lambda) \\ \Pr[\text{ExptSubgroup}_{\mathcal{B}}(\lambda, 1) = 1] &= \rho_2 + \rho_3 \end{aligned}$$

The advantage of \mathcal{B} is thus $\rho_3(1 - \text{negl}(\lambda)) \geq \varepsilon - \text{negl}(\lambda)$, and the claim holds. \square

Returning to the proof of Lemma 6.34, we can now write

$$\begin{aligned} \Pr[\text{Hyb}_2(\mathcal{A}) = 1] &= \Pr[\text{E}_{2,1}] + \Pr[\text{E}_{2,2}] \\ &\leq \Pr[\text{E}_{2,1}] + \text{negl}(\lambda) && \text{by Claim 6.36} \\ &\leq \Pr[\text{E}_{3,1}] + \text{negl}(\lambda) && \text{by Claim 6.35} \\ &\leq \Pr[\text{E}_{3,1}] + \Pr[\text{E}_{3,2}] + \text{negl}(\lambda) \\ &= \Pr[\text{Hyb}_3(\mathcal{A}) = 1] + \text{negl}(\lambda), \end{aligned}$$

and the lemma holds. \square

Lemma 6.37. $\Pr[\text{Hyb}_3(\mathcal{A}) = 1] = \text{negl}(\lambda)$.

Proof. In order for the output in Hyb_3 to be 1, the adversary \mathcal{A} must output (h, u, σ) such that

$$e(h, B^*) = e(g_1, u) \quad \text{and} \quad e(h, g_2 g_3) \neq 1.$$

We claim that this can only happen with negligible probability over the choice of β^* . Specifically, if β^* is non-zero in the p_2 and p_3 subgroups, and $e(h, g_2 g_3) \neq 1$, then $e(h, B^*)$ is non-zero in the order $p_2 p_3$ subgroup. However $e(g_1, u)$ is always zero in the order $p_2 p_3$ subgroup, so the verification relation is unsatisfiable. \square

By Lemmas 6.30, 6.31 and 6.34, we have that $\Pr[\text{Hyb}_0(\mathcal{A}) = 1] \leq \Pr[\text{Hyb}_3(\mathcal{A}) = 1] + \text{negl}(\lambda)$. By Lemma 6.37, we have that $\Pr[\text{Hyb}_3(\mathcal{A}) = 1] \leq \text{negl}(\lambda)$. We conclude that $\Pr[\text{Hyb}_0(\mathcal{A}) = 1] \leq \text{negl}(\lambda)$ and extractor validity holds. \square

7 Monotone-Policy Aggregate Signatures

In this section, we formalize our construction of monotone policy aggregate signatures from a *non-adaptively* sound monotone policy BARG for NP together with a “puncturable” signature scheme (called an all-but-one signature scheme in [GVW19]).

Definition 7.1 (Puncturable Signature [GVW19, adapted]). An puncturable (or all-but-one) signature scheme with message space $\{0, 1\}^\lambda$ is a tuple of efficient algorithms $\Pi_{\text{PunctSig}} = (\text{Gen}, \text{GenPunc}, \text{Sign}, \text{Verify})$ with the following syntax:

- $\text{Gen}(1^\lambda) \rightarrow (\text{vk}, \text{sk})$: On input the security parameter λ , the key-generation algorithm outputs a key pair (vk, sk) .

- $\text{GenPunc}(1^\lambda, m^*) \rightarrow (\text{vk}, \text{sk})$: On input a security parameter λ and a message $m^* \in \{0, 1\}^\lambda$, the punctured key generation algorithm outputs a key pair (vk, sk) .
- $\text{Sign}(\text{sk}, m) \rightarrow \sigma$: On input a signing key sk and a message $m \in \{0, 1\}^\lambda$, the signing algorithm outputs a signature σ .
- $\text{Verify}(\text{vk}, m, \sigma) \rightarrow b$: On input a verification key vk , a message $m \in \{0, 1\}^\lambda$, and a signature σ , the verification algorithm outputs a bit $b \in \{0, 1\}$.

Moreover, the puncturable signature scheme should satisfy the following properties:

- **Correctness:** For all $\lambda \in \mathbb{N}$ and all $m \in \{0, 1\}^\lambda$, it holds that

$$\Pr \left[\text{Verify}(\text{vk}, m, \sigma) = 1 \quad : \quad \begin{array}{l} (\text{vk}, \text{sk}) \leftarrow \text{Gen}(1^\lambda) \\ \sigma \leftarrow \text{Sign}(\text{sk}, m) \end{array} \right] = 1.$$

- **Punctured correctness:** For all $\lambda \in \mathbb{N}$, all $m^* \in \{0, 1\}^\lambda$, and all $\sigma^* \in \{0, 1\}^*$, it holds that

$$\Pr \left[\text{Verify}(\text{vk}, m^*, \sigma^*) = 1 \quad : \quad (\text{vk}, \text{sk}) \leftarrow \text{GenPunc}(1^\lambda, m^*) \right] = 0.$$

- **Verification key indistinguishability:** For any adversary \mathcal{A} and any $b \in \{0, 1\}$, we define the verification key indistinguishability experiment $\text{ExptVKI}_{\mathcal{A}}(\lambda, b)$ as follows:

1. On input a security parameter λ , the adversary \mathcal{A} outputs a message $m^* \in \{0, 1\}^\lambda$ and sends it to the challenger.
2. The challenger samples $(\text{vk}_0, \text{sk}_0) \leftarrow \text{Gen}(1^\lambda)$ and $(\text{vk}_1, \text{sk}_1) \leftarrow \text{GenPunc}(1^\lambda, m^*)$ and gives vk_b to the adversary.
3. Next, the challenger can make signing queries on messages $m \in \{0, 1\}^\lambda \setminus \{m^*\}$. On each signing query, the challenger replies with $\sigma \leftarrow \text{Sign}(\text{sk}_b, m)$.
4. The adversary outputs a bit $b' \in \{0, 1\}$, which is the output of the experiment.

We say that Π_{PunctSig} satisfies verification key indistinguishability if for any efficient adversary \mathcal{A} there exists a negligible function $\text{negl}(\cdot)$ such that

$$\left| \Pr[\text{ExptVKI}_{\mathcal{A}}(\lambda, 0) = 1] - \Pr[\text{ExptVKI}_{\mathcal{A}}(\lambda, 1) = 1] \right| = \text{negl}(\lambda).$$

Remark 7.2 (Multiple Verification Keys). By a standard hybrid argument, we can show that any puncturable signature scheme that satisfies verification key indistinguishability also satisfies a stronger multi-key version of the definition where the adversary can ask for multiple verification keys (punctured at the same message m^*) and signatures on messages $m \neq m^*$ with respect to those keys. We define this formally below:

- **Multiple verification keys indistinguishability:** For any adversary \mathcal{A} and any $b \in \{0, 1\}$, we define the multiple verification key indistinguishability experiment $\text{ExptMVKI}_{\mathcal{A}}(\lambda, b)$ as follows:

1. On input a security parameter λ , the adversary \mathcal{A} outputs the number of challenge keys 1^n together with a message $m^* \in \{0, 1\}^\lambda$.
2. For each $i \in [n]$, the challenger samples the key pairs $(\text{vk}_0^{(i)}, \text{sk}_0^{(i)}) \leftarrow \text{Gen}(1^\lambda)$ and a punctured key pair $(\text{vk}_1^{(i)}, \text{sk}_1^{(i)}) \leftarrow \text{GenPunc}(1^\lambda, m^*)$. It gives the verification keys $\text{vk}_b^{(1)}, \dots, \text{vk}_b^{(n)}$ to the adversary.
3. The adversary can now make signature queries. Each signing query consists of an index $i \in [n]$ and a message $m \in \{0, 1\}^\lambda \setminus \{m^*\}$. The challenger responds with $\sigma \leftarrow \text{Sign}(\text{sk}_b^{(i)}, m)$.
4. The adversary outputs a bit $b' \in \{0, 1\}$, which is the output of the experiment.

We say that Π_{PunctSig} satisfies multiple verification key indistinguishability if for all efficient adversaries \mathcal{A} , there exists a negligible function $\text{negl}(\cdot)$ such that

$$\left| \Pr[\text{ExptMVKI}_{\mathcal{A}}(\lambda, 0) = 1] - \Pr[\text{ExptMVKI}_{\mathcal{A}}(\lambda, 1) = 1] \right| = \text{negl}(\lambda).$$

Fact 7.3 (Puncturable Signatures [GVW19]). Assuming either (1) the plain LWE assumption, or (2) the decision linear assumption in a pairing group, there exists a puncturable signature scheme.

Monotone policy aggregate signature. We now define the notion of a monotone policy aggregate (multi)-signature.

Definition 7.4 (Monotone Policy Aggregate Signatures). Let $\Pi_{\text{Sig}} = (\text{Gen}, \text{Sign}, \text{Verify})$ be a digital signature scheme with message space $\{0, 1\}^\lambda$. A monotone policy aggregation scheme for Π_{Sig} is a tuple of polynomial time algorithms $\Pi_{\text{Agg}} = (\text{Setup}, \text{Aggregate}, \text{AggVerify})$ with the following syntax:

- $\text{Setup}(1^\lambda, 1^k, 1^{s_p}) \rightarrow \text{crs}$: On input a security parameter λ , a bound on the number of signers k , and a bound s_p on the policy size, the setup algorithm outputs a common reference string crs .
- $\text{Aggregate}(\text{crs}, m, P, (\text{vk}_1, \sigma_1), \dots, (\text{vk}_k, \sigma_k)) \rightarrow \sigma_{\text{agg}}$: On input the common reference string crs , a message $m \in \{0, 1\}^\lambda$, a policy circuit $P: \{0, 1\}^k \rightarrow \{0, 1\}$, a collection of verification key/signature pairs (vk_i, σ_i) , the aggregation algorithm produces an aggregate signature σ_{agg} .
- $\text{AggVerify}(\text{crs}, m, P, (\text{vk}_1, \dots, \text{vk}_k), \sigma_{\text{agg}}) \rightarrow b$: On input the common reference string crs , a message m , a policy circuit $P: \{0, 1\}^k \rightarrow \{0, 1\}$, a tuple of k verification keys and an aggregate signature σ_{agg} , the aggregate verification algorithm outputs a bit $b \in \{0, 1\}$.

Moreover, Π_{Agg} must satisfy the following properties:

- **Correctness:** For all $\lambda, k, s_p \in \mathbb{N}$, all messages $m \in \{0, 1\}^\lambda$, all monotone circuits $P: \{0, 1\}^k \rightarrow \{0, 1\}$ and all key and signature tuples $\{(i, \text{vk}_i, \sigma_i)\}_{i \in [k]}$ where $P(\text{Verify}(\text{vk}_1, m, \sigma_1), \dots, \text{Verify}(\text{vk}_k, m, \sigma_k)) = 1$, it holds that

$$\Pr \left[\text{AggVerify}(\text{crs}, m, P, (\text{vk}_1, \dots, \text{vk}_k), \sigma_{\text{agg}}) = 1 \quad : \quad \begin{array}{l} \text{crs} \leftarrow \text{Setup}(1^\lambda, 1^k, 1^{s_p}) \\ \sigma_{\text{agg}} \leftarrow \text{Aggregate}(\text{crs}, m, P, (\text{vk}_1, \sigma_1), \dots, (\text{vk}_k, \sigma_k)) \end{array} \right] = 1.$$

- **Succinctness:** There exists a fixed polynomial $\text{poly}(\cdot)$ such that for all $\lambda, k, s_p \in \mathbb{N}$, all messages $m \in \{0, 1\}^\lambda$, all monotone circuits $P: \{0, 1\}^k \rightarrow \{0, 1\}$ and all pairs $\{(\text{vk}_i, \sigma_i)\}_{i \in [k]}$, the size of the aggregate signature σ_{agg} in the correctness experiment satisfies $|\sigma_{\text{agg}}| = \text{poly}(\lambda + \log |P|)$.

- **Static security:** For any adversary \mathcal{A} define the static unforgeability experiment $\text{ExptSU}_{\mathcal{A}}(\lambda)$ as follows:
 1. On input the security parameter λ , the adversary \mathcal{A} outputs the number of parties 1^k , a number of verification keys 1^n , the bound on the policy size 1^{s_p} , a challenge message $m^* \in \{0, 1\}^\lambda$, and a monotone policy $P: \{0, 1\}^k \rightarrow \{0, 1\}$.
 2. The challenger samples key-pairs $(\text{vk}_i, \text{sk}_i) \leftarrow \text{Gen}(1^\lambda)$ for all $i \in [n]$ and sends $\text{vk}_1, \dots, \text{vk}_k$ to the adversary.
 3. The adversary \mathcal{A} can now issue signing queries. Each signing query consists of an index $i \in [n]$ and a message $m \in \{0, 1\}^\lambda \setminus \{m^*\}$. The challenger responds with $\sigma \leftarrow \text{Sign}(\text{sk}_i, m)$.
 4. After the adversary is finished making signing queries, it outputs a tuple of verification keys $(\text{vk}_1^*, \dots, \text{vk}_k^*)$.
 5. The challenger replies with the common reference string $\text{crs} \leftarrow \text{Setup}(1^\lambda, 1^k, 1^{s_p})$.
 6. The adversary \mathcal{A} can continue to make signing queries. The challenger responds to these exactly as before.
 7. The adversary outputs the aggregate signature σ_{agg}^* .

8. The output of the experiment is 1 if all of the following holds:
- For each $i \in [k]$, let $b_i = 0$ if $vk_i^* = vk_j$ for some $j \in [n]$. Otherwise, let $b_i = 1$. Then, it holds that $P(b_1, \dots, b_k) = 0$.
 - $\text{AggVerify}(\text{crs}, m^*, P, (vk_1^*, \dots, vk_k^*), \sigma_{\text{agg}}^*) = 1$.
- Otherwise, the output is 0.

We say that Π_{Agg} satisfies static security if for every efficient adversary \mathcal{A} , there exists a negligible function $\text{negl}(\cdot)$ such that $\Pr[\text{ExptSU}_{\mathcal{A}}(\lambda) = 1] = \text{negl}(\lambda)$.

Aggregating puncturable signatures. We now show that we can combine any monotone policy BARG (satisfying non-adaptive soundness) with a puncturable signature scheme to obtain a statically-secure monotone policy aggregate signature scheme.

Construction 7.5 (Monotone Policy Aggregate Signature). The construction uses the following ingredients: let $\Pi_{\text{BARG}} = (\text{BARG.Gen}, \text{BARG.Prove}, \text{BARG.Verify})$ be a monotone policy BARG for NP and let $\Pi_{\text{PunctSig}} = (\text{PS.Gen}, \text{PS.GenPunc}, \text{PS.Sign}, \text{PS.Verify})$ be a puncturable signature scheme with message space $\{0, 1\}^\lambda$. Let $\ell_{vk} = \ell_{vk}(\lambda)$ be a bound on the length of the verification keys of Π_{PunctSig} . For any message $m \in \{0, 1\}^\lambda$, define the binary relation $\mathcal{R}[m]$ where

$$\mathcal{R}[m](vk, \sigma) = \begin{cases} 1 & \text{PS.Verify}(vk, m, \sigma) = 1 \\ 0 & \text{otherwise.} \end{cases}$$

Let C_m be the Boolean circuit that computes the relation $\mathcal{R}[m]$, and let $s_c = s_c(\lambda)$ be a bound on the size of C_m . We construct a monotone aggregate scheme $\Pi_{\text{Agg}} = (\text{Setup}, \text{Aggregate}, \text{AggVerify})$ for Π_{PunctSig} as follows:

- $\text{Setup}(1^\lambda, 1^k, 1^{s_p})$: Sample $\text{crs} \leftarrow \text{BARG.Gen}(1^\lambda, 1^{\ell_{vk}}, 1^{s_c}, 1^{s_p})$ and output crs .
- $\text{Aggregate}(\text{crs}, m, P, (vk_1, \sigma_1), \dots, (vk_k, \sigma_k))$: Output $\text{BARG.Prove}(\text{crs}, C_m, P, (vk_1, \dots, vk_k), (\sigma_1, \dots, \sigma_k))$.
- $\text{AggVerify}(\text{crs}, m, P, (vk_1, \dots, vk_k), \sigma_{\text{agg}})$: Output $\text{BARG.Verify}(\text{crs}, C_m, P, (vk_1, \dots, vk_k), \sigma_{\text{agg}})$.

Theorem 7.6 (Correctness). *If Π_{BARG} is complete, then [Construction 7.5](#) is correct.*

Proof. Fix $\lambda, k, s_p \in \mathbb{N}$, message $m \in \{0, 1\}^\lambda$, a monotone policy $P: \{0, 1\}^k \rightarrow \{0, 1\}$ and k tuples $\{(i, vk_i, \sigma_i)\}_{i \in [k]}$ such that $P(\text{Verify}(vk_1, m, \sigma_1), \dots, \text{Verify}(vk_k, m, \sigma_k)) = 1$. By construction of C_m , it holds that $C_m(vk_i, \sigma_i) = \text{Verify}(vk_i, m, \sigma_i)$ for all $i \in [k]$. Thus, $P(C_m(vk_1, \sigma_1), \dots, C_m(vk_k, \sigma_k)) = 1$. The theorem now follows by completeness of Π_{BARG} . \square

Theorem 7.7 (Succinctness). *If Π_{BARG} is succinct then [Construction 7.5](#) has succinct aggregate signatures.*

Proof. This follows directly from the succinctness of Π_{BARG} and the fact that the aggregate signature is simply a BARG proof. Fix $\lambda, k, s_p \in \mathbb{N}$, message $m \in \{0, 1\}^\lambda$, a monotone policy $P: \{0, 1\}^k \rightarrow \{0, 1\}$ and k tuples $\{(i, vk_i, \sigma_i)\}_{i \in [k]}$. The aggregate signature σ_{agg} is a BARG proof for circuit C_m , policy P , the statements (vk_1, \dots, vk_k) and the signatures $(\sigma_1, \dots, \sigma_k)$. By succinctness of the BARG, the length of σ_{agg} is $\text{poly}(\lambda + s_c + \log |P|)$. For every message $m \in \{0, 1\}^\lambda$, the circuit C_m simply checks whether the input verification key and signature verify the message m , so $s_c(\lambda) = \text{poly}(\lambda)$. Hence, the overall proof size is $\text{poly}(\lambda + \log |P|)$ and the claim follows. \square

Remark 7.8 (Fast Verification via RAM Delegation). Similar to [Remark 2.13](#) it is possible to use a RAM delegation scheme [[CJJ21b](#), [WW22](#), [KLVW23](#), [CGJ⁺23](#)] to delegate the aggregate signature verification to the aggregator. Currently, the aggregate verification algorithm AggVerify in [Construction 7.5](#) runs in time $\text{poly}(\lambda + |P|)$. This is because the aggregation algorithm needs to read the policy as well as the verification keys vk_1, \dots, vk_k . If the policy P and the verification keys are known in advance, the aggregator can include a proof π that the function $F_{\text{crs}, P, (vk_1, \dots, vk_k)}(m, \sigma_{\text{agg}}) := \text{AggVerify}(\text{crs}, m, P, (vk_1, \dots, vk_k), \sigma_{\text{agg}})$ satisfies $F_{\text{crs}, P, (vk_1, \dots, vk_k)}(m, \sigma_{\text{agg}}) = 1$. In this case, the common reference string would also contain a CRS for the RAM delegation scheme. The new aggregate verification algorithm would only check the RAM delegation proof (with respect to the function $F_{\text{crs}, P, (vk_1, \dots, vk_k)}$); formally, the

RAM delegation scheme would take as input a *hash* h of the parameters $(\text{crs}, P, (\text{vk}_1, \dots, \text{vk}_k))$, and the verification algorithm for the RAM program only needs to take the (honestly-precomputed) hash h , the message m , and the signature σ_{agg} . With this modification, the aggregate verification algorithm (given the precomputed hash h) runs in time $\text{poly}(\lambda + \log |P|)$.

Theorem 7.9 (Static Security). *If Π_{PunctSig} satisfies (multiple) verification key indistinguishability and Π_{BARG} satisfies non-adaptive soundness, then [Construction 7.5](#) is statically unforgeable.*

Proof. Let \mathcal{A} be any efficient adversary for the static security game. We begin by defining a sequence of hybrid experiments:

- Hyb_0 : This is the static unforgeability experiment:
 1. On input the security parameter λ , the adversary \mathcal{A} outputs the number of parties 1^k , a number of verification keys 1^n , a bound on the policy size 1^{s_p} , a challenge message $m^* \in \{0, 1\}^\lambda$, and a monotone policy $P: \{0, 1\}^k \rightarrow \{0, 1\}$.
 2. The challenger samples key-pairs $(\text{vk}_i, \text{sk}_i) \leftarrow \text{PS.Gen}(1^\lambda)$ for all $i \in [n]$ and sends $\text{vk}_1, \dots, \text{vk}_k$ to the adversary.
 3. The adversary \mathcal{A} can now issue signing queries. Each signing query consists of an index $i \in [n]$ and a message $m \in \{0, 1\}^\lambda \setminus \{m^*\}$. The challenger responds with $\sigma \leftarrow \text{PS.Sign}(\text{sk}_i, m)$.
 4. After the adversary is finished making signing queries, it outputs a tuple of verification keys $(\text{vk}_1^*, \dots, \text{vk}_k^*)$.
 5. The challenger replies with the common reference string $\text{crs} \leftarrow \text{BARG.Gen}(1^\lambda, 1^{\ell_{\text{vk}}}, 1^{s_c}, 1^{s_p})$.
 6. The adversary \mathcal{A} can continue to make signing queries. The challenger responds to these exactly as before. When \mathcal{A} finishes making signing queries, the adversary outputs the aggregate signature σ_{agg}^* .
 7. The output of the experiment is 1 if all of the following holds:
 - For each $i \in [k]$, let $b_i = 0$ if $\text{vk}_i^* = \text{vk}_j$ for some $j \in [n]$. Otherwise, let $b_i = 1$. Then, it holds that $P(b_1, \dots, b_k) = 0$.
 - $\text{BARG.Verify}(\text{crs}, C_{m^*}, P, (\text{vk}_1^*, \dots, \text{vk}_k^*), \sigma_{\text{agg}}^*) = 1$.
 Otherwise, the output is 0.
- Hyb_1 : Same as Hyb_0 , except the challenger uses the following modified procedure to sample key-pairs ([Step 2](#)):
 - For all $i \in [n]$, sample $(\text{vk}_i, \text{sk}_i) \leftarrow \text{PS.GenPunc}(1^\lambda, m^*)$.

For an adversary \mathcal{A} , we write $\text{Hyb}_i(\mathcal{A})$ to denote the output of Hyb_i with adversary \mathcal{A} . We now show that the output distributions of Hyb_0 and Hyb_1 are computationally indistinguishable, and moreover, that for all efficient adversaries \mathcal{A} , the output of $\text{Hyb}_1(\mathcal{A})$ is 1 with negligible probability.

Lemma 7.10. *If Π_{PunctSig} satisfies verification key indistinguishability, then there exists a negligible function $\text{negl}(\cdot)$ such that $|\Pr[\text{Hyb}_0(\mathcal{A}) = 1] - \Pr[\text{Hyb}_1(\mathcal{A}) = 1]| = \text{negl}(\lambda)$.*

Proof. Suppose $|\Pr[\text{Hyb}_0(\mathcal{A}) = 1] - \Pr[\text{Hyb}_1(\mathcal{A}) = 1]| = \varepsilon$ for some non-negligible ε . We construct an adversary \mathcal{B} for the multiple verification key indistinguishability game of Π_{PunctSig} as follows:

1. On input the security parameter 1^λ , algorithm \mathcal{B} computes $(1^k, 1^n, 1^{s_p}, m^*, P) \leftarrow \mathcal{A}(1^\lambda)$. Algorithm \mathcal{B} forwards m^* and 1^n , to the challenger.
2. The challenger replies with a tuple of verification keys $(\text{vk}_1, \dots, \text{vk}_n)$. Algorithm \mathcal{B} forwards $(\text{vk}_1, \dots, \text{vk}_n)$ to \mathcal{A} .
3. Whenever \mathcal{A} makes a signing query on an index $i \in [n]$ and a message $m \in \{0, 1\}^\lambda \setminus \{m^*\}$, algorithm \mathcal{B} forwards (i, m) to the challenger to obtain a signature σ . Algorithm \mathcal{B} replies to \mathcal{A} with σ .

4. When the adversary \mathcal{A} outputs a tuple (vk_1^*, \dots, vk_k^*) , algorithm \mathcal{B} computes the common reference string $crs \leftarrow \text{BARG.Gen}(1^\lambda, 1^{\ell_{vk}}, 1^{s_c}, 1^{s_p})$ and gives crs to \mathcal{A} .
5. Whenever algorithm \mathcal{A} makes additional signing queries, algorithm \mathcal{B} responds in the same manner as before.
6. When \mathcal{A} outputs a signature σ_{agg}^* , algorithm \mathcal{B} checks the following:
 - For each $i \in [k]$, let $b_i = 0$ if $vk_i^* = vk_j^*$ for some $j \in [n]$. Otherwise, let $b_i = 1$. Then, check that $P(b_1, \dots, b_k) = 0$.
 - $\text{AggVerify}(crs, m^*, P, (vk_1^*, \dots, vk_k^*), \sigma_{agg}^*) = 1$.

Algorithm \mathcal{B} outputs 1 if both checks pass and 0 otherwise.

By construction, algorithm \mathcal{B} constructs the key-pairs (vk_i, sk_i) for $i \in [k] \setminus \mathcal{I}$ exactly as required in Hyb_0 and Hyb_1 . It suffices to consider the distribution of the verification keys vk_i for $i \in \mathcal{I}$ and the responses to the signing queries. We consider the two possibilities:

- If the challenger responds according to the specification of $\text{ExptMVKI}_{\mathcal{B}}(\lambda, 0)$, then it samples $(vk_i, sk_i) \leftarrow \text{Gen}(1^\lambda)$. Moreover, the challenger responds to a signing query on (j, m) where $j \in [n]$ and $m \in \{0, 1\}^\lambda \setminus \{m^*\}$ with $\sigma \leftarrow \text{Sign}(sk_i, m)$. This is precisely the distribution in $\text{Hyb}_0(\mathcal{A})$. Finally, algorithm \mathcal{B} computes the output using the same procedure as in Hyb_0 and Hyb_1 . Therefore, $\Pr[\text{Hyb}_0(\mathcal{A}) = 1] = \Pr[\text{ExptMVKI}_{\mathcal{B}}(\lambda, 0) = 1]$.
- If the challenger responds according to the specification of $\text{ExptMVKI}_{\mathcal{B}}(\lambda, 1)$, then it samples $(vk_i, sk_i) \leftarrow \text{GenPunc}(1^\lambda, m^*)$. Moreover, the challenger responds to a signing query on (j, m) where $j \in [n]$ and $m \in \{0, 1\}^\lambda \setminus \{m^*\}$ with $\sigma \leftarrow \text{Sign}(sk_i, m)$. This is precisely the distribution in $\text{Hyb}_1(\mathcal{A})$. We conclude that $\Pr[\text{Hyb}_1(\mathcal{A}) = 1] = \Pr[\text{ExptMVKI}_{\mathcal{B}}(\lambda, 1) = 1]$.

We conclude that algorithm \mathcal{B} wins the multiple verification key indistinguishability game with the same non-negligible advantage ε , and the claim follows. \square

Lemma 7.11. *If Π_{BARG} satisfies non-adaptive soundness, and Π_{PunctSig} satisfies punctured correctness, then there exists a negligible function $\text{negl}(\cdot)$ such that $\Pr[\text{Hyb}_1(\mathcal{A}) = 1] = \text{negl}(\lambda)$.*

Proof. Suppose $\Pr[\text{Hyb}_1(\mathcal{A}) = 1] = \varepsilon$ for some non-negligible ε . We construct an adversary \mathcal{B} for the non-adaptive soundness game as follows:

1. On input the security parameter 1^λ , algorithm \mathcal{B} runs $(1^k, 1^n, 1^{s_p}, m^*, P) \leftarrow \mathcal{A}(1^\lambda)$. Algorithm \mathcal{B} then samples $(vk_i, sk_i) \leftarrow \text{PS.GenPunc}(1^\lambda, m^*)$ for all $i \in [n]$. It forwards the verification keys (vk_1, \dots, vk_n) .
2. Whenever algorithm \mathcal{A} makes a signing query on an index $i \in [n]$ and a message $m \in \{0, 1\}^\lambda \setminus \{m^*\}$, algorithm \mathcal{B} replies with a signature $\sigma \leftarrow \text{Sign}(sk_i, m)$.
3. When the adversary \mathcal{A} outputs a tuple (vk_1^*, \dots, vk_k^*) , algorithm \mathcal{B} forwards the instance size $1^{\ell_{vk}}$, the circuit size 1^{s_c} , the monotone policy size bound 1^{s_p} , the circuit C_{m^*} , the monotone policy P , and the instance (vk_1^*, \dots, vk_k^*) to the BARG challenger. The challenger replies with a common reference string crs which \mathcal{B} forwards to \mathcal{A} .
4. Whenever algorithm \mathcal{A} makes additional signing queries, algorithm \mathcal{B} responds in the same manner as before.
5. At the end of the game, algorithm \mathcal{A} outputs an aggregate signature σ_{agg} . Algorithm \mathcal{B} forwards $\pi = \sigma_{agg}$ to the challenger.

The challenger constructs the common reference string as $crs \leftarrow \text{BARG.Gen}(1^\lambda, 1^{\ell_{vk}}, 1^{s_c}, 1^{s_p})$. Thus, algorithm \mathcal{B} perfectly simulates an execution of Hyb_1 for \mathcal{A} . Thus, with probability at least ε , algorithm \mathcal{A} outputs an aggregate signature σ_{agg} where

$$\text{BARG.Verify}(crs, C_{m^*}, P, (vk_1^*, \dots, vk_k^*), \sigma_{agg}) = 1,$$

and $P(b_1, \dots, b_k) = 0$ where $b_i = 0$ if $vk_i^* = vk_j^*$ for some $j \in [n]$ and $b_i = 1$ otherwise. We argue that algorithm \mathcal{B} wins the non-adaptive soundness game when this happens:

- By punctured correctness of Π_{PunctSig} , for all $\sigma \in \{0, 1\}^*$, it holds that $\text{Verify}(\text{vk}_i, m^*, \sigma) = 0$ for all $i \in [n]$. Correspondingly, this means that for all $i \in [n]$, it holds that $C_{m^*}(\text{vk}_i, \sigma) = 0$ for all inputs $\sigma \in \{0, 1\}^*$.
- Thus for all $i \in [k]$, if $\text{vk}_i^* = \text{vk}_j$ for some $j \in [n]$, then $C_{m^*}(\text{vk}_i^*, \sigma) = 0$ for all $\sigma \in \{0, 1\}^*$. Next $P(b_1, \dots, b_k) = 0$ where $b_i = 0$ whenever $\text{vk}_i^* = \text{vk}_j$ for some $j \in [n]$, and $b_i = 1$ otherwise. This means $(C_{m^*}, P, (\text{vk}_1^*, \dots, \text{vk}_k^*)) \notin \mathcal{L}_{\text{MP-CSAT}}$.
- If $\text{BARG.Verify}(\text{crs}, C_{m^*}, P, (\text{vk}_1^*, \dots, \text{vk}_k^*), \sigma_{\text{agg}}) = 1$, and $(C_{m^*}, P, (\text{vk}_1^*, \dots, \text{vk}_k^*)) \notin \mathcal{L}_{\text{MP-CSAT}}$, then algorithm \mathcal{B} wins the non-adaptive soundness game.

Thus, algorithm \mathcal{B} breaks the non-adaptive soundness of Π_{BARG} with the same advantage ϵ . \square

[Theorem 7.9](#) now follows by [Lemmas 7.10](#) and [7.11](#) and a hybrid argument. \square

8 Semi-Somewhere Extractability of Monotone Policy BARGs

In this section, we show that our proof of non-adaptive soundness for our monotone policy BARG in [Section 4](#) easily extends to achieve a notion of *extractability*.¹⁰ Our notion of extractability is a relaxed version of the somewhere extractability notion from [\[BBK⁺23\]](#). In the notion from [\[BBK⁺23\]](#), there is a trapdoor setup algorithm that takes as input a set of indices S and outputs an extraction trapdoor. The guarantee is that whenever the prover produces a proof for a tuple of statements (x_1, \dots, x_k) with respect to a circuit C and policy P for which S is “critical,” then the extraction algorithm will output a witness w_i where $C(x_i, w_i) = 1$ for some index $i \in S$. In this setting, a set S is critical for a policy P if every input $(b_1, \dots, b_k) \in \{0, 1\}^k$ where $P(b_1, \dots, b_k) = 1$ has an index $i \in S$ where $b_i = 1$ (i.e., every input that satisfies the policy P must set some index in the critical set S to 1). In addition, the trapdoor CRS should hide the set S .

Semi-somewhere extractability. To extract a witness from the critical set S , the [\[BBK⁺23\]](#) construction program S into the CRS and then rely on an FHE-based hash function to homomorphically “propagate” one of the witnesses in S into the hash digest. This enables an efficient extraction procedure. In our setting, we do not use FHE. Instead, we observe that our existing proof in [Section 4](#) already achieves a notion of extractability by relying only on somewhere extractability of the underlying (vanilla) BARG. The caveat of our notion is that there is a $1/k$ loss in the success probability of our extractor. Namely, if an adversary produces a proof on (x_1, \dots, x_k) with probability ϵ , then the extractor will output a witness w_i for some $i \in S$ with probability ϵ/k . We refer to our notion as semi-somewhere extractability. We give the formal definition below:

Definition 8.1 (Semi-Somewhere Extractable Monotone BARG). A semi-somewhere extractable monotone policy BARG for Boolean circuit satisfiability is a tuple of polynomial time algorithms $\Pi_{\text{MP-BARG}} = (\text{Gen}, \text{Prove}, \text{Verify}, \text{TrapGen}, \text{Extract})$ such that $(\text{Gen}, \text{Prove}, \text{Verify})$ is monotone policy BARG for Boolean circuit satisfiability and the two additional algorithms $(\text{TrapGen}, \text{Extract})$ have the following syntax:

- $\text{TrapGen}(1^\lambda, 1^n, 1^{s_c}, 1^{s_p}, 1^k, S) \rightarrow (\text{crs}, \text{td})$: On input the security parameter $\lambda \in \mathbb{N}$, the instance size $n \in \mathbb{N}$, the number of instances $k \in \mathbb{N}$, a bound on the size of the Boolean circuit $s_c \in \mathbb{N}$, a bound on the size of the policy $s_p \in \mathbb{N}$, and a subset $S \subseteq [k]$, the indexed generator algorithm outputs a common reference string crs and a trapdoor td .
- $\text{Extract}(\text{td}, C, P, (x_1, \dots, x_k), \pi) \rightarrow (i, w_i)$: On input a trapdoor td , a Boolean circuit C , a monotone policy P , instances x_1, \dots, x_k , a proof π , and an index i , the extraction algorithm outputs an index i and an NP witness w_i .

Moreover, $\Pi_{\text{MP-BARG}}$ should satisfy the following properties:

¹⁰As we discussed in [Section 1.1](#), it is not clear what the right or most useful notion of extraction is in the context of monotone policy BARGs. The desired notion of extractability may in fact be application-dependent. For this reason, we focus on non-adaptive soundness for the main construction and include this section primarily as an illustration that our approach can support some non-trivial form of extractability.

- **Set hiding:** For an adversary \mathcal{A} and a bit $b \in \{0, 1\}$, define the set hiding experiment $\text{ExptSH}_{\mathcal{A}}(\lambda, b)$ as follows:
 1. On input a security parameter λ , algorithm \mathcal{A} starts by outputting the instance size 1^n , the bound on the size of the NP relation 1^{s_c} , the bound on the size of the policy 1^{s_p} , the number of instances 1^k , and a set $S \subseteq [k]$.
 2. If $b = 0$, the challenger samples $\text{crs} \leftarrow \text{Gen}(1^\lambda, 1^n, 1^{s_c}, 1^{s_p})$. Otherwise, if $b = 1$, the challenger samples $(\text{crs}, \text{td}) \leftarrow \text{TrapGen}(1^\lambda, 1^n, 1^{s_c}, 1^{s_p}, 1^k, S)$. The challenger sends crs to \mathcal{A} .
 3. Algorithm \mathcal{A} outputs a bit $b' \in \{0, 1\}$ which is the output of the experiment.

We say that $\Pi_{\text{MP-BARG}}$ satisfies set hiding if for every efficient adversary \mathcal{A} there exists a negligible function $\text{negl}(\cdot)$ such that

$$|\Pr[\text{ExptSH}_{\mathcal{A}}(\lambda, 0) = 1] - \Pr[\text{ExptSH}_{\mathcal{A}}(\lambda, 1) = 1]| = \text{negl}(\lambda).$$

- **Semi-somewhere extractability:** For an integer $k \in \mathbb{N}$ and an adversary \mathcal{A} , define the semi-somewhere extractability experiment $\text{ExptSE}_{\mathcal{A}}(\lambda, k)$ as follows:
 1. On input the security parameter 1^λ , algorithm \mathcal{A} starts by outputting the instance size 1^n , the bound on the size of the NP relation 1^{s_c} , the bound on the size of the policy 1^{s_p} , a monotone Boolean circuit $P: \{0, 1\}^k \rightarrow \{0, 1\}$ of size at most s_p , and a set $S \subseteq [k]$.
 2. The challenger samples $(\text{crs}, \text{td}) \leftarrow \text{Gen}(1^\lambda, 1^n, 1^{s_c}, 1^{s_p}, 1^k, S)$ and sends crs to \mathcal{A} .
 3. Algorithm \mathcal{A} outputs a Boolean circuit $C: \{0, 1\}^n \times \{0, 1\}^h \rightarrow \{0, 1\}$ of size at most s_c , statements $x_1, \dots, x_k \in \{0, 1\}^n$, and a proof π .
 4. The challenger extracts a witness $(i, w_i) \leftarrow \text{Extract}(\text{td}, C, P, (x_1, \dots, x_k), \pi)$.
 5. The output of the experiment is 1 if $C(x_i, w_i) = 1$ and $i \in S$. Otherwise, the output is 0.

An adversary \mathcal{A} is admissible if it outputs a set $S \subseteq [k]$ and a policy P such that $P(b_1, \dots, b_k) = 0$ where $b_i = 0$ if $i \in S$ and $b_i = 1$ otherwise. Let

$$\varepsilon_{\mathcal{A}}(\lambda, k) := \Pr[\text{Verify}(\text{crs}, C, P, (x_1, \dots, x_k), \pi) = 1] \tag{8.1}$$

in an execution of $\text{ExptSE}_{\mathcal{A}}(\lambda, k)$. We say that Π_{BARG} is semi-somewhere extractable if for every polynomial $k = k(\lambda)$ and every efficient and admissible adversary \mathcal{A} , there exists a negligible function $\text{negl}(\cdot)$ such that

$$\Pr[\text{ExptSE}_{\mathcal{A}}(\lambda, k) = 1] \geq \frac{1}{k} \cdot \varepsilon_{\mathcal{A}}(\lambda, k) - \text{negl}(\lambda).$$

Remark 8.2 (On Semi-Somewhere Extractability). An important caveat of the semi-somewhere extractability notion in [Definition 8.1](#) is that we allow the extractor to succeed with *smaller* probability (by an inverse polynomial factor) than the honest prover. While this is still sufficient for applications to monotone policy aggregate signatures (see [Section 8.1](#)), this may not be the case in all settings where an extraction guarantee might be employed. As an example, suppose we have an adversary \mathcal{A} that samples statements from one of two distributions \mathcal{D}_1 and \mathcal{D}_2 (with equal probability) and produces a valid proof on the statement with probability ε . Normally, we would hope that the extractor algorithm would be able to extract witnesses for statements sampled from *both* \mathcal{D}_1 and \mathcal{D}_2 . However, since we allow for an inverse polynomial loss in the extractor's success probability, it could be the case that the extractor *only* works for instances sampled from \mathcal{D}_1 and never outputs witnesses for instances sampled from \mathcal{D}_2 . If this were to happen in a security proof which relies on the ability to extract witnesses from instances drawn from \mathcal{D}_2 , then the proof would no longer go through. Thus, using the semi-somewhere extractability notion in the context of a security proof could require some extra care.

Adapting Construction 4.4. We now show how to extend [Construction 4.4](#) to support semi-somewhere extractability. The construction relies on the fact that the proof of [Construction 4.4 \(Section 4.2\)](#) implicitly achieves a notion of extractability.

Construction 8.3 (Semi-Somewhere Extractable Monotone BARG). Let $(\text{Gen}, \text{Prove}, \text{Verify})$ be the monotone policy BARG of [Construction 4.4](#). Let $\Pi'_{\text{BARG}} = (\text{Gen}', \text{Prove}', \text{Verify}', \text{TrapGen}', \text{Extract}')$ be the underlying somewhere extractable BARG for Boolean circuit satisfiability, and $\Pi_{\text{H}} = (\text{H.Setup}, \text{H.Hash}, \text{H.ProveOpen}, \text{H.VerOpen}, \text{H.Extract}, \text{H.ValidateDigest})$ be the underlying zero-fixing hash function. We extend [Construction 4.4](#) with the following algorithms:

- $\text{TrapGenIndex}(1^\lambda, 1^n, 1^{s_c}, 1^{s_p}, 1^k, S, j) \rightarrow (\text{crs}, \text{td}_j)$: On input the security parameter $\lambda \in \mathbb{N}$, the instance size $n \in \mathbb{N}$, a bound on the size of the Boolean circuit $s_c \in \mathbb{N}$, a bound on the size of the policy $s_p \in \mathbb{N}$, the number of instances k , a set $S \subseteq [k]$, and an index $j \in S$, the indexed trapdoor generator algorithm proceeds as follows:
 - Let $j_1, \dots, j_{|S|} \in S$ be the elements of S in ascending order. Let $t \in [|S|]$ be the index where $j = j_t$. Define the set $S_t = \{j_1, \dots, j_{t-1}\}$ if $t > 1$ and $S_t = \emptyset$ otherwise.
 - Sample two hash keys $(\text{hk}_0, \text{vk}_0, \text{td}_0) \leftarrow \text{H.Setup}(1^\lambda, 1^{s_p}, \emptyset)$ and $(\text{hk}_1, \text{vk}_1, \text{td}_1) \leftarrow \text{H.Setup}(1^\lambda, 1^{s_p}, S_t)$.
 - Let s' be a bound on the size of the circuit that computes the relation $\mathcal{R}[C, k, s_p, \text{vk}_0, \text{vk}_1, \text{dig}_0, \text{dig}_1]$ from [Fig. 1](#) when instantiated with an arbitrary Boolean circuit C of size at most s_c , an input length $k \leq s_p$ and digests $\text{dig}_0, \text{dig}_1$ associated with the hash and verification keys $(\text{hk}_0, \text{vk}_0)$ and $(\text{hk}_1, \text{vk}_1)$. Let $n' = 3 \cdot \lceil \log s_p \rceil + 1$ be the bound on the statement length. Sample $(\text{crs}_{\text{BARG}}, \text{vk}_{\text{BARG}}, \text{td}_{\text{BARG}}) \leftarrow \text{TrapGen}'(1^\lambda, 1^{s_p}, 1^{n'}, 1^{s'}, j)$.
 - Outputs the common reference string $\text{crs} = (\text{crs}_{\text{BARG}}, \text{vk}_{\text{BARG}}, \text{hk}_0, \text{hk}_1, \text{vk}_0, \text{vk}_1)$ and the trapdoor $\text{td}_j = \text{td}_{\text{BARG}}$.

Looking ahead, the helper algorithm $\text{TrapGenIndex}(1^\lambda, 1^n, 1^{s_c}, 1^{s_p}, 1^k, S, j_t)$ implements the setup algorithm according to the specification of the hybrid experiments $\text{Hyb}_{0,t,3}$ and $\text{Hyb}_{0,t,4}$ in the proof of [Theorem 8.7](#). These are the analogs of the hybrid experiments $\text{Hyb}_{0,t,3}$ and $\text{Hyb}_{0,t,4}$ from the proof of [Theorem 4.7 in Section 4.2](#).

- $\text{TrapGen}(1^\lambda, 1^n, 1^{s_c}, 1^{s_p}, 1^k, S) \rightarrow (\text{crs}, \text{td})$: On input the security parameter $\lambda \in \mathbb{N}$, the instance size $n \in \mathbb{N}$, a bound on the size of the Boolean circuit $s_c \in \mathbb{N}$, a bound on the size of the policy $s_p \in \mathbb{N}$, the number of instances k , and a set $S \subseteq [k]$, the generator algorithm proceeds as follows:
 - Sample a random $j \xleftarrow{\mathbb{R}} S$.
 - Compute $(\text{crs}, \text{td}_j) \leftarrow \text{TrapGenIndex}(1^\lambda, 1^n, 1^{s_c}, 1^{s_p}, 1^k, j)$.
 - Output the common reference string crs , and the trapdoor $\text{td} = (j, \text{td}_j)$.
- $\text{Extract}(\text{td}, C, P, (x_1, \dots, x_k), \pi) \rightarrow (j, w_j)$: On input a trapdoor $\text{td} = (j, \text{td}_j)$, a Boolean circuit C , a monotone policy P , instances x_1, \dots, x_k , and a proof π , the algorithm computes $\hat{w}_j \leftarrow \text{Extract}(\text{td}_j, C_{\text{aug}}, (\hat{x}_1, \dots, \hat{x}_{s_p}), \pi, j)$, where the circuit C_{aug} and the instances $\hat{x}_1, \dots, \hat{x}_{s_p}$ are computed from C, P, x_1, \dots, x_k as in [Construction 4.4](#). Parse $\hat{w}_j = (b, \sigma^{(0)}, \sigma^{(1)}, w)$ and if $b = 1$, output (j, w) . If $b \neq 1$, output \perp .

Theorem 8.4 (Set Hiding). *If Π_{BARG} satisfies index hiding and Π_{H} satisfies set hiding, then [Construction 8.3](#) satisfies set hiding.*

Proof. Let \mathcal{A} be an efficient non-uniform adversary for the set hiding game of $\Pi_{\text{MP-BARG}}$. We proceed via a hybrid argument:

- Hyb_0 : This is experiment $\text{ExptSH}_{\mathcal{A}}[\lambda, 0]$:
 1. On input a security parameter λ , algorithm \mathcal{A} starts by outputting the instance size 1^n , the bound on the size of the NP relation 1^{s_c} , the bound on the size of the policy 1^{s_p} , the number of instances 1^k , and a set $S \subseteq [k]$.

2. The challenger samples $\text{crs} \leftarrow \text{Gen}(1^\lambda, 1^n, 1^{s_c}, 1^{s_p})$. Namely, it samples
 - $(\text{hk}_0, \text{vk}_0, \text{td}_0) \leftarrow \text{H.Setup}(1^\lambda, 1^{s_p}, \emptyset)$.
 - $(\text{hk}_1, \text{vk}_1, \text{td}_1) \leftarrow \text{H.Setup}(1^\lambda, 1^{s_p}, \emptyset)$.
 - $(\text{crs}_{\text{BARG}}, \text{vk}_{\text{BARG}}) \leftarrow \text{Gen}'(1^\lambda, 1^{s_p}, 1^{n'}, 1^{s'})$, where n', s' are defined as in [Construction 8.3](#).

The challenger sends $\text{crs} = (\text{crs}_{\text{BARG}}, \text{vk}_{\text{BARG}}, \text{hk}_0, \text{hk}_1, \text{vk}_0, \text{vk}_1)$ to \mathcal{A} .

3. Algorithm \mathcal{A} outputs a bit $b' \in \{0, 1\}$ which is the output of the experiment.

- Hyb_1 : Same as Hyb_0 , except the challenger samples crs_{BARG} to bind to a random index $j_t \in S$. Concretely, the challenger samples a random index $j \xleftarrow{\mathcal{R}} S$ and samples $(\text{crs}_{\text{BARG}}, \text{vk}_{\text{BARG}}, \text{td}_{\text{BARG}}) \leftarrow \text{TrapGen}'(1^\lambda, 1^{s_p}, 1^{n'}, 1^{s'}, j)$.
- Hyb_2 : Same as Hyb_1 , except the challenger samples hk_1, vk_1 to be zero-fixing on the set S_t . This is $\text{ExptSH}_{\mathcal{A}}[\lambda, 1]$. Specifically, in this experiment, the challenger samples the keys for the zero-fixing hash $(\text{hk}_1, \text{vk}_1, \text{td}_1) \leftarrow \text{H.Setup}(1^\lambda, 1^{s_p}, S_t)$, where $S_t = \{j_1, \dots, j_{t-1}\}$, the indices $j_1, \dots, j_{|S|}$ are the elements of S in ascending order, and $t \in [|S|]$ is the index where $j = j_t$.

We write $\text{Hyb}_i(\mathcal{A})$ to denote the output of Hyb_i with adversary \mathcal{A} . We now show that each pair of adjacent output distributions are computationally indistinguishable.

Claim 8.5. *If Π_{BARG} satisfies index hiding, then there exists a negligible function $\text{negl}(\cdot)$ such that*

$$|\Pr[\text{Hyb}_0(\mathcal{A}) = 1] - \Pr[\text{Hyb}_1(\mathcal{A}) = 1]| = \text{negl}(\lambda).$$

Proof. Suppose $|\Pr[\text{Hyb}_0(\mathcal{A}) = 1] - \Pr[\text{Hyb}_1(\mathcal{A}) = 1]| = \varepsilon$ for some non-negligible ε . We use \mathcal{A} to construct an adversary \mathcal{B} for the index hiding game of Π_{BARG} as follows:

1. On input the security parameter 1^λ , algorithm \mathcal{B} runs \mathcal{A} on input 1^λ to obtain $(1^n, 1^{s_c}, 1^{s_p}, 1^k, S)$.
2. Algorithm \mathcal{B} samples $j \xleftarrow{\mathcal{R}} S$ and send $(1^{s_p}, 1^{n'}, 1^{s'}, j)$ to the BARG challenger. The challenger replies with $(\text{crs}_{\text{BARG}}, \text{vk}_{\text{BARG}})$.
3. Algorithm \mathcal{B} computes $(\text{hk}_0, \text{vk}_0, \text{td}_0) \leftarrow \text{H.Setup}(1^\lambda, 1^{s_p}, \emptyset)$ and $(\text{hk}_1, \text{vk}_1, \text{td}_1) \leftarrow \text{H.Setup}(1^\lambda, 1^{s_p}, \emptyset)$.
4. Algorithm \mathcal{B} sets $\text{crs} = (\text{crs}_{\text{BARG}}, \text{vk}_{\text{BARG}}, \text{hk}_0, \text{hk}_1, \text{vk}_0, \text{vk}_1)$ and gives crs to \mathcal{A} . It outputs whatever \mathcal{A} outputs.

If the challenger samples $(\text{crs}_{\text{BARG}}, \text{vk}_{\text{BARG}}) \leftarrow \text{Gen}'(1^\lambda, 1^{s_p}, 1^{n'}, 1^{s'})$, then algorithm \mathcal{B} perfectly simulates Hyb_0 for \mathcal{A} . Conversely, if it samples $(\text{crs}_{\text{BARG}}, \text{vk}_{\text{BARG}}, \text{td}_{\text{BARG}}) \leftarrow \text{TrapGen}'(1^\lambda, 1^{s_p}, 1^{n'}, 1^{s'}, j)$, algorithm \mathcal{B} perfectly simulates Hyb_1 for \mathcal{A} . We conclude that the advantage of algorithm \mathcal{B} is ε , and the claim holds. \square

Claim 8.6. *If Π_{H} satisfies set hiding, then there exists a negligible function $\text{negl}(\cdot)$ such that*

$$|\Pr[\text{Hyb}_1(\mathcal{A}) = 1] - \Pr[\text{Hyb}_2(\mathcal{A}) = 1]| = \text{negl}(\lambda).$$

Proof. Suppose $|\Pr[\text{Hyb}_1(\mathcal{A}) = 1] - \Pr[\text{Hyb}_2(\mathcal{A}) = 1]| = \varepsilon$ for some non-negligible ε . We construct an attacker \mathcal{B} for the set hiding game of Π_{H} as follows:

1. On input the security parameter 1^λ , algorithm \mathcal{B} runs \mathcal{A} on input 1^λ to obtain $(1^n, 1^{s_c}, 1^{s_p}, 1^k, S)$.
2. Let $j_1, \dots, j_{|S|} \in S$ be the elements of S in ascending order. Algorithm \mathcal{B} samples $j \xleftarrow{\mathcal{R}} S$ and sets $t \in [|S|]$ to be the index where $j = j_t$. Algorithm \mathcal{B} send $(1^{s_p}, S_t)$ to the challenger, where $S_t = \{j_1, \dots, j_{t-1}\}$ if $t > 1$ and $S_t = \emptyset$ otherwise. The challenger replies with $(\text{hk}_1, \text{vk}_1)$.
3. Algorithm \mathcal{B} samples a hash key $(\text{hk}_0, \text{vk}_0, \text{td}_0) \leftarrow \text{H.Setup}(1^\lambda, 1^{s_p}, \emptyset)$ along with the BARG parameters $(\text{crs}_{\text{BARG}}, \text{vk}_{\text{BARG}}, \text{td}_{\text{BARG}}) \leftarrow \text{TrapGen}'(1^\lambda, 1^{s_p}, 1^{n'}, 1^{s'}, j)$.
4. Algorithm \mathcal{B} sets $\text{crs} = (\text{crs}_{\text{BARG}}, \text{vk}_{\text{BARG}}, \text{hk}_0, \text{hk}_1, \text{vk}_0, \text{vk}_1)$ and gives crs to \mathcal{A} . It outputs whatever \mathcal{A} outputs.

If the challenger samples $(hk_1, vk_1, td_1) \leftarrow H.Setup(1^\lambda, 1^{s_p}, \emptyset)$, then algorithm \mathcal{B} perfectly simulates Hyb_1 for \mathcal{A} . Conversely, if the challenger samples $(hk_1, vk_1, td_1) \leftarrow H.Setup(1^\lambda, 1^{s_p}, S_t)$, then algorithm \mathcal{B} perfectly simulates Hyb_2 for \mathcal{A} . Correspondingly, the advantage of algorithm \mathcal{B} is ε , and the claim follows. \square

Theorem 8.4 now follows from **Claims 8.5** and **8.6**. \square

Theorem 8.7. *If Π_H satisfies set hiding, index hiding with extracted guess, zero fixing and extractor validity against non-uniform adversaries, and Π_{BARG} is somewhere extractable and satisfies set hiding against non-uniform adversaries, then **Construction 8.3** is semi-somewhere extractable against non-uniform adversaries.*

Proof. To prove **Theorem 8.7**, we use a similar strategy as in the proof of **Theorem 4.7** (**Section 4.2**). Here, we give a high-level overview. Specifically, we start by defining sequence of hybrids $\text{Hyb}_0, \dots, \text{Hyb}_d$, where d is the depth of the monotone circuit P . These are essentially the same experiments from the proof of **Theorem 4.7** in **Section 4.2**. The initial hybrid corresponds to the semi-somewhere extractability experiment where the output is 1 if the adversary outputs an accepting proof (i.e., the output in the initial hybrid is 1 with probability $\varepsilon_{\mathcal{A}}(\lambda, k)$ as defined in **Eq. (8.1)**). In the final hybrid, we show that the output is 1 probability 0. Finally, we argue that any difference in advantage between adjacent hybrids can *only* occur in settings where the extractor is successful. There are a maximum of k such experiments (one associated with each of the inputs to P). Since the probability of an experiments drops from ε to 0, in at least one of these intermediary experiments, the probability must decrease by ε/k (up to negligible differences); this directly translates into the extractor succeeding with probability at least ε/k (up to negligible differences). We give the formal argument below.

Outer hybrids. Take any polynomial $k = k(\lambda)$ and any efficient (non-uniform) and admissible adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$. Let $P: \{0, 1\}^k \rightarrow \{0, 1\}$ and $S \subseteq [k]$ be the monotone policy and the challenge set that algorithm \mathcal{A}_1 outputs (on input the security parameter λ). Let d be the depth of P and s be its size. We now define the sequence of outer hybrids:

- $\widetilde{\text{Hyb}}_0$: This is the analog of Hyb_0 from the proof of **Theorem 4.7** (**Section 4.2**). We define it here:
 - **Phase 1:** On input the security parameter 1^λ , algorithm \mathcal{A}_1 outputs $1^n, 1^{s_c}, 1^{s_p}$, a monotone Boolean circuit $P: \{0, 1\}^k \rightarrow \{0, 1\}$ of size $s \leq s_p$, a set $S \subseteq [k]$ and the state $\text{st}_{\mathcal{A}}$. **The experiment outputs 0 if $P(b_1, \dots, b_k) = 1$ where $b_i = 0$ if $i \in S$ and $b_i = 1$ otherwise.**
 - **Phase 2:** The challenger computes $\text{crs} \leftarrow \text{Gen}(1^\lambda, 1^n, 1^{s_c}, 1^{s_p})$. Specifically, the challenger samples the following components:
 - * $(\text{crs}_{\text{BARG}}, \text{vk}_{\text{BARG}}) \leftarrow \text{Gen}'(1^\lambda, 1^{s_p}, 1^{n'}, 1^{s'})$.
 - * $(\text{hk}_0, \text{vk}_0, \text{td}_0) \leftarrow H.Setup(1^\lambda, 1^{s_p}, \emptyset)$.
 - * $(\text{hk}_1, \text{vk}_1, \text{td}_1) \leftarrow H.Setup(1^\lambda, 1^{s_p}, \emptyset)$.
The challenger sets $\text{crs} = (\text{crs}_{\text{BARG}}, \text{hk}_0, \text{hk}_1, \text{vk}_0, \text{vk}_1)$ and runs \mathcal{A}_2 on input $(\text{crs}, \text{st}_{\mathcal{A}})$. Algorithm \mathcal{A}_2 outputs **a Boolean circuit C of size at most s_c , an instance $\mathbf{x} = (x_1, \dots, x_k)$, and a proof string $\pi = (\text{dig}_0, \text{dig}_1, \pi_{\text{BARG}})$** . Let $\hat{\mathbf{x}} = (\hat{x}_1, \dots, \hat{x}_{s_p})$ and C_{aug} be as defined in **Prove and Verify** in **Construction 4.4**. The output of the experiment is 1 if all of the following conditions hold (and 0 otherwise):
 - * $H.ValidateDigest(\text{vk}_0, \text{dig}_0) = H.ValidateDigest(\text{vk}_1, \text{dig}_1) = 1$.
 - * $\text{Verify}'(\text{vk}_{\text{BARG}}, C_{\text{aug}}, \hat{\mathbf{x}}, \pi_{\text{BARG}}) = 1$.
- $\widetilde{\text{Hyb}}_i$ for $i \in [d]$: Same as Hyb_0 , but hk_{low} binds on J_i , where $\text{low} = i \bmod 2$ and $\text{high} = 1 - \text{low}$. This is the analog of Hyb_1 from the proof of **Theorem 4.7** (**Section 4.2**). Specifically, the game proceeds as follows:
 - **Phase 1:** On input the security parameter 1^λ , algorithm \mathcal{A}_1 outputs $1^n, 1^{s_c}, 1^{s_p}$, a monotone Boolean circuit $P: \{0, 1\}^k \rightarrow \{0, 1\}$ of size $s \leq s_p$, a set $S \subseteq [k]$ and the state $\text{st}_{\mathcal{A}}$. The experiment outputs 0 if $P(b_1, \dots, b_k) = 1$ where $b_i = 0$ if $i \in S$ and $b_i = 1$ otherwise. The challenger then computes the following quantities:
 - * **For each $j \in [k]$, let $\beta_j = 1$ if $j \in S$ and $\beta_j = 0$ otherwise.**

- * For $j \in [k+1, s]$ let β_j be the value of the wire j in the evaluation of P on input $(\beta_1, \dots, \beta_k)$.
 - * For each layer $\ell \in [d]$, let $J_\ell = \{j \in \text{layer}_\ell(P) : \beta_j = 0\}$.
- **Phase 2:** The challenger samples the following components:
- * $(\text{crs}_{\text{BARG}}, \text{vk}_{\text{BARG}}) \leftarrow \text{Gen}'(1^\lambda, 1^{s_p}, 1^{n'}, 1^{s'})$.
 - * $(\text{hk}_{\text{low}}, \text{vk}_{\text{low}}, \text{td}_{\text{low}}) \leftarrow \text{H.Setup}(1^\lambda, 1^{s_p}, J_i)$.
 - * $(\text{hk}_{\text{high}}, \text{vk}_{\text{high}}, \text{td}_{\text{high}}) \leftarrow \text{H.Setup}(1^\lambda, 1^{s_p}, \emptyset)$.

The challenger sets $\text{crs} = (\text{crs}_{\text{BARG}}, \text{hk}_0, \text{hk}_1, \text{vk}_0, \text{vk}_1)$ and runs \mathcal{A}_2 on input $(\text{crs}, \text{st}_{\mathcal{A}})$. Algorithm \mathcal{A}_2 outputs a proof string $\pi = (\text{dig}_0, \text{dig}_1, \pi_{\text{BARG}})$. Let $\hat{\mathbf{x}} = (\hat{x}_1, \dots, \hat{x}_{s_p})$ and C_{aug} be as defined in Prove and Verify in [Construction 4.4](#). The output of the experiment is 1 if all of the following conditions hold (and 0 otherwise):

- * $\text{H.ValidateDigest}(\text{vk}_0, \text{dig}_0) = \text{H.ValidateDigest}(\text{vk}_1, \text{dig}_1) = 1$.
- * $\text{Verify}'(\text{vk}_{\text{BARG}}, C_{\text{aug}}, \hat{\mathbf{x}}, \pi_{\text{BARG}}) = 1$.
- * $\text{H.Extract}(\text{td}_{\text{low}}, \text{dig}_{\text{low}}) = \text{Matching}$.

Inner games. To argue that each pair of outer hybrids is computationally indistinguishable, we define a sequence of “inner hybrids” exactly as in the proof of [Theorem 4.7](#) in [Section 4.2](#). Specifically, for each $i \in \{0, \dots, d\}$, each $t \in [|J_i|]$ and each $\ell \in \{1, \dots, 7\}$, we define $\widetilde{\text{Hyb}}_{i,t,\ell}$ as follows:

- **Phase 1:** Same as $\widetilde{\text{Hyb}}_i$. Note that algorithm \mathcal{A}_1 does not output the Boolean circuit C or the statements \mathbf{x} in this phase.
- **Phase 2:** Same as in $\widetilde{\text{Hyb}}_{i,t,\ell}$ from [Section 4.2](#), *except* the adversary additionally outputs the Boolean circuit C and the instances \mathbf{x} in this phase (as in $\widetilde{\text{Hyb}}_i$).

We now analyze each pair of hybrid experiments. With the exception of one of the transitions (from $\widetilde{\text{Hyb}}_{0,t,3}$ to $\widetilde{\text{Hyb}}_{0,t,4}$), each transition follows by a similar argument as the corresponding transition in the proof of [Theorem 4.7](#).

Claim 8.8. *If Π_{H} satisfies extractor validity against efficient non-uniform adversaries, then there exists a negligible function $\text{negl}(\cdot)$ such that for every $i \in \{0, \dots, d-1\}$, it holds that*

$$\left| \Pr[\widetilde{\text{Hyb}}_i(\mathcal{A}) = 1] - \Pr[\widetilde{\text{Hyb}}_{i,1,1}(\mathcal{A}) = 1] \right| \leq \text{negl}(\lambda).$$

Proof. Follows by a similar argument as the proof of [Claim 4.8](#). □

Claim 8.9. *If Π_{BARG} satisfies set hiding against efficient non-uniform adversaries, then there exists a negligible function $\text{negl}(\cdot)$ such that for every $i \in \{0, \dots, d-1\}$ and $t \in [|J_{i+1}|]$, it holds that*

$$\left| \Pr[\widetilde{\text{Hyb}}_{i,t,1}(\mathcal{A}) = 1] - \Pr[\widetilde{\text{Hyb}}_{i,t,2}(\mathcal{A}) = 1] \right| \leq \text{negl}(\lambda).$$

Proof. Follows by a similar argument as the proof of [Claim 4.9](#). □

Claim 8.10. *If Π_{BARG} satisfies somewhere extractability in trapdoor mode against efficient non-uniform adversaries, then there exists a negligible function $\text{negl}(\cdot)$ such that for every $i \in \{0, \dots, d-1\}$, $t \in [|J_{i+1}|]$, it holds that*

$$\left| \Pr[\widetilde{\text{Hyb}}_{i,t,2}(\mathcal{A}) = 1] - \Pr[\widetilde{\text{Hyb}}_{i,t,3}(\mathcal{A}) = 1] \right| \leq \text{negl}(\lambda).$$

Proof. Follows by a similar argument as the proof of [Claim 4.10](#). □

Extracting valid witnesses. As mentioned above the transition from $\widetilde{\text{Hyb}}_{0,t,3}$ to $\widetilde{\text{Hyb}}_{0,t,4}$ diverges from the corresponding analysis (Claim 4.11) in the proof of Theorem 4.7. In Claim 4.11, the relevant statement $x_{J_1[t]}$ was false, and thus, by somewhere extractability of the BARG, we were able to argue that the outputs of $\text{Hyb}_{0,t,3}$ to $\text{Hyb}_{0,t,4}$ could only change by a negligible amount. Upon closer inspection, the proof of Claim 4.11 actually shows a *stronger* property: the difference between these two hybrids is exactly equal to the probability of extracting a valid witness for the instance $x_{J_1[t]}$. In the case of Claim 4.11, the statement $x_{J_1[t]}$ was false, so this probability was identically 0. In the somewhere extractability game, this probability could be noticeable. But that means our extractor succeeds with noticeable probability. To formalize this, we start with a full specification of $\widetilde{\text{Hyb}}_{0,t,3}$:

- **Phase 1:** Same as $\widetilde{\text{Hyb}}_i$.
- **Phase 2:** The challenger samples the following components.
 - $(\text{crs}_{\text{BARG}}, \text{td}_{\text{BARG}}) \leftarrow \text{TrapGen}'(1^\lambda, 1^{s_p}, 1^{n'}, 1^{s'}, J_1[t])$.
 - $(\text{hk}_{\text{low}}, \text{vk}_{\text{low}}, \text{td}_{\text{low}}) \leftarrow \text{H.Setup}(1^\lambda, 1^{s_p}, \emptyset)$.
 - $(\text{hk}_{\text{high}}, \text{vk}_{\text{high}}, \text{td}_{\text{high}}) \leftarrow \text{H.Setup}(1^\lambda, 1^{s_p}, J_1[1, \dots, t-1])$.

The challenger sets $\text{crs} = (\text{crs}_{\text{BARG}}, \text{hk}_0, \text{hk}_1, \text{vk}_0, \text{vk}_1)$ and runs \mathcal{A}_2 on input $(\text{crs}, \text{st}_{\mathcal{A}})$. Algorithm \mathcal{A}_2 outputs a proof string $\pi = (\text{dig}_0, \text{dig}_1, \pi_{\text{BARG}})$. Let $\hat{x} = (\hat{x}_1, \dots, \hat{x}_{s_p})$ and C_{aug} be as defined in Prove and Verify in Construction 4.4. The challenger then computes $\hat{w} = (b, \sigma^{(0)}, \sigma^{(1)}, w) \leftarrow \text{Extract}'(\text{td}_{\text{BARG}}, C_{\text{aug}}, \hat{x}, \pi_{\text{BARG}}, J_1[t])$. The output is 1 if the following conditions hold (and 0 otherwise):

- $\text{H.ValidateDigest}(\text{vk}_0, \text{dig}_0) = \text{H.ValidateDigest}(\text{vk}_1, \text{dig}_1) = 1$.
- $\text{Verify}'(\text{vk}_{\text{BARG}}, C_{\text{aug}}, \hat{x}, \pi_{\text{BARG}}) = 1$.
- $\text{H.Extract}(\text{td}_{\text{high}}, \text{dig}_{\text{high}}) = \text{Matching}$.
- $C_{\text{aug}}(\hat{x}_{j_t}, \hat{w}) = 1$.

Since $S = J_1$ by construction, the challenger in $\widetilde{\text{Hyb}}_{0,t,3}$ is sampling the common reference string crs according to the specification of $\text{TrapGenIndex}(1^\lambda, 1^n, 1^{s_c}, 1^{s_p}, 1^k, S, t)$. Similarly, the witness $\hat{w} = (b, \sigma^{(0)}, \sigma^{(1)}, w)$ in $\widetilde{\text{Hyb}}_{0,t,3}$ is computed using the same procedure as $\text{Extract}(\text{td}, C, P, (x_1, \dots, x_k), \pi)$ where $\text{td} = (J_1[t], \text{td}_{\text{BARG}})$. For $t \in [J_1]$, let

$$p_t := \Pr \left[\widetilde{\text{Hyb}}_{0,t,3}(\mathcal{A}) = 1 \wedge b = 1 \text{ in the execution of } \widetilde{\text{Hyb}}_{0,t,3} \right].$$

By definition of C_{aug} , if $\widetilde{\text{Hyb}}_{0,t,3}(\mathcal{A})$ outputs 1 and the extracted bit b satisfies $b = 1$, this means that $C(x_{J_1[t]}, w) = 1$. In this case, $\text{Extract}(\text{td}, C, P, (x_1, \dots, x_k), \pi)$ outputs w such that $C(x_{J_1[t]}, w) = 1$. In particular, this means that

$$p_t \geq \Pr \left[C(x_{S[t]}, w) = 1 : \begin{array}{l} (1^n, 1^{s_c}, 1^{s_p}, P, S, \text{st}_{\mathcal{A}}) \leftarrow \mathcal{A}_1(1^\lambda) \\ (\text{crs}, \text{td}) \leftarrow \text{TrapGenIndex}(1^\lambda, 1^n, 1^{s_c}, 1^{s_p}, 1^k, S, S[t]) \\ (C, (x_1, \dots, x_k), \pi) \leftarrow \mathcal{A}_2(\text{crs}, \text{st}_{\mathcal{A}}) \\ (t', w) \leftarrow \text{Extract}(\text{td}, C, P, (x_1, \dots, x_k), \pi) \end{array} \right], \quad (8.2)$$

where we write $S[t]$ to denote the t^{th} value in S in ascending order. We now show that the difference between the outputs of $\widetilde{\text{Hyb}}_{0,t,3}$ and $\widetilde{\text{Hyb}}_{0,t,4}$ is exactly p_t :

Claim 8.11. *It holds that $\Pr[\widetilde{\text{Hyb}}_{0,t,3}(\mathcal{A}) = 1] = \Pr[\widetilde{\text{Hyb}}_{0,t,4}(\mathcal{A}) = 1] + p_t$.*

Proof. By construction, the only difference between $\widetilde{\text{Hyb}}_{0,t,3}$ and $\widetilde{\text{Hyb}}_{0,t,4}$ is the *additional* check in $\text{Hyb}_{0,t,4}$ that the extracted bit b satisfies $b = 0$. Let E_0 be the event that $b = 0$ in the execution of $\widetilde{\text{Hyb}}_{0,t,3}$ and E_1 be the event that $b = 1$. Then,

$$\begin{aligned} \Pr[\widetilde{\text{Hyb}}_{0,t,4}(\mathcal{A}) = 1] &= \Pr[\widetilde{\text{Hyb}}_{0,t,3}(\mathcal{A}) = 1 \wedge E_0] \\ &= \Pr[\widetilde{\text{Hyb}}_{0,t,3}(\mathcal{A}) = 1] - \Pr[\widetilde{\text{Hyb}}_{0,t,3}(\mathcal{A}) = 1 \wedge E_1] \\ &= \Pr[\widetilde{\text{Hyb}}_{0,t,3}(\mathcal{A}) = 1] - p_t. \end{aligned}$$

The claim follows. \square

Claim 8.12. If Π_H satisfies zero-fixing against efficient non-uniform adversaries, then there exists a negligible function $\text{negl}(\cdot)$ such that for every $i \in \{1, \dots, d-1\}$, $t \in [|J_{i+1}|]$, it holds that

$$\left| \Pr[\widetilde{\text{Hyb}}_{i,t,3}(\mathcal{A}) = 1] - \Pr[\widetilde{\text{Hyb}}_{i,t,4}(\mathcal{A}) = 1] \right| \leq \text{negl}(\lambda).$$

Proof. Follows by a similar argument as the proof of [Claim 4.11](#) (for the case where $i > 0$). \square

Claim 8.13. If Π_H satisfies one-sided index hiding with extracted guess security against efficient non-uniform adversaries, then there exists a negligible function $\text{negl}(\cdot)$ such that for every $i \in \{0, \dots, d-1\}$, $t \in [|J_{i+1}|]$, it holds that

$$\Pr[\widetilde{\text{Hyb}}_{i,t,4}(\mathcal{A}) = 1] \leq \Pr[\widetilde{\text{Hyb}}_{i,t,5}(\mathcal{A}) = 1] + \text{negl}(\lambda).$$

Proof. Follows by a similar argument as the proof of [Claim 4.12](#). \square

Claim 8.14. For every $i \in \{0, \dots, d-1\}$, $t \in [|J_{i+1}|]$, it holds that

$$\Pr[\widetilde{\text{Hyb}}_{i,t,6}(\mathcal{A}) = 1] \geq \Pr[\widetilde{\text{Hyb}}_{i,t,5}(\mathcal{A}) = 1].$$

Proof. Follows by a similar argument as the proof of [Claim 4.13](#). \square

Claim 8.15. If Π_{BARG} satisfies set hiding against efficient non-uniform adversaries, then there exists a negligible function $\text{negl}(\cdot)$ such that for every $i \in \{0, \dots, d-1\}$ and $t \in [|J_{i+1}|]$, it holds that

$$\left| \Pr[\widetilde{\text{Hyb}}_{i,t,6}(\mathcal{A}) = 1] - \Pr[\widetilde{\text{Hyb}}_{i,t,7}(\mathcal{A}) = 1] \right| \leq \text{negl}(\lambda).$$

Proof. Follows by a similar argument as the proof of [Claim 4.14](#). \square

Claim 8.16. For every $i \in \{0, \dots, d-1\}$, it holds that

$$\Pr[\widetilde{\text{Hyb}}_{i,|J_{i+1}|,7}(\mathcal{A}) = 1] \leq \Pr[\widetilde{\text{Hyb}}_{i,\text{final}}(\mathcal{A}) = 1].$$

Proof. Follows by a similar argument as the proof of [Claim 4.15](#). \square

Claim 8.17. If Π_H satisfies set hiding property against efficient non-uniform adversaries then there exists a negligible function $\text{negl}(\cdot)$ such that for every $i \in \{0, \dots, d-1\}$, it holds that

$$\left| \Pr[\widetilde{\text{Hyb}}_{i,\text{final}}(\mathcal{A}) = 1] - \Pr[\widetilde{\text{Hyb}}_{i+1}(\mathcal{A}) = 1] \right| \leq \text{negl}(\lambda).$$

Proof. Follows by a similar argument as the proof of [Claim 4.16](#). \square

Claim 8.18. If \mathcal{A} is admissible, then $\Pr[\widetilde{\text{Hyb}}_{d-1,1,4}(\mathcal{A}) = 1] = 0$.

Proof. The proof is almost identical to that of [Claim 4.17](#), but relies on the fact that \mathcal{A} is an admissible adversary. Namely, if \mathcal{A} is admissible for the semi-somewhere extractability game, then it outputs a set S and a policy P such that $P(b_1, \dots, b_k) = 0$ where $b_i = 0$ if $i \in S$ and $b_i = 1$ otherwise. By construction, the challenger in $\widetilde{\text{Hyb}}_{d-1,1,4}(\mathcal{A}) = 1$ sets $\beta_i = b_i$ for all $i \in [k]$, and $\beta_{k+1}, \dots, \beta_s$ to be the wire values for $P(b_1, \dots, b_k)$. In particular, this means that $\beta_s = 0$, and therefore $J_d[1] = \{\beta_s\}$. In this case, $\widetilde{\text{Hyb}}_{d-1,1,4}$ cannot always output 1 since the following two conditions must simultaneously hold:

- On the one hand, there must exist a witness $\hat{w} = (b, \sigma^{(0)}, \sigma^{(1)}, w)$ for instance \hat{x}_s (of the relation in [Fig. 1](#)) where $b = 0$.
- On the other hand, by definition of instance \hat{x}_s , since s is the output wire, it must be that $b = 1$.

Therefore the output in $\widetilde{\text{Hyb}}_{d-1,1,4}$ is always 0. \square

Combining [Claims 8.8 to 8.18](#), we conclude that

$$\Pr[\widetilde{\text{Hyb}}_0(\mathcal{A}) = 1] \leq \sum_{t \in [|S|]} p_t + \text{negl}(\lambda). \quad (8.3)$$

To complete the proof, we relate the probability $\varepsilon_{\mathcal{A}}(\lambda, k)$ from [Eq. \(8.1\)](#) that \mathcal{A} outputs a valid proof in the semi-somewhere extractability game to the probability that $\widetilde{\text{Hyb}}_0(\mathcal{A})$ outputs 1. The only difference between $\widetilde{\text{Hyb}}_0$ and the semi-somewhere extractability game is the fact that in Hyb_0 , the common reference string is norm (output by Gen) whereas in the semi-somewhere extractability game, it is output by TrapGen. We give the formal reduction below:

Claim 8.19. *If [Construction 8.3](#) satisfies set hiding, then there exists a negligible function $\text{negl}(\cdot)$ such that*

$$|\Pr[\widetilde{\text{Hyb}}_0(\mathcal{A}) = 1] - \varepsilon_{\mathcal{A}}(\lambda, k)| \leq \text{negl}(\lambda)$$

where $\varepsilon_{\mathcal{A}}$ is the probability that \mathcal{A} outputs a verifying proof in the ExptSE experiment from [Definition 8.1](#).

Proof. Suppose $|\Pr[\widetilde{\text{Hyb}}_0(\mathcal{A}) = 1] - \varepsilon_{\mathcal{A}}(\lambda, k)| \geq \varepsilon$ for some non-negligible ε . We use \mathcal{A} to construct an adversary \mathcal{B} for the set hiding game:

1. On input the security parameter 1^λ , algorithm \mathcal{B} runs $(1^n, 1^{s_c}, 1^{s_p}, P, S, \text{st}_{\mathcal{A}}) \leftarrow \mathcal{A}_1(1^\lambda)$. Algorithm \mathcal{B} then forwards $1^n, 1^{s_c}, 1^{s_p}, 1^{k(\lambda)}$, and S to the challenger.
2. The challenger replies with crs . Algorithm \mathcal{B} runs $(C, (x_1, \dots, x_k), \pi) \leftarrow \mathcal{A}_2(\text{crs}, \text{st}_{\mathcal{A}})$. Algorithm \mathcal{B} then outputs $\text{Verify}(\text{crs}, C, P, (x_1, \dots, x_k), \pi)$.

If the challenger samples $\text{crs} \leftarrow \text{Gen}(1^\lambda, 1^n, 1^{s_c}, 1^{s_p})$, then \mathcal{B} perfectly simulates the distribution $\widetilde{\text{Hyb}}_0$ and outputs 1 with probability $\Pr[\widetilde{\text{Hyb}}_0(\mathcal{A}) = 1]$. If the challenger samples $\text{crs} \leftarrow \text{TrapGen}(1^\lambda, 1^n, 1^{s_c}, 1^{s_p}, 1^k, S)$, then \mathcal{B} perfectly simulates the distribution $\text{ExptSE}_{\mathcal{A}}(\lambda, k)$ and outputs 1 with probability $\varepsilon_{\mathcal{A}}(\lambda, k)$. We conclude that algorithm \mathcal{B} succeeds with the same advantage ε . \square

Combining [Eq. \(8.3\)](#) with [Claim 8.19](#), we have that

$$\varepsilon_{\mathcal{A}}(\lambda, k) \leq \sum_{t \in [|S|]} p_t + \text{negl}(\lambda).$$

Next, since TrapGen invokes TrapGenIndex on a random $t \in S$ and appealing to [Eq. \(8.2\)](#), we conclude that

$$\Pr[\text{ExptSE}_{\mathcal{A}}(\lambda) = 1] = \sum_{t \in [|S|]} \frac{p_t}{|S|} \geq \sum_{t \in [|S|]} \frac{p_t}{k} \geq \frac{1}{k} \cdot \varepsilon_{\mathcal{A}}(\lambda, k) - \text{negl}(\lambda). \quad \square$$

8.1 Monotone Policy Aggregate Signature via Semi-Somewhere Extractability

In [Section 7](#), we showed how to combine a non-adaptively-sound monotone BARG with a puncturable signature scheme to obtain a monotone policy aggregate signature scheme. In this section, we show that the same construction is also secure for *any* signature scheme (not necessarily puncturable) if we rely on semi-somewhere extractability instead. We first recall the notion of a standard (non-puncturable) signature scheme:

Definition 8.20 (Digital Signature). An digital signature scheme with message space $\{0, 1\}^\lambda$ is a tuple of efficient algorithms $\Pi_{\text{Sig}} = (\text{Gen}, \text{Sign}, \text{Verify})$ with the following syntax:

- $\text{Gen}(1^\lambda) \rightarrow (\text{vk}, \text{sk})$: On input the security parameter λ , the key-generation algorithm outputs a key pair (vk, sk) .
- $\text{Sign}(\text{sk}, m) \rightarrow \sigma$: On input a signing key sk and a message $m \in \{0, 1\}^\lambda$, the signing algorithm outputs a signature σ .

- $\text{Verify}(\text{vk}, m, \sigma) \rightarrow b$: On input a verification key vk , a message $m \in \{0, 1\}^\lambda$, and a signature σ , the verification algorithm outputs a bit $b \in \{0, 1\}$.

Moreover, the signature scheme should satisfy the following properties:

- **Correctness:** For all $\lambda \in \mathbb{N}$ and all $m \in \{0, 1\}^\lambda$, it holds that

$$\Pr \left[\text{Verify}(\text{vk}, m, \sigma) = 1 : \begin{array}{l} (\text{vk}, \text{sk}) \leftarrow \text{Gen}(1^\lambda) \\ \sigma \leftarrow \text{Sign}(\text{sk}, m) \end{array} \right] = 1.$$

- **Unforgeability:** For all efficient and admissible adversaries \mathcal{A} , there exists a negligible function $\text{negl}(\cdot)$ such that

$$\Pr \left[\text{Verify}(\text{vk}, m^*, \sigma^*) = 1 : \begin{array}{l} (\text{vk}, \text{sk}) \leftarrow \text{Gen}(1^\lambda) \\ (m^*, \sigma^*) \leftarrow \mathcal{A}^{\text{Sign}(\text{sk}, \cdot)}(1^\lambda, \text{vk}) \end{array} \right] = \text{negl}(\lambda),$$

where we say \mathcal{A} is admissible if it does not query the signing oracle $\text{Sign}(\text{sk}, \cdot)$ on the message m^* in the above security game.

Theorem 8.21 (Static Unforgeability). *Consider an instantiation of [Construction 4.4](#) where the puncturable signature scheme Π_{PunctSig} is replaced by a standard digital signature scheme $\Pi_{\text{Sig}} = (\text{Sig.Gen}, \text{Sig.Sign}, \text{Sig.Verify})$ ([Definition 8.20](#)). If $\Pi_{\text{MP-BARG}}$ is semi-somewhere extractable and satisfies set hiding, and Π_{Sig} is unforgeable, then [Construction 7.5](#) is statically unforgeable.*

Proof. Let \mathcal{A} be any efficient adversary for the static security game. We begin by defining a sequence of hybrid experiments:

- Hyb_0 : This is the static unforgeability experiment:
 1. On input the security parameter λ , the adversary \mathcal{A} outputs the number of parties 1^k , a number of verification keys 1^n , a bound on the policy size 1^{s_p} , a challenge message $m^* \in \{0, 1\}^\lambda$, and a monotone policy $P: \{0, 1\}^k \rightarrow \{0, 1\}$.
 2. The challenger samples key-pairs $(\text{vk}_i, \text{sk}_i) \leftarrow \text{Sig.Gen}(1^\lambda)$ for all $i \in [n]$ and sends $\text{vk}_1, \dots, \text{vk}_n$ to the adversary.
 3. The adversary \mathcal{A} can now issue signing queries. Each signing query consists of an index $i \in [n]$ and a message $m \in \{0, 1\}^\lambda \setminus \{m^*\}$. The challenger responds with $\sigma \leftarrow \text{Sig.Sign}(\text{sk}_i, m)$.
 4. After the adversary is finished making signing queries, it outputs a tuple of verification keys $(\text{vk}_1^*, \dots, \text{vk}_k^*)$.
 5. The challenger replies with the common reference string $\text{crs} \leftarrow \text{BARG.Gen}(1^\lambda, 1^{\ell_{\text{vk}}}, 1^{s_c}, 1^{s_p})$.
 6. The adversary \mathcal{A} can continue to make signing queries. The challenger responds to these exactly as before. When \mathcal{A} finishes making signing queries, the adversary outputs the aggregate signature σ_{agg}^* .
 7. The output of the experiment is 1 if all of the following holds:
 - For each $i \in [k]$, let $b_i = 0$ if $\text{vk}_i^* = \text{vk}_j$ for some $j \in [n]$. Otherwise, let $b_i = 1$. Then, it holds that $P(b_1, \dots, b_k) = 0$.
 - $\text{BARG.Verify}(\text{crs}, C_{m^*}, P, (\text{vk}_1^*, \dots, \text{vk}_k^*), \sigma_{\text{agg}}^*) = 1$.
Otherwise, the output is 0.
- Hyb_1 : Same as Hyb_0 , except the challenger uses the following modified procedure to sample the BARG common reference string:
 - Sample the common reference string $(\text{crs}, \text{td}) \leftarrow \text{BARG.TrapGen}(1^\lambda, 1^{\ell_{\text{vk}}}, 1^{s_c}, 1^{s_p}, 1^k, S)$, where $S = \{i \in [k] \mid \exists j \in [n] : \text{vk}_i^* = \text{vk}_j\}$.
- Hyb_2 : Same as Hyb_1 , except the experiment outputs 1 if all of the following holds:

- For each $i \in [k]$, let $b_i = 0$ if $\text{vk}_i^* = \text{vk}_j$ for some $j \in [n]$. Otherwise, let $b_i = 1$. Then, it holds that $P(b_1, \dots, b_k) = 0$.
- Extract the instance and witness $(i, \sigma_i) \leftarrow \text{BARG.Extract}(\text{td}_{\text{BARG}}, C_{m^*}, P, (\text{vk}_1^*, \dots, \text{vk}_k^*), \sigma_{\text{agg}}^*)$ and check that $C_{m^*}(\text{vk}_i^*, \sigma_i) = 1$ and $i \in S$.

If either check fails, then the challenger outputs 0. Notably, the experiment **no longer checks the condition** $\text{BARG.Verify}(\text{crs}, C_{m^*}, P, (\text{vk}_1^*, \dots, \text{vk}_k^*), \sigma_{\text{agg}}^*) = 1$.

- Hyb_3 : Same as Hyb_2 , except the challenger **samples a random index** $i^* \xleftarrow{\mathbb{R}} [n]$ at the beginning of the security game. After computing $(i, \sigma_i) \leftarrow \text{BARG.Extract}(\text{td}_{\text{BARG}}, C_{m^*}, P, (\text{vk}_1^*, \dots, \text{vk}_k^*), \sigma_{\text{agg}}^*)$, the challenger **additionally** checks that $\text{vk}_i^* = \text{vk}_{i^*}$. If the check fails, the challenger outputs 0.

Lemma 8.22. *If $\Pi_{\text{MP-BARG}}$ satisfies set hiding, then there exists a negligible function $\text{negl}(\cdot)$ such that*

$$|\Pr[\text{Hyb}_0(\mathcal{A}) = 1] - \Pr[\text{Hyb}_1(\mathcal{A}) = 1]| = \text{negl}(\lambda).$$

Proof. Suppose $|\Pr[\text{Hyb}_0(\mathcal{A}) = 1] - \Pr[\text{Hyb}_1(\mathcal{A}) = 1]| \geq \varepsilon$ for some non-negligible ε . We use \mathcal{A} to construct an adversary \mathcal{B} for the set hiding game as follows:

1. On input a security parameter 1^λ , compute $(1^k, 1^n, 1^{s_p}, m^*, P) \leftarrow \mathcal{A}(1^\lambda)$.
2. Algorithm \mathcal{B} samples key-pairs $(\text{vk}_i, \text{sk}_i) \leftarrow \text{Sig.Gen}(1^\lambda)$ for all $i \in [n]$ and send $\text{vk}_1, \dots, \text{vk}_n$ to \mathcal{A} .
3. Whenever algorithm \mathcal{A} makes a signing query on an index $i \in [n]$ and a message $m \in \{0, 1\}^\lambda \setminus \{m^*\}$, algorithm \mathcal{B} responds with $\sigma \leftarrow \text{Sig.Sign}(\text{sk}_i, m)$.
4. After \mathcal{A} is finished making signing queries, it outputs a tuple of verification keys $(\text{vk}_1^*, \dots, \text{vk}_k^*)$.
5. Let $S = \{i \in [k] \mid \exists j \in [n] : \text{vk}_i^* = \text{vk}_j\}$. Algorithm \mathcal{B} sends the tuple $(1^{\ell_{\text{vk}}}, 1^{s_c}, 1^{s_p}, 1^k, S)$ to the $\Pi_{\text{MP-BARG}}$ challenger. The challenger replies with a common reference string crs , which algorithm \mathcal{B} forwards to algorithm \mathcal{A} .
6. Algorithm \mathcal{B} responds to additional signing queries exactly as before.
7. Upon receiving an aggregate signature σ_{agg}^* from \mathcal{A} , algorithm \mathcal{B} outputs 1 if all of the following holds:
 - (a) $P(b_1, \dots, b_k) = 0$ where $b_i = 0$ if $i \in S$ and $b_i = 1$ otherwise.
 - (b) $\text{BARG.Verify}(\text{crs}, C_{m^*}, P, (\text{vk}_1^*, \dots, \text{vk}_k^*), \sigma_{\text{agg}}^*) = 1$.

Otherwise, algorithm \mathcal{B} outputs 0.

If the set hiding challenger samples $\text{crs} \leftarrow \text{BARG.Gen}(1^\lambda, 1^{\ell_{\text{vk}}}, 1^{s_c}, 1^{s_p})$, then \mathcal{B} perfectly simulates Hyb_0 for \mathcal{A} and outputs 1 with probability $\Pr[\text{Hyb}_0(\mathcal{A}) = 1]$. On the other hand, if the set hiding challenger samples $(\text{crs}, \text{td}) \leftarrow \text{BARG.TrapGen}(1^\lambda, 1^{\ell_{\text{vk}}}, 1^{s_c}, 1^{s_p}, 1^k, S)$, then \mathcal{B} perfectly simulates an execution of Hyb_1 for \mathcal{A} and outputs 1 with probability $\Pr[\text{Hyb}_1(\mathcal{A}) = 1]$. We conclude that algorithm \mathcal{B} breaks set hiding with the same advantage ε . \square

Lemma 8.23. *If $\Pi_{\text{MP-BARG}}$ is semi-somewhere extractable, then there exists a negligible function such that*

$$\Pr[\text{Hyb}_2(\mathcal{A}) = 1] \geq \frac{1}{k} \cdot \Pr[\text{Hyb}_1(\mathcal{A}) = 1] - \text{negl}(\lambda).$$

Proof. Let $k = k(\lambda)$ be the number of parties that algorithm \mathcal{A} outputs.¹¹ We use \mathcal{A} to construct an adversary \mathcal{B} for the semi-somewhere extractability game (with parameter $k = k(\lambda)$) as follows:

1. On input a security parameter 1^λ , algorithm \mathcal{B} computes $(1^k, 1^n, 1^{s_p}, m^*, P) \leftarrow \mathcal{A}(1^\lambda)$.

¹¹We can assume that for each value of $\lambda \in \mathbb{N}$, algorithm \mathcal{A} always picks a *fixed* value of k . This can be the value that maximizes its success probability for each value of λ (formally, we can take this “maximizing” value of k to be non-uniform advice provided to \mathcal{A}).

2. Algorithm \mathcal{B} samples key-pairs $(vk_i, sk_i) \leftarrow \text{Sig.Gen}(1^\lambda)$ for all $i \in [n]$ and send vk_1, \dots, vk_n to \mathcal{A} .
3. Whenever algorithm \mathcal{A} makes a signing query on an index $i \in [n]$ and a message $m \in \{0, 1\}^\lambda \setminus \{m^*\}$, algorithm \mathcal{B} responds with $\sigma \leftarrow \text{Sig.Sign}(sk_i, m)$.
4. After \mathcal{A} is finished making signing queries, it outputs a tuple of verification keys (vk_1^*, \dots, vk_k^*) .
5. Let $S = \{i \in [k] \mid \exists j \in [n] : vk_i^* = vk_j^*\}$. Let $b_i = 0$ if $i \in S$ and $b_i = 1$ otherwise. If $P(b_1, \dots, b_k) = 0$, then algorithm \mathcal{B} aborts with output \perp . Otherwise, algorithm \mathcal{B} sends the tuple $(1^{\ell_{vk}}, 1^{s_c}, 1^{s_p}, P, S)$ to the challenger and receives a common reference string crs . Algorithm \mathcal{B} forwards crs to \mathcal{A} .
6. Algorithm \mathcal{B} responds to additional signing queries exactly as before.
7. Upon receiving an aggregate signature σ_{agg}^* from \mathcal{A} , algorithm \mathcal{B} outputs the circuit C_{m^*} , the instances (vk_1^*, \dots, vk_k^*) and the proof σ_{agg}^* .

Algorithm \mathcal{B} is admissible by construction (since $P(b_1, \dots, b_k) = 0$ where $b_i = 0$ if $i \in S$ and $b_i = 1$ otherwise). Next, algorithm \mathcal{B} perfectly simulates the view of \mathcal{A} in the hybrids Hyb_1 and Hyb_2 . By assumption, with probability $\Pr[\text{Hyb}_1(\mathcal{A}) = 1]$, algorithm \mathcal{A} outputs σ_{agg} where

$$\text{BARG.Verify}(\text{crs}, C_{m^*}, P, (vk_1^*, \dots, vk_k^*), \sigma_{\text{agg}}^*) = 1.$$

This means that

$$\varepsilon_{\mathcal{B}}(\lambda, k) = \Pr[\text{BARG.Verify}(\text{crs}, C_{m^*}, P, (vk_1^*, \dots, vk_k^*), \sigma_{\text{agg}}^*) = 1] = \Pr[\text{Hyb}_1(\mathcal{A}) = 1], \quad (8.4)$$

where $\varepsilon_{\mathcal{B}}(\lambda, k)$ is the quantity from Eq. (8.1). By somewhere extractability of $\Pi_{\text{MP-BARG}}$, there exists a negligible function $\text{negl}(\cdot)$ such that

$$\Pr[\text{ExptSE}_{\mathcal{B}}(\lambda, k) = 1] \geq \frac{1}{k} \cdot \varepsilon_{\mathcal{B}}(\lambda, k) - \text{negl}(\lambda), \quad (8.5)$$

Next, the output of $\text{ExptSE}_{\mathcal{B}}(\lambda, k)$ is 1 if $C_{m^*}(vk_i^*, \sigma_i) = 1$ and $i \in S$ where

$$(i, \sigma_i) \leftarrow \text{BARG.Extract}(\text{td}_{\text{BARG}}, C_{m^*}, P, (vk_1^*, \dots, vk_k^*), \sigma_{\text{agg}}^*).$$

This is the same set of conditions checked in Hyb_2 and we conclude that

$$\Pr[\text{ExptSE}_{\mathcal{B}}(\lambda, k) = 1] = \Pr[\text{Hyb}_2(\mathcal{A}) = 1]. \quad (8.6)$$

Combining Eqs. (8.4) to (8.6), we have that

$$\Pr[\text{Hyb}_2(\mathcal{A}) = 1] \geq \frac{1}{k} \cdot \Pr[\text{Hyb}_1(\mathcal{A}) = 1] - \text{negl}(\lambda). \quad \square$$

Lemma 8.24. $\Pr[\text{Hyb}_3(\mathcal{A}) = 1] = \frac{1}{n} \cdot \Pr[\text{Hyb}_2(\mathcal{A}) = 1]$.

Proof. By construction, the adversary's view in Hyb_2 and Hyb_3 is identical. The only difference is how the output of the experiment is computed. Suppose $\Pr[\text{Hyb}_2(\mathcal{A}) = 1] = \varepsilon$. Then, with probability ε , the extracted value (i, σ_i) in Hyb_2 and Hyb_3 satisfies $i \in S$. By definition of S , there exists some $j \in [n]$ such that $vk_i^* = vk_j^*$. Since $i^* \xleftarrow{\mathcal{R}} [n]$ (and is entirely independent of the view of the adversary), $\Pr[j = i^*] = 1/n$. When $j = i^*$, the output in Hyb_3 is also 1 (and otherwise, it is 0). As such, $\Pr[\text{Hyb}_3(\mathcal{A}) = 1] = \frac{1}{n} \Pr[\text{Hyb}_2(\mathcal{A}) = 1]$ and the claim follows. \square

Lemma 8.25. If Π_{Sig} is unforgeable, then there exists a negligible function $\text{negl}(\cdot)$ such that $\Pr[\text{Hyb}_3(\mathcal{A}) = 1] = \text{negl}(\lambda)$.

Proof. Suppose $\Pr[\text{Hyb}_3(\mathcal{A}) = 1] > \varepsilon$ for some non-negligible ε . We use \mathcal{A} to construct an efficient adversary \mathcal{B} that breaks unforgeability of Π_{Sig} :

1. At the beginning of the game, algorithm \mathcal{B} receives a verification key vk from the challenger.

2. On input a security parameter 1^λ , algorithm \mathcal{B} computes $(1^k, 1^n, 1^{s_p}, m^*, P) \leftarrow \mathcal{A}(1^\lambda)$.
3. Algorithm \mathcal{B} samples a random index $i^* \xleftarrow{R} [n]$. For $i \in [n] \setminus \{i^*\}$, algorithm \mathcal{B} samples $(vk_i, sk_i) \leftarrow \text{Sig.Gen}(1^\lambda)$. It sets $vk_{i^*} := vk$ and send vk_1, \dots, vk_n to \mathcal{A} .
4. Whenever \mathcal{A} makes a signing query on an index $i \in [n]$ and a message $m \in \{0, 1\}^\lambda \setminus \{m^*\}$, if $i = i^*$, then algorithm \mathcal{B} forwards the query to the challenger and receives a signature σ . Algorithm \mathcal{B} replies to \mathcal{A} with σ . If $i \neq i^*$, then algorithm \mathcal{B} replies with $\sigma \leftarrow \text{Sign}(sk_i, m)$.
5. At some point, algorithm \mathcal{A} outputs a collection of verification keys (vk_1^*, \dots, vk_k^*) .
6. Let $S = \{i \in [k] \mid \exists j \in [n] : vk_i^* = vk_j\}$. For each $i \in S$, set $b_i = 0$ and set $b_i = 1$ otherwise. If $P(b_1, \dots, b_k) = 1$, then algorithm \mathcal{B} aborts. Otherwise, algorithm \mathcal{B} samples the common reference string $(\text{crs}, \text{td}) \leftarrow \text{BARG.TrapGen}(1^\lambda, 1^{\ell_{\text{vk}}}, 1^{s_c}, 1^{s_p}, 1^k, S)$ and send crs to \mathcal{A} .
7. Algorithm \mathcal{B} responds to additional signing queries exactly as before.
8. Upon receiving an aggregate signature σ_{agg}^* from \mathcal{A} , algorithm \mathcal{B} computes

$$(i, \sigma_i) \leftarrow \text{BARG.Extract}(\text{td}_{\text{BARG}}, C_{m^*}, P, (vk_1^*, \dots, vk_k^*), \sigma_{\text{agg}}^*).$$

If $vk_i^* \neq vk_{i^*} = vk$ then algorithm \mathcal{B} aborts. Otherwise, algorithm output $\sigma^* = \sigma_i$ and m^* .

Algorithm \mathcal{B} is admissible since it never needs to query its challenger for a signature on m^* . Next, algorithm \mathcal{B} perfectly simulates an execution of Hyb_3 for \mathcal{A} . Thus, with probability ε , algorithm \mathcal{A} outputs a tuple of verification keys (vk_1^*, \dots, vk_k^*) and a signature σ_{agg} such that the extracted index-signature pair (i, σ_i) satisfies $vk_i^* = vk_{i^*} = vk$ and $C_{m^*}(vk_i^*, \sigma_i) = 1$. By definition of C_{m^*} , this means that

$$1 = \text{Verify}(vk_i^*, m^*, \sigma_i) = \text{Verify}(vk, m^*, \sigma_i),$$

in which case \mathcal{B} wins the unforgeability game. □

[Theorem 8.21](#) now follows immediately from [Lemmas 8.22](#) to [8.25](#). □

Acknowledgments

We thank Yuval Ishai for helpful pointers on batch arguments. Brent Waters is supported by NSF CNS-1908611, CNS-2318701, and a Simons Investigator award. David J. Wu is supported by NSF CNS-2151131, CNS-2140975, CNS-2318701, a Microsoft Research Faculty Fellowship, and a Google Research Scholar award.

References

- [BBK⁺23] Zvika Brakerski, Maya Farber Brodsky, Yael Tauman Kalai, Alex Lombardi, and Omer Paneth. SNARGs for monotone policy batch NP. In *CRYPTO*, pages 252–283, 2023.
- [BCJP24] Maya Farber Brodsky, Arka Rai Choudhuri, Abhishek Jain, and Omer Paneth. Monotone-policy aggregate signatures. In *EUROCRYPT*, 2024.
- [BGN05] Dan Boneh, Eu-Jin Goh, and Kobbi Nissim. Evaluating 2-DNF formulas on ciphertexts. In *TCC*, pages 325–341, 2005.
- [BKP22] Shany Ben-David, Yael Tauman Kalai, and Omer Paneth. Verifiable private information retrieval. In *TCC*, pages 3–32, 2022.
- [BWY11] Mihir Bellare, Brent Waters, and Scott Yilek. Identity-based encryption secure against selective opening attack. In *TCC*, pages 235–252, 2011.

- [CF13] Dario Catalano and Dario Fiore. Vector commitments and their applications. In *PKC*, pages 55–72, 2013.
- [CGJ⁺23] Arka Rai Choudhuri, Sanjam Garg, Abhishek Jain, Zhengzhong Jin, and Jiaheng Zhang. Correlation intractability and snargs from sub-exponential DDH. In *CRYPTO*, pages 635–668, 2023.
- [CJJ21a] Arka Rai Choudhuri, Abhishek Jain, and Zhengzhong Jin. Non-interactive batch arguments for NP from standard assumptions. In *CRYPTO*, pages 394–423, 2021.
- [CJJ21b] Arka Rai Choudhuri, Abhishek Jain, and Zhengzhong Jin. SNARGs for P from LWE. In *FOCS*, pages 68–79, 2021.
- [DGKV22] Lalita Devadas, Rishab Goyal, Yael Kalai, and Vinod Vaikuntanathan. Rate-1 non-interactive arguments for batch-NP and applications. In *FOCS*, pages 1057–1068, 2022.
- [FWW23] Cody Freitag, Brent Waters, and David J. Wu. How to use (plain) witness encryption: Registered abe, flexible broadcast, and more. In *CRYPTO*, pages 498–531, 2023.
- [Gam84] Taher El Gamal. A public key cryptosystem and a signature scheme based on discrete logarithms. In *CRYPTO*, pages 10–18, 1984.
- [GVW19] Rishab Goyal, Satyanarayana Vusirikala, and Brent Waters. Collusion resistant broadcast and trace from positional witness encryption. In *PKC*, pages 3–33, 2019.
- [GZ21] Alonso González and Alexandros Zacharakis. Succinct publicly verifiable computation. *IACR Cryptol. ePrint Arch.*, page 353, 2021.
- [HJKS22] James Hulett, Ruta Jawale, Dakshita Khurana, and Akshayaram Srinivasan. SNARGs for P from sub-exponential DDH and QR. In *EUROCRYPT*, pages 520–549, 2022.
- [HW15] Pavel Hubáček and Daniel Wichs. On the communication complexity of secure function evaluation with long output. In *ITCS*, pages 163–172, 2015.
- [KLVW23] Yael Kalai, Alex Lombardi, Vinod Vaikuntanathan, and Daniel Wichs. Boosting batch arguments and RAM delegation. In *STOC*, pages 1545–1552, 2023.
- [KPY19] Yael Tauman Kalai, Omer Paneth, and Lisa Yang. How to delegate computations publicly. In *STOC*, pages 1115–1124, 2019.
- [KVZ21] Yael Tauman Kalai, Vinod Vaikuntanathan, and Rachel Yun Zhang. Somewhere statistical soundness, post-quantum security, and SNARGs. In *TCC*, pages 330–368, 2021.
- [Mer87] Ralph C. Merkle. A digital signature based on a conventional encryption function. In *CRYPTO*, pages 369–378, 1987.
- [NY90] Moni Naor and Moti Yung. Public-key cryptosystems provably secure against chosen ciphertext attacks. In *STOC*, pages 427–437, 1990.
- [Pai99] Pascal Paillier. Public-key cryptosystems based on composite degree residuosity classes. In *EUROCRYPT*, pages 223–238, 1999.
- [PP22] Omer Paneth and Rafael Pass. Incrementally verifiable computation via rate-1 batch arguments. In *FOCS*, pages 1045–1056, 2022.
- [PR17] Omer Paneth and Guy N. Rothblum. On zero-testable homomorphic encryption and publicly verifiable non-interactive arguments. In *TCC*, pages 283–315, 2017.
- [Reg05] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In *STOC*, pages 84–93, 2005.

- [Vad06] Salil P. Vadhan. An unconditional study of computational zero knowledge. *SIAM J. Comput.*, 36(4):1160–1214, 2006.
- [WW22] Brent Waters and David J. Wu. Batch arguments for NP and more from standard bilinear group assumptions. In *CRYPTO*, pages 433–463, 2022.

A Set Hiding with Extraction for BARGs

We now show to construct a BARG satisfying [Definition 2.8](#) from any somewhere extractable BARG that supports extraction on a single instance (e.g., [CJJ21b, WW22, HJKS22, DGKV22, KLVW23]). Our construction is a direct parallel of the analogous constructions from [GZ21, CJJ21b] in the setting of somewhere extractable commitments.

Construction A.1 (BARGs Satisfying Set Hiding with Extraction). Let $\Pi'_{\text{BARG}} = (\text{Gen}', \text{Prove}', \text{Verify}', \text{TrapGen}', \text{Extract}')$ be a somewhere-extractable BARG for Boolean circuit satisfiability that supports extraction on a single instance. We use Π'_{BARG} to construct a new BARG $\Pi_{\text{BARG}} = (\text{Gen}, \text{Prove}, \text{Verify}, \text{TrapGen}, \text{Extract})$ that supports extraction on multiple instances and which satisfies [Definition 2.8](#):

- $\text{Gen}(1^\lambda, 1^k, 1^n, 1^s, 1^\ell)$: On input the security parameter λ , the number of instances k , the instance length n , the bound on the size of the Boolean circuit s , and the bound on the size of the extraction set ℓ , the generator algorithm samples $(\text{crs}'_i, \text{vk}'_i) \leftarrow \text{Gen}'(1^\lambda, 1^k, 1^n, 1^s, 1^1)$ for each $i \in [\ell]$. Then, it samples a random permutation $\tau: [\ell] \rightarrow [\ell]$ and outputs $\text{crs} = (\text{crs}'_{\tau(1)}, \dots, \text{crs}'_{\tau(\ell)})$ and $\text{vk} = (\text{vk}'_{\tau(1)}, \dots, \text{vk}'_{\tau(\ell)})$.
- $\text{Prove}(\text{crs}, C, (x_1, \dots, x_k), (w_1, \dots, w_k))$: On input the common reference string $\text{crs} = (\text{crs}'_1, \dots, \text{crs}'_\ell)$, a Boolean circuit $C: \{0, 1\}^n \times \{0, 1\}^h \rightarrow \{0, 1\}$, statements $x_1, \dots, x_k \in \{0, 1\}^n$, and witnesses $w_1, \dots, w_k \in \{0, 1\}^h$, the prove algorithm computes $\pi'_i \leftarrow \text{Prove}'(\text{crs}'_i, C, (x_1, \dots, x_k), (w_1, \dots, w_k))$ for all $i \in [k]$ and outputs $\pi = (\pi'_1, \dots, \pi'_\ell)$.
- $\text{Verify}(\text{vk}, C, (x_1, \dots, x_k), \pi)$: On input the verification key $\text{vk} = (\text{vk}'_1, \dots, \text{vk}'_\ell)$, a Boolean circuit $C: \{0, 1\}^n \times \{0, 1\}^h \rightarrow \{0, 1\}$, statements $x_1, \dots, x_k \in \{0, 1\}^n$, and a proof $\pi = (\pi'_1, \dots, \pi'_\ell)$, the verification algorithm outputs 1 if for all $i \in [\ell]$, it holds that $\text{Verify}'(\text{vk}'_i, C, (x_1, \dots, x_k), \pi'_i) = 1$. Otherwise, it outputs 0.
- $\text{TrapGen}(1^\lambda, 1^k, 1^n, 1^s, 1^\ell, S)$: On input the security parameter λ , the number of instances k , the instance size n , the bound on the size of the Boolean circuit s , the bound on the size of the extraction set ℓ , and a set of indices $S \subseteq [k]$ of size at most ℓ , the trapdoor-generator algorithm proceeds as follows:
 - Let $S = \{j_1, \dots, j_d\}$ where $j_1 < j_2 < \dots < j_d$ are in sorted order.
 - For each $i \in [d]$, sample $(\text{crs}'_i, \text{vk}'_i, \text{td}'_i) \leftarrow \text{TrapGen}'(1^\lambda, 1^k, 1^n, 1^s, 1^1, \{j_i\})$. For each $i \in \{d+1, \dots, \ell\}$, sample $\text{crs}'_i \leftarrow \text{Gen}'(1^\lambda, 1^k, 1^n, 1^s, 1^1)$.
 - Sample a random permutation $\tau: [\ell] \rightarrow [\ell]$, and define the dictionary D where $D[j_i] \mapsto (\tau^{-1}(i), \text{td}'_i)$ for all $i \in S$. Output $\text{crs} = (\text{crs}'_{\tau(1)}, \dots, \text{crs}'_{\tau(\ell)})$, $\text{vk} = (\text{vk}'_{\tau(1)}, \dots, \text{vk}'_{\tau(\ell)})$, and $\text{td} = D$.
- $\text{Extract}(\text{td}, C, (x_1, \dots, x_k), \pi, i)$: On input the trapdoor $\text{td} = D$, a Boolean circuit $C: \{0, 1\}^n \times \{0, 1\}^h \rightarrow \{0, 1\}$, statements $x_1, \dots, x_k \in \{0, 1\}^n$, a proof $\pi = (\pi'_1, \dots, \pi'_\ell)$, and an index i , the extraction algorithm outputs \perp if $i \notin D$. Otherwise, let $(t, \text{td}') \leftarrow D[i]$ and output $\text{Extract}'(\text{td}', C, (x_1, \dots, x_k), \pi'_t, i)$.

Theorem A.2 (Completeness). *If Π'_{BARG} satisfies completeness, then Π_{BARG} in [Construction A.1](#) is also complete.*

Proof. This follows by construction. Specifically, take any $\lambda, k, n, s, \ell \in \mathbb{N}$, any Boolean circuit $C: \{0, 1\}^n \times \{0, 1\}^h \rightarrow \{0, 1\}$ of size at most s , any sequence of statements $x_1, \dots, x_k \in \{0, 1\}^n$ along with witnesses $w_1, \dots, w_k \in \{0, 1\}^h$ where $C(x_i, w_i) = 1$ for all $i \in [k]$. Then, the following properties hold:

- Suppose $(\text{crs}, \text{vk}) \leftarrow \text{Gen}(1^\lambda, 1^k, 1^n, 1^s, 1^\ell)$. By construction, this means $\text{crs} = (\text{crs}'_1, \dots, \text{crs}'_\ell)$ and $\text{vk} = (\text{vk}'_1, \dots, \text{vk}'_\ell)$. Moreover, for all $i \in [\ell]$, we have that $(\text{crs}'_i, \text{vk}'_i) \leftarrow \text{Gen}'(1^\lambda, 1^k, 1^n, 1^s, 1^1)$.

- Let $\pi \leftarrow \text{Prove}(\text{crs}, C, (x_1, \dots, x_k), (w_1, \dots, w_k))$. By construction, $\pi = (\pi'_1, \dots, \pi'_\ell)$ where for all $i \in [\ell]$, we have that $\pi'_i \leftarrow \text{Prove}'(\text{crs}'_i, C, (x_1, \dots, x_k), (w_1, \dots, w_k))$.
- By completeness of Π'_{BARG} , $\text{Verify}'(\text{vk}'_i, C, (x_1, \dots, x_k), \pi'_i) = 1$ for all $i \in [\ell]$. Thus $\text{Verify}(\text{vk}, C, (x_1, \dots, x_k), \pi)$ outputs 1 and completeness holds. \square

Theorem A.3 (Set Hiding). *If Π'_{BARG} satisfies set hiding, then Π_{BARG} in [Construction A.1](#) also satisfies set hiding.*

Proof. We start by defining a sequence of hybrid experiments.

- Hyb_0 : This is experiment $\text{ExptSH}_{\mathcal{A}}(\lambda, 0)$. Namely, after the adversary chooses the parameters $k, n, s, \ell \in \mathbb{N}$ and the set $S \subseteq [k]$, the challenger replies with $\text{crs} = (\text{crs}'_{\tau(1)}, \dots, \text{crs}'_{\tau(\ell)})$ and $\text{vk} = (\text{vk}'_{\tau(1)}, \dots, \text{vk}'_{\tau(\ell)})$ where $(\text{crs}'_i, \text{vk}'_i) \leftarrow \text{Gen}'(1^\lambda, 1^k, 1^n, 1^s, 1^1)$ for all $i \in [\ell]$ and τ is a random permutation. At the end of the experiment, the adversary outputs a bit $b' \in \{0, 1\}$ which is the output of the experiment.
- Hyb_i for $i \in [d]$: Same as Hyb_0 , except for indices $t \leq i$, the challenger computes $(\text{crs}'_t, \text{vk}'_t, \text{td}'_t) \leftarrow \text{TrapGen}'(1^\lambda, 1^k, 1^n, 1^s, 1^1, \{j_t\})$, where $S = \{j_1, \dots, j_d\}$ and $j_1 < \dots < j_d$.

For an adversary \mathcal{A} , we write $\text{Hyb}_i(\mathcal{A})$ to denote the output of an execution of hybrid Hyb_i with adversary \mathcal{A} . By construction, $\text{Hyb}_0(\mathcal{A}) \equiv \text{ExptSH}_{\mathcal{A}}(\lambda, 0)$ and $\text{Hyb}_d(\mathcal{A}) \equiv \text{ExptSH}_{\mathcal{A}}(\lambda, 1)$. We now show that if Π'_{BARG} satisfies set hiding, then for all $i \in [d]$, and for all efficient adversaries \mathcal{A} , the output distributions $\text{Hyb}_{i-1}(\mathcal{A})$ and $\text{Hyb}_i(\mathcal{A})$ are computationally indistinguishable. To see this, suppose there exists an efficient adversary \mathcal{A} such that

$$|\Pr[\text{Hyb}_{i-1}(\mathcal{A}) = 1] - \Pr[\text{Hyb}_i(\mathcal{A}) = 1]| \geq \varepsilon(\lambda),$$

for some non-negligible ε . We use \mathcal{A} to construct an efficient algorithm \mathcal{B} that breaks set hiding of Π'_{BARG} :

1. On input the security parameter 1^λ , algorithm \mathcal{B} starts by running adversary \mathcal{A} on the same security parameter. Algorithm \mathcal{A} outputs $1^k, 1^n, 1^s, 1^\ell$, and a set $S = \{j_1, \dots, j_d\}$, where $j_1 < \dots < j_d$.
2. Algorithm \mathcal{B} sends $1^k, 1^n, 1^s, 1^\ell$, and $\{j_i\}$ to the set hiding challenger for Π'_{BARG} and receives a pair $(\text{crs}^*, \text{vk}^*)$. It sets $\text{crs}'_i = \text{crs}^*$ and $\text{vk}'_i = \text{vk}^*$.
3. For $t < i$, algorithm \mathcal{B} samples $(\text{crs}'_t, \text{vk}'_t, \text{td}'_t) \leftarrow \text{TrapGen}'(1^\lambda, 1^k, 1^n, 1^s, 1^1, \{j_t\})$. For $t \in [i+1, \ell]$, it samples $(\text{crs}'_t, \text{vk}'_t) \leftarrow \text{Gen}'(1^\lambda, 1^k, 1^n, 1^s, 1^1)$.
4. Finally, algorithm \mathcal{B} samples a random permutation $\tau: [\ell] \rightarrow [\ell]$ and gives $\text{crs} = (\text{crs}'_{\tau(1)}, \dots, \text{crs}'_{\tau(\ell)})$ and $\text{vk} = (\text{vk}'_{\tau(1)}, \dots, \text{vk}'_{\tau(\ell)})$ to \mathcal{A} . It outputs whatever \mathcal{A} outputs.

By design, if $(\text{crs}^*, \text{vk}^*) \leftarrow \text{Gen}'(1^\lambda, 1^k, 1^n, 1^s, 1^1)$, then algorithm \mathcal{B} perfectly simulates Hyb_{i-1} for \mathcal{A} and outputs 1 with probability $\Pr[\text{Hyb}_{i-1}(\mathcal{A}) = 1]$. If $(\text{crs}^*, \text{vk}^*, \text{td}^*) \leftarrow \text{TrapGen}'(1^\lambda, 1^k, 1^n, 1^s, 1^1, \{j_i\})$, then algorithm \mathcal{B} perfectly simulates Hyb_i for \mathcal{A} and outputs 1 with probability $\Pr[\text{Hyb}_i(\mathcal{A}) = 1]$. We conclude that algorithm \mathcal{B} breaks set hiding of Π'_{BARG} with the same non-negligible advantage ε . The claim now follows by a hybrid argument. \square

Theorem A.4 (Set Hiding with Extraction). *If Π'_{BARG} satisfies set hiding, then Π_{BARG} in [Construction A.1](#) satisfies set hiding with extraction.*

Proof. We begin by defining a sequence of hybrid experiments:

- Hyb_0 : This is experiment $\text{ExptSHwE}_{\mathcal{A}}(\lambda, 0)$:
 - At the beginning of the game, the adversary chooses the parameters $k, n, s, \ell \in \mathbb{N}$, the set $S \subseteq [k]$, and the index $i^* \in S$. Let $S = \{j_1, \dots, j_d\}$ where $j_1 < \dots < j_d$. Let $t^* \in [d]$ be the index where $i^* = j_{t^*}$.
 - Then, for each $i \in [d]$, the challenger samples $(\text{crs}'_i, \text{vk}'_i, \text{td}'_i) \leftarrow \text{TrapGen}'(1^\lambda, 1^k, 1^n, 1^s, 1^1, \{j_i\})$. For each $i \in \{d+1, \dots, \ell\}$, sample $(\text{crs}'_i, \text{vk}'_i) \leftarrow \text{Gen}'(1^\lambda, 1^k, 1^n, 1^s, 1^1)$. Then, it samples a random permutation $\tau: [\ell] \rightarrow [\ell]$. It gives $\text{crs} = (\text{crs}'_{\tau(1)}, \dots, \text{crs}'_{\tau(\ell)})$ and $\text{vk} = (\text{vk}'_{\tau(1)}, \dots, \text{vk}'_{\tau(\ell)})$ to \mathcal{A} . The challenger also sets $\text{td}' = \text{td}'_{t^*}$ and $z = \tau^{-1}(t^*)$.

- Algorithm \mathcal{A} outputs a Boolean circuit $C: \{0, 1\}^n \times \{0, 1\}^h \rightarrow \{0, 1\}$, statements $x_1, \dots, x_k \in \{0, 1\}^n$, and a proof $\pi = (\pi'_1, \dots, \pi'_\ell)$.
 - The challenger checks that for all $i \in [\ell]$, it holds that $\text{Verify}'(\text{vk}_{\tau(i)}, C, (x_1, \dots, x_k), \pi'_i) = 1$. If not, the challenger halts with output 0. Otherwise, the challenger replies with $\text{Extract}'(\text{td}', C, (x_1, \dots, x_k), \pi'_2, i^*)$.
 - Algorithm \mathcal{A} outputs a bit $b' \in \{0, 1\}$, which is the output of the experiment.
- Hyb_1 : Same as Hyb_0 , except the challenger swaps $(\text{crs}'_1, \text{vk}'_1)$ with $(\text{crs}'_{t^*}, \text{vk}'_{t^*})$. In more detail, the challenger samples $(\text{crs}'_1, \text{vk}'_1, \text{td}'_1) \leftarrow \text{TrapGen}'(1^\lambda, 1^k, 1^n, 1^s, 1^1, \{j_{t^*}\})$ and $(\text{crs}'_{t^*}, \text{vk}'_{t^*}, \text{td}'_{t^*}) \leftarrow \text{TrapGen}'(1^\lambda, 1^k, 1^n, 1^s, 1^1, \{j_1\})$. In addition, it sets $\text{td}' = \text{td}'_1$ and $z = \tau^{-1}(1)$. The remainder of the experiment proceeds as in Hyb_0 .
 - Hyb_i for $i \in \{2, \dots, d\}$: Same as Hyb_1 except for indices $t \in \{2, \dots, i\}$, the challenger now computes $(\text{crs}'_t, \text{vk}'_t) \leftarrow \text{Gen}'(1^\lambda, 1^k, 1^n, 1^s, 1^1)$.

For an adversary \mathcal{A} , we write $\text{Hyb}_i(\mathcal{A})$ to denote the output distribution of an execution of hybrid Hyb_i with adversary \mathcal{A} . By construction $\text{Hyb}_0(\mathcal{A}) \equiv \text{ExptSHwE}_{\mathcal{A}}(\lambda, 0)$ while $\text{Hyb}_d(\mathcal{A}) \equiv \text{ExptSHwE}_{\mathcal{A}}(\lambda, 1)$. To complete the proof, we now show that the output of each adjacent pair of hybrid experiments are indistinguishable.

Lemma A.5. *For all adversaries \mathcal{A} , we have that $\text{Hyb}_0(\mathcal{A}) \equiv \text{Hyb}_1(\mathcal{A})$.*

Proof. The view of the adversary in the two experiments is identical since τ is a random permutation. More precisely, the distribution in Hyb_1 corresponds to the distribution in Hyb_0 where the permutation τ is replaced by $\tau \circ \sigma$ where $\sigma: [\ell] \rightarrow [\ell]$ is the elementary permutation that transposes j_1 with j_{t^*} (and fixes all other inputs). Since τ is uniform, the distributions of τ and $\tau \circ \sigma$ are identical. \square

Lemma A.6. *If Π'_{BARG} satisfies set hiding, then for all $i \in \{2, \dots, d\}$ and all efficient adversaries \mathcal{A} , it holds that $|\Pr[\text{Hyb}_{i-1}(\mathcal{A}) = 1] - \Pr[\text{Hyb}_i(\mathcal{A}) = 1]| = \text{negl}(\lambda)$.*

Proof. Suppose there exists an efficient adversary \mathcal{A} where $|\Pr[\text{Hyb}_{i-1}(\mathcal{A}) = 1] - \Pr[\text{Hyb}_i(\mathcal{A}) = 1]| \geq \varepsilon(\lambda)$ for some non-negligible ε . We use \mathcal{A} to construct an efficient adversary \mathcal{B} that breaks set hiding of Π'_{BARG} :

1. On input the security parameter 1^λ , algorithm \mathcal{B} starts running adversary \mathcal{A} on the same security parameter. Algorithm \mathcal{A} outputs $1^k, 1^n, 1^s, 1^\ell$, a set $S = \{j_1, \dots, j_d\}$ where $j_1 < \dots < j_d$, and an index $i^* \in S$. Let $t^* \in [d]$ be the index where $i^* = j_{t^*}$. Let $\rho_1 = j_{t^*}$, $\rho_{t^*} = j_1$, and $\rho_i = j_i$ for all $i \notin \{1, t^*\}$.
2. Algorithm \mathcal{B} sends $1^k, 1^n, 1^s, 1^1$, and $\{\rho_i\}$ to the set hiding challenger for Π'_{BARG} and receives a common reference string crs^* and verification key vk^* .
3. Algorithm \mathcal{B} samples $(\text{crs}'_1, \text{vk}'_1, \text{td}'_1) \leftarrow \text{TrapGen}'(1^\lambda, 1^k, 1^n, 1^1, \{\rho_1\})$. For $2 \leq t < i$, algorithm \mathcal{B} samples $(\text{crs}'_t, \text{vk}'_t) \leftarrow \text{Gen}'(1^\lambda, 1^k, 1^n, 1^s, 1^1)$. Finally, for all indices $t \in \{i+1, \dots, d\}$, it samples $(\text{crs}'_t, \text{vk}'_t, \text{td}'_t) \leftarrow \text{TrapGen}'(1^\lambda, 1^k, 1^n, 1^s, 1^1, \{\rho_t\})$. For each $t \in \{d+1, \dots, \ell\}$, it samples $(\text{crs}'_t, \text{vk}'_t) \leftarrow \text{Gen}'(1^\lambda, 1^k, 1^n, 1^s, 1^1)$. It sets $\text{crs}'_i = \text{crs}^*$ and $\text{vk}'_i = \text{vk}^*$.
4. Algorithm \mathcal{B} samples a random permutation $\tau: [\ell] \rightarrow [\ell]$ and gives $\text{crs} = (\text{crs}'_{\tau(1)}, \dots, \text{crs}'_{\tau(\ell)})$ and $\text{vk} = (\text{vk}'_{\tau(1)}, \dots, \text{vk}'_{\tau(\ell)})$ to \mathcal{A} . It also sets $\text{td}' = \text{td}'_1$ and $z = \tau^{-1}(1)$.
5. Algorithm \mathcal{A} outputs a Boolean circuit $C: \{0, 1\}^n \times \{0, 1\}^h \rightarrow \{0, 1\}$, statements $x_1, \dots, x_k \in \{0, 1\}^n$, and a proof $\pi = (\pi'_1, \dots, \pi'_\ell)$.
6. Algorithm \mathcal{B} first checks that for all $i \in [\ell]$, it holds that $\text{Verify}'(\text{vk}_{\tau(i)}, C, (x_1, \dots, x_k), \pi'_i) = 1$. If not, it halts with output 0. Otherwise, algorithm \mathcal{B} replies with $\text{Extract}'(\text{td}', C, (x_1, \dots, x_k), \pi'_2, i^*)$.
7. Algorithm \mathcal{A} outputs a bit $b' \in \{0, 1\}$, which \mathcal{B} also outputs.

By design, if $(\text{crs}^*, \text{vk}^*, \text{td}^*) \leftarrow \text{TrapGen}'(1^\lambda, 1^k, 1^n, 1^s, 1^1, \{\rho_i\})$, then algorithm \mathcal{B} perfectly simulates Hyb_{i-1} for \mathcal{A} , so algorithm \mathcal{B} outputs 1 with probability $\Pr[\text{Hyb}_{i-1}(\mathcal{A}) = 1]$. Conversely, if $(\text{crs}^*, \text{vk}^*) \leftarrow \text{Gen}'(1^\lambda, 1^k, 1^n, 1^s, 1^1)$, then algorithm \mathcal{B} perfectly simulates Hyb_i for \mathcal{A} and algorithm \mathcal{B} outputs 1 with probability $\Pr[\text{Hyb}_i(\mathcal{A}) = 1]$. Thus, the distinguishing advantage of algorithm \mathcal{B} is at least ε , which is non-negligible by assumption. \square

Security now follows by combining [Lemmas A.5](#) and [A.6](#) and appealing to the fact that $d \leq \ell = \text{poly}(\lambda)$. \square

Theorem A.7 (Somewhere Extraction). *If Π'_{BARG} is somewhere extractable in trapdoor mode, then Π_{BARG} in [Construction A.1](#) is also somewhere extractable in trapdoor mode.*

Proof. Suppose there exists an efficient adversary \mathcal{A} that breaks the somewhere extractability of [Construction A.1](#) with non-negligible probability $\varepsilon(\lambda)$. We use \mathcal{A} to construct an efficient adversary \mathcal{B} that breaks the somewhere extractability of Π_{BARG} :

1. On input the security parameter 1^λ , algorithm \mathcal{B} starts by running algorithm \mathcal{A} on the same parameter. Algorithm \mathcal{A} outputs $1^k, 1^n, 1^s, 1^\ell$, and a set $S = \{j_1, \dots, j_d\}$, where $j_1 < \dots < j_d$. Algorithm \mathcal{B} samples a random index $t^* \xleftarrow{R} [d]$ and sends $1^k, 1^n, 1^s, 1^1$, and $\{j_{t^*}\}$ to its challenger. It receives a common reference string crs^* and a verification key vk^* .
2. For $t \in [d] \setminus \{t^*\}$, algorithm \mathcal{B} samples $(\text{crs}'_t, \text{vk}'_t, \text{td}'_t) \leftarrow \text{TrapGen}'(1^\lambda, 1^k, 1^n, 1^s, 1^1, \{j_t\})$. For each $i \in \{d+1, \dots, \ell\}$, it samples $(\text{crs}'_i, \text{vk}'_i) \leftarrow \text{Gen}'(1^\lambda, 1^k, 1^n, 1^s, 1^1)$. Finally, it sets $\text{crs}'_{t^*} = \text{crs}^*$.
3. Finally, algorithm \mathcal{B} samples a random permutation $\tau: [\ell] \rightarrow [\ell]$ and gives $\text{crs} = (\text{crs}'_{\tau(1)}, \dots, \text{crs}'_{\tau(\ell)})$ and $\text{vk} = (\text{vk}'_{\tau(1)}, \dots, \text{vk}'_{\tau(\ell)})$ to \mathcal{A} .
4. Algorithm \mathcal{A} outputs a Boolean circuit $C: \{0, 1\}^n \times \{0, 1\}^h \rightarrow \{0, 1\}$, statements $x_1, \dots, x_k \in \{0, 1\}^n$ and a proof $\pi = (\pi'_1, \dots, \pi'_\ell)$. Algorithm \mathcal{B} outputs the circuit C , the statement $x_{j_{t^*}}$, and the proof $\pi'_{\tau^{-1}(t^*)}$.

First, algorithm \mathcal{B} perfectly simulates the common reference string crs for \mathcal{A} , so with probability at least ε , algorithm \mathcal{A} outputs $(C, x_1, \dots, x_k, \pi)$ such that there exists some $t \in [d]$ such that the following two conditions hold:

- $\text{Verify}'(\text{vk}'_{\tau(z)}, C, (x_1, \dots, x_k), \pi'_z) = 1$ where $z = \tau^{-1}(t) \in [\ell]$; and
- $C(x_{j_t}, w_{j_t}) \neq 1$ where $w_{j_t} \leftarrow \text{Extract}'(\text{td}'_t, C, (x_1, \dots, x_k), \pi'_z, j_t)$.

Moreover, the special index t^* is perfectly hidden from the view of \mathcal{A} , so with probability $1/|S| \geq 1/\ell$, it will be the case that $t^* = t$. In this case, $\tau(z) = \tau(\tau^{-1}(t)) = t$, so we have that π'_z verifies with respect to $\text{vk}'_{\tau(z)} = \text{vk}'_{t^*}$, but the extracted witness $w_{j_{t^*}}$ is such that $C(x_{j_{t^*}}, w_{j_{t^*}}) \neq 1$. This breaks somewhere extractability of Π_{BARG} . Thus, if \mathcal{A} succeeds with advantage ε , then algorithm \mathcal{B} succeeds with advantage at least ε/ℓ , which is non-negligible as $\ell = \text{poly}(\lambda)$. \square

Theorem A.8 (Succinctness). *If Π'_{BARG} is succinct, then Π_{BARG} in [Construction A.1](#) is also succinct.*

Proof. Take any $\lambda, k, n, s, \ell \in \mathbb{N}$ and any (crs, vk) in the support of $\text{Gen}(1^\lambda, 1^k, 1^n, 1^s, 1^\ell)$. Then, $\text{crs} = (\text{crs}'_1, \dots, \text{crs}'_\ell)$ and $\text{vk} = (\text{vk}'_1, \dots, \text{vk}'_\ell)$, where $(\text{crs}'_i, \text{vk}'_i)$ is in the support of $\text{Gen}'(1^\lambda, 1^k, 1^n, 1^s, 1^1)$. Take any Boolean circuit $C: \{0, 1\}^n \times \{0, 1\}^h \rightarrow \{0, 1\}$. We consider each condition separately:

- **Succinct proof:** By succinctness of Π_{BARG} , the proofs π' output by $\text{Prove}'(\text{crs}', C, \cdot, \cdot)$ satisfy $|\pi'_i| \leq \text{poly}(\lambda + \log k + s)$. Then, the proofs output by $\text{Prove}(\text{crs}, C, \cdot, \cdot)$ satisfy $|\pi| \leq \ell \cdot |\pi'_i| \leq \text{poly}(\lambda + \log k + s + \ell)$.
- **Succinct CRS:** By succinctness of Π'_{BARG} , each crs'_i satisfies $|\text{crs}'_i| \leq \text{poly}(\lambda + k + n) + \text{poly}(\lambda + \log m + s)$. The total size of the CRS is a factor of ℓ larger which satisfies the succinctness requirement.
- **Succinct verification key:** By succinctness of Π'_{BARG} , each vk'_i satisfies $|\text{vk}'_i| \leq \text{poly}(\lambda + \log k + s)$. The verification key vk output by Setup is a factor ℓ larger, which satisfies the succinctness requirement. \square

Remark A.9 (Index BARGs). While we described [Construction A.1](#) for the case of standard BARGs, the same construction directly extends to the case of index BARGs, and moreover, the construction preserves the efficiency requirements of an index BARG (since it is simply a concatenation of ℓ copies of the underlying BARG).