# Non-Interactive Zero-Knowledge from Non-Interactive Batch Arguments 

Jeffrey Champion<br>UT Austin<br>jchampion@utexas.edu

David J. Wu<br>UT Austin<br>dwu4@cs.utexas.edu


#### Abstract

Zero-knowledge and succinctness are two important properties that arise in the study of non-interactive arguments. Previously, Kitagawa et al. (TCC 2020) showed how to obtain a non-interactive zero-knowledge (NIZK) argument for NP from a succinct non-interactive argument (SNARG) for NP. In particular, their work demonstrates how to leverage the succinctness property from an argument system and transform it into a zero-knowledge property.

In this work, we study a similar question of leveraging succinctness for zero-knowledge. Our starting point is a batch argument for NP, a primitive that allows a prover to convince a verifier of $T$ NP statements $x_{1}, \ldots, x_{T}$ with a proof whose size scales sublinearly with $T$. Unlike SNARGs for NP, batch arguments for NP can be built from group-based assumptions in both pairing and pairing-free groups and from lattice-based assumptions. The challenge with batch arguments is that the proof size is only amortized over the number of instances, but can still encode full information about the witness to a small number of instances.

We show how to combine a batch argument for NP with a local pseudorandom generator (i.e., a pseudorandom generator where each output bit only depends on a small number of input bits) and a dual-mode commitment scheme to obtain a NIZK for NP. Our work provides a new generic approach of realizing zero-knowledge from succinctness and highlights a new connection between succinctness and zero-knowledge.


## 1 Introduction

In a non-interactive argument system for an NP language $\mathcal{L}$, a prover sends a single message $\pi$ to try and convince an efficient verifier that an NP statement $x$ is true (i.e., $x \in \mathcal{L}$ ). In an argument system [BCC88], we require soundness to hold against computationally-bounded provers (i.e., a computationally-bounded prover should not be able to convince the verifier of a false statement). Two of the most important properties considered in the context of cryptographic argument systems are zero-knowledge and succinctness:

- Zero-knowledge: We say a non-interactive argument satisfies zero-knowledge [GMR85] if the proof $\pi$ for an NP statement $x$ reveals nothing more about $x$ other than the fact that the statement is true. We refer to such arguments as non-interactive zero-knowledge (NIZK) arguments [BFM88]. While NIZKs for NP are unlikely to exist in the plain model (unless NP $\subseteq$ BPP), a long line of works have constructed NIZKs in the common reference string (CRS) model from a broad range of algebraic assumptions including factoring [FLS90], assumptions on pairing groups [CHK03, GOS06, GOS12, LPWW20] and pairing-free groups [JJ21], lattice-based assumptions [ $\mathrm{CCH}^{+} 19$, PS19], and combinations of multiple assumptions [BKM20].
- Succinctness: A second property of interest is the length of the proof (and by correspondence, the verification complexity). We say that an argument system for $\mathcal{L}$ is succinct if the length of the proof and the running time of the verifier is sublinear (or more commonly, polylogarithmic) in the size of the NP relation associated with $\mathcal{L}$. Such arguments are referred to as succinct non-interactive arguments, or SNARGs [GW11]. We refer to [Mic95, Gro10, BCCT12, DFH12, Lip13, PHGR13, GGPR13, BCI ${ }^{+}$13, BCPR14, Gro16, BISW17, BCC ${ }^{+} 17$, BISW18, BBHR18, COS20, $\mathrm{CHM}^{+} 20$, Set20, $\mathrm{ACL}^{+} 22, \mathrm{BS} 23, \mathrm{CBBZ} 23$ ] and the references therein for a survey of succinct arguments.

Many applications require non-interactive arguments that are simultaneously succinct and zero-knowledge (i.e., zkSNARGs). Given a SNARG and a NIZK (argument of knowledge), it is straightforward to obtain a zkSNARG by direct composition (i.e., the zkSNARG is a NIZK argument of knowledge of a SNARG proof of the statement). A natural question is whether there is a formal connection between these two fundamental properties of cryptographic arguments. Previously Kitagawa et al. [KMY20] showed that SNARGs for NP and one-way functions together imply NIZKs for NP (and by composition, a zkSNARG). ${ }^{1}$ The intuition underlying this result is that since the length of a SNARG proof for an NP statement $x$ is much shorter than the length of the associated witness $w$, the SNARG proof simply cannot reveal too many bits of the witness information-theoretically. Then, by composing the SNARG with leakage-resilient cryptography, this intuition can be leveraged to obtain a NIZK.

Batch arguments and zero knowledge. In this work, we continue this line of inquiry of studying the relationship between succinctness and zero-knowledge. Instead of focusing on SNARGs for NP, which have a very strong succinctness property and necessitates constructions in either idealized models or based on non-falsifiable assumptions [GW11], we start with the weaker notion of SNARGs for batch NP languages. This is a notion that has received extensive study recently [KVZ21, CJJ21a, CJJ21b, HJKS22, WW22, DGKV22, GSWW22, CGJ+ 22, KLVW23] and can be realized from standard cryptographic assumptions. At a high level, in a non-interactive batch argument (BARG), a prover can convince a verifier of a collection of $T$ NP statements $\left(x_{1}, \ldots, x_{T}\right)$ with a proof of size poly $(\lambda, s) \cdot o(T)$, where $s$ is the size of the circuit computing the NP relation. In particular, a batch argument amortizes the cost of NP verification across multiple instances. While this amortization still confers some succinctness, it is certainly possible for a BARG proof to leak one or more of the underlying witnesses associated with the statements used to construct it. Thus, we ask the question:

## Can we construct a NIZK argument for NP from a non-interactive batch argument for NP?

Our results. In this work, we give a generic construction of NIZKs for NP from a non-interactive batch argument for NP in conjunction with a dual-mode commitment scheme and a (sub-exponentially-hard) low-locality pseudorandom generator (PRG) with super-linear stretch. Dual-mode commitment schemes can be built from any lossy/dual-mode public-key encryption scheme, which is known from most standard assumptions [PW08, HLOV11, AFMP20]. A low-locality PRG is a PRG where each output bit only depends on a small number of the seed bits. PRGs with constant locality and super-linear stretch are a notable ingredient in the recent constructions of indistinguishability obfuscation from well-studied assumptions [JLS21, JLS22]. Our instantiations can rely on locality as high as $c \log \lambda$, where $c<1$ is a constant and $\lambda$ is the seed length; this is a much weaker requirement compared to the constant locality PRGs required for indistinguishability obfuscation. The local PRG can in turn be instantiated using Goldreich's family of PRGs [Gol00, CM01].

While the additional ingredients we rely on for constructing a NIZK for NP are (much) stronger than one-way functions, we emphasize that no combination of the underlying primitives by themselves are known to imply NIZKs for NP. We summarize our main result with the following theorem (see Corollary 3.15 for a formal description and parameter specification):
Theorem 1.1 (Informal). Let $G:\{0,1\}^{\lambda} \rightarrow\{0,1\}^{\lambda^{\delta}}$ be a $P R G$ with locality $c \log \lambda$ and super-linear stretch $\delta>1$ for a (sufficiently-small) constant $c<1$ and (sufficiently-large) $\delta>1$. Then, assuming the existence of a dual-mode commitment scheme, a non-interactive batch argument for NP with (sufficiently-small) proof size, ${ }^{2}$ and sub-exponential hardness of $G$, there exists a NIZK for NP.

Our work highlights a new connection between the succinctness of an argument system and zero-knowledge. It also provides a new generic approach for constructing NIZKs for NP. Finally, by composing a batch argument satisfying certain efficiency properties (satisfied by most existing constructions [CJJ21b, WW22, CGJ ${ }^{+} 22$ ]) with a NIZK for NP, we obtain a zero-knowledge batch argument. Thus, our approach can also be used to generically upgrade a

[^0]non-interactive batch argument for NP into a zero-knowledge batch argument by relying only on a low-locality PRG and a dual-mode commitment (neither of which are known to imply NIZKs for NP).

Concurrent work. In a concurrent and independent work, Bitansky et al. [ $\left.\mathrm{BKP}^{+} 23 \mathrm{a}, \mathrm{BK} \mathrm{P}^{+} 23 \mathrm{~b}\right]$ studied the relationship between batch arguments (both interactive and non-interactive) and statistical witness indistinguishability. Notably, their work also shows how to leverage batch arguments to construct statistical NIZK arguments (using a different set of techniques relative to this work). In the initial version of their work [BKP ${ }^{+}$23a], they showed how to construct a statistical NIZK argument for NP with a non-uniform prover from non-interactive batch arguments and lossy public-key encryption. This result is technically incomparable to ours. On the one hand, the prover in our construction is uniform while on the other, we require an additional assumption of sub-exponentially hard local PRGs; moreover, we achieve computational zero knowledge rather than statistical zero knowledge. Subsequent to the initial posting of this work and [ $\mathrm{BKP}^{+} 23 \mathrm{a}$ ], the authors have improved their construction [ $\mathrm{BKP}^{+} 23 \mathrm{~b}$ ] to obtain statistical NIZK arguments with a uniform prover by only relying on non-interactive batch arguments and lossy public-key encryption. This is a strict strengthening over Theorem 1.1. An interesting open question that remains is whether NIZK arguments can be built from just non-interactive batch arguments and vanilla public-key encryption.

### 1.1 Technical Overview

Our starting point is the generic approach of Kitagawa et al. [KMY20] who show how to generically transform a SNARG for NP into a NIZK for NP. The approach of [KMY20] instantiates the hidden-bits paradigm of Feige et al. [FLS90] of combining a NIZK in the idealized hidden-bits model with a hidden-bits generator (HBG) [QRW19].

NIZKs in the hidden-bits model. The hidden-bits model is an idealized model for constructing non-interactive zero-knowledge proofs. In this model, a trusted party first generates a string of uniformly random bits $r_{1}, \ldots, r_{m}{ }_{\leftarrow}^{\mathrm{R}}$ $\{0,1\}$ and gives them to the prover. To construct a proof for a statement $x$, the prover selects a subset of indices $I \subseteq[m]$ along with a proof $\pi$. The verifier then receives $\left\{r_{i}\right\}_{i \in I}$ and $\pi$ from the trusted party. The model ensures that the prover cannot influence the choice of bits $r_{1}, \ldots, r_{m}$ and that the verifier cannot learn the value of any unrevealed bit $r_{i}$ for $i \notin I$. Feige et al. [FLS90] previously showed how to construct a NIZK with statistical soundness and perfect zero-knowledge in the hidden-bits model for the NP-complete problem of graph Hamiltonicity.

Hidden-bits generators. Given a NIZK in the idealized hidden-bits model, a number of works have shown how to transform it into a NIZK in the CRS model through a cryptographic compiler [FLS90, BY92, CHK03, GR13, CL18, QRW19, LPWW20, KMY20]. In this work, we focus on the abstraction based on hidden-bits generators introduced by Quach et al. [QRW19]. At a high-level, a hidden-bits generator is a cryptographic primitive that generates a (pseudorandom) sequence of hidden bits. The prover can then open up a subset of the bits while ensuring the unopened bits remain hidden. Moreover, the hidden-bits generator ensures that the prover has limited control over the output sequence of bits. In a sense, hidden-bits generators provide a cryptographic realization of the trusted sampling of the hidden-bits string in the hidden-bits model. Thus, combined with the (unconditional) NIZK for NP in the hidden-bits model, a hidden-bits generator immediately implies a NIZK for NP in the CRS model. We now describe the syntax of a hidden-bits generator more formally; we specifically consider the adaptation from [KMY20]:

- The setup algorithm Setup takes as input the security parameter $\lambda$ and an output length $m$ and outputs a common reference string crs.
- The generator algorithm GenBits takes the common reference string crs and outputs a bit-string $\mathbf{r} \in\{0,1\}^{m}$ of length $m$ along with a generator state st. Here, $\mathbf{r}$ is the "hidden-bits string."
- The prove algorithm Prove takes the generator state st and a subset of indices $I \subseteq[m]$, and outputs a succinct proof $\pi$. The proof $\pi$ is an "opening" to the bits of $\mathbf{r}$ indexed by $I$; we denote these bits by $\mathbf{r}_{I} \in\{0,1\}^{|I|}$.
- The verification algorithm Verify takes as input the common reference string crs, a set of indices $I \subseteq[m]$, a collection of bits $\mathbf{r}_{I} \in\{0,1\}^{|I|}$, and an opening $\pi$. The verification algorithms decides whether $\pi$ is a valid opening or not to the bits indexed by $I$ (with respect to crs).

The hidden-bits generator must in turn satisfy the following properties:

- Correctness: Correctness says that if crs $\leftarrow \operatorname{Setup}\left(1^{\lambda}, 1^{m}\right)$ and ( $\left.\mathbf{r}, \mathrm{st}\right) \leftarrow \operatorname{GenBits}(\mathrm{crs})$, then for all sets of indices $I \subseteq[m]$, the opening $\pi$ output by Prove(st, $I$ ) is valid with respect to Verify.
- Binding: The binding property restricts the set of possible openings that can be computed by a computationallybounded algorithm. Namely, for each crs in the support of Setup, there exists a subset $\mathcal{V}^{\text {crs }} \subset\{0,1\}^{m}$ of "valid" hidden-bits strings. Namely, no efficient adversary can come up with an accepting proof $\pi$ for a set of indices $I \subseteq[m]$ and an assignment $\mathbf{r}_{I} \in\{0,1\}^{|I|}$ that is inconsistent with every $\mathbf{r}^{\prime} \in \mathcal{V}^{\text {crs }}$ (i.e., an assignment $\mathbf{r}_{I}$ such that for all $\mathbf{r}^{\prime} \in \mathcal{V}^{\text {crs }}, \mathbf{r}_{I}^{\prime} \neq \mathbf{r}_{I}$ ). Moreover, the set of possible hidden-bits strings induced by a particular CRS must be sparse: $\left|\mathcal{V}^{\text {crs }}\right| \leq 2^{m^{\gamma} \operatorname{poly}(\lambda)}$ for some constant $\gamma<1$ and where $\lambda$ is a security parameter.
- Hiding: The hiding property says that the unopened bits of $\mathbf{r}$ are pseudorandom. Namely, for any set $I \subseteq[m]$ and honestly-generated $\mathbf{r}$ and $\pi$, the distribution ( $\operatorname{crs}, I, \mathbf{r}_{I}, \mathbf{r}_{I}, \pi$ ) is computationally indistinguishable from the distribution (crs, $I, \mathbf{r}_{I}, \hat{\mathbf{r}}_{\bar{I}}, \pi$ ) where $\hat{\mathbf{r}} \stackrel{\mathrm{R}}{\leftarrow}\{0,1\}^{m}$ and $\bar{I}=[m] \backslash I$.

Kitagawa et al. [KMY20] show how to construct a hidden-bits generator satisfying the above properties by combining a SNARG for NP with a leakage-resilient (weak) pseudorandom function (PRF):

- The CRS contains $m$ random points in the domain of the PRF $x_{1}, \ldots, x_{m}$ and the CRS for the SNARG.
- The hidden-bits string is constructed by sampling a PRF key $k$ and setting $r_{i} \leftarrow \operatorname{PRF}\left(k, x_{i}\right)$ for each $i \in[m]$.
- The opening for a subset $I \subseteq[m]$ is a SNARG proof that there exists $k$ such that for all $i \in I, r_{i}=\operatorname{PRF}\left(k, x_{i}\right)$.

In this case, binding follows from security of the (weak) PRF (as long as the length of the PRF key is smaller than the output length $m$ ) in conjunction with soundness of the SNARG (i.e., the only possible openings are to those consistent with an evaluation under a PRF key $k$ on the inputs $x_{1}, \ldots, x_{m}$ ). The hiding property follows by treating the SNARG proof in the opening as "leakage" on the PRF key and then appealing to leakage-resilient pseudorandomness of the underlying PRF. Critically, this latter step relies on the length of the SNARG being sublinear in the length of the PRF key.

Replacing the SNARG with a batch argument. We first observe that the SNARG proof in the opening is almost a batch language. Namely, the proof is showing that for each index $i \in I$, the bit $r_{i}$ satisfies $r_{i}=\operatorname{PRF}\left(k, x_{i}\right)$. Each instance is described by a tuple ( $i, x_{i}, r_{i}$ ) and the witness is the PRF key $k$. The caveat is that in a batch language, there is no requirement that the prover uses the same witness (i.e., the PRF key $k$ ) for each instance. Namely, if we use replace the SNARG in the [KMY20] construction with a BARG, then the proof only suffices to argue "local consistency" (i.e., there exists some key $k_{i}$ that explains each output bit $r_{i}$ ) rather than "global consistency" (i.e., there exists a single key $k$ that explains each output bit $r_{i}$ ). Certainly, local consistency is insufficient as it is trivial to find a tuple of keys $\left(k_{1}, \ldots, k_{m}\right)$ that explains any candidate hidden-bits string $\mathbf{r} \in\{0,1\}^{m}$.

Enforcing consistency. To force the prover to use a consistent PRF key $k$ across all of the instances when constructing the batch argument, we have the prover include a commitment $c$ to the PRF key $k$ as part of the opening. Each instance of the batch NP language is now

$$
\exists k: c \text { is a commitment to } k \text { and } \operatorname{PRF}\left(k, x_{i}\right)=r_{i} .
$$

In fact, we note that we can replace the PRF with a pseudorandom generator PRG: $\{0,1\}^{\lambda} \rightarrow\{0,1\}^{m}$, and indeed, the (weak) PRF in the [KMY20] construction is essentially used as a PRG. We will write $\operatorname{PRG}_{i}:\{0,1\}^{\lambda} \rightarrow\{0,1\}$ to denote the function that takes as input the seed $s \in\{0,1\}^{\lambda}$ and outputs the $i^{\text {th }}$ bit of $\operatorname{PRG}(s)$. To generate the hidden-bits string, the generator now samples $s \stackrel{R}{\leftarrow}\{0,1\}^{\lambda}$ and commits to $s$ with a commitment $c$. The hidden-bits string is $\mathbf{r} \leftarrow \mathrm{PRG}(\mathrm{s})$ and the opening to $\mathbf{r}_{I}$ is a batch argument $\pi$ for the following language:

$$
\forall i \in I, \exists \mathbf{s} \in\{0,1\}^{\lambda}: c \text { is a commitment to } \mathbf{s} \text { and } \mathrm{PRG}_{i}(\mathbf{s})=r_{i} .
$$

As long as the commitment is statistically binding (i.e., the commitment $c$ can be opened to at most one seed s) and the batch argument is computationally sound, the scheme satisfies the binding requirement. In our security analysis (Theorem 3.3), we technically require a stronger extractability property on the commitment, which allows us to base binding on semi-adaptive soundness of the underlying BARG; this is the notion achieved by most recent constructions from standard assumptions [CJJ21b, WW22, HJKS22, DGKV22, HJKS22]. ${ }^{3}$ In contrast, the construction of [KMY20] relied on a SNARG with adaptive soundness. This is a stronger requirement that cannot be proven under a black-box reduction to a falsifiable assumption [GW11]. However, this approach for constructing a hidden bits generator does not satisfy hiding. There are two issues:

- Length of the commitment: The opening now contains a commitment $c$ to the PRG seed s. Since $c$ is statistically binding, the length of $c$ is at least as long as the seed $s$.
- Length of the proof: Succinctness of the batch argument says that the length of the proof $\pi$ satisfies $|\pi|=$ $\operatorname{poly}(\lambda,|C|, \log |I|)$, where $C$ is the circuit that takes as input $\left(i, \mathrm{~s}, r_{i}\right)$ and checks that $c$ is a commitment to $\mathbf{s}$ and $\operatorname{PRG}_{i}(\mathbf{s})=r_{i}$. Unlike the case of a SNARG, the length of $\pi$ scales polynomially with the size of the circuit $|C|$. Since $C$ takes the PRG seed as input, $|C| \geq|\mathbf{s}|$, so the length of $\pi$ is at least as long as the seed s.

The [KMY20] construction argues hiding by relying on leakage resilience of the underlying weak PRF. In their setting, the only leakage on the PRF key is from the SNARG, whose length is smaller than the length of the PRF key. As such, the analysis reduces to a standard leakage-resilience argument. In our setting, both the commitment to the PRG seed and the length of the BARG proof potentially leak too much information about the PRG seed, and we cannot directly leverage leakage resilience to argue hiding.

Leveraging locality. Our first observation is that each individual instance in the batch language is checking a single output bit of the PRG. Since the length of the BARG proof scales with the size of the circuit checking a single instance, this means that if the circuit for validating a single output bit of the PRG is much smaller than the length of the overall PRG seed, we can rely on BARG succinctness. One way to construct PRGs with this property is by relying on locality. We say that a PRG is $k$-local if each output bit only depends on at most $k$ bits of the seed. If a PRG is $k$-local, then each output bit can be verified with a circuit of size at most $2^{k} \cdot p o l y(\lambda)$. In this case, to check that output bit $i$ is correctly computed, the relation only needs to check (local) openings for the $k$ bits of $s$ that determine $\operatorname{PRG}_{i}(\mathbf{s})$. For instance, if the PRG has constant locality [Gol00, CM01], and we take the commitment $c$ to be a bit-by-bit commitment to the bits of $s$, then verifying a single output bit only requires a circuit of size $\lambda^{\delta}$, for some fixed constant $\delta>0$ that depends on the BARG scheme and the commitment scheme (but not the seed length of the PRG). Here $\lambda$ is the main security parameter (for the BARG and for the commitment scheme). If we set the length of the PRG seed to be at least $n>\lambda^{\delta}$, then we can hope to rely on leakage resilience of the PRG to argue that the output still has high min-entropy even given the BARG proof. In our constructions (Theorem 1.1 and Corollary 3.15), we can use $k$-local PRGs with locality as high as $k=c \log \lambda$ for some constant $c<1$.

We additionally require that our $k$-local PRGs be leakage resilient. Here, we rely on sub-exponential hardness and the Gentry-Wichs leakage-simulation lemma [GW11]. Roughly speaking, it says that if PRG(s) is computationally indistinguishable from $t \stackrel{R}{\leftarrow}\{0,1\}^{m}$ against (non-uniform) adversaries of size at most $s$, then there exists an auxiliary distribution over strings $(\mathbf{t}$, aux*) such that $(\operatorname{PRG}(\mathbf{s})$, aux) is computationally indistinguishable from ( $\mathbf{t}$, aux*) against (non-uniform) adversaries of size at most $s / 2^{|a u x|}$. Here aux is a string that can be arbitrarily correlated with $\mathbf{s}$. Thus, as long as the leakage aux is sufficiently short (as a function of the seed length) and the PRG satisfies sub-exponential security, we can argue that the outputs are still pseudorandom. Finally, we can apply a standard randomness extractor to $t$ to obtain a sequence of bits that are statistically close to uniform (even given aux*).

Dual-mode commitments. The only remaining challenge is to ensure that the (statistically-binding) commitment to the PRG seed s does not leak information about the seed. While it is tempting to rely on computational hiding of the commitment scheme and replace the commitment to the seed with a commitment to the all-zeroes string, this hybrid strategy does not work. The BARG proof (in the opening) is generated using the openings to the commitment

[^1]scheme (i.e., the openings to the commitment are part of the witness for the BARG). Alternatively, we can apply the Gentry-Wichs leakage lemma to argue that the joint distribution (Commit(s), $\operatorname{PRG}(\mathbf{s})$, aux) is computationally indistinguishable from (Commit( $\mathbf{0}$ ), $\mathbf{t}, \mathrm{aux}$ ). As long as the commitment scheme is hiding even for adversaries of size $2^{|a u x|}$, then security follows. However, there is a circular dependency here, as the length of aux is the length of the BARG proof, which is at least as long as the commitment (since the commitment is an input to the BARG relation). As a result, we cannot use complexity leveraging on the commitment as we could with the PRG.

We instead take a different "dual-mode" strategy [GOS06, PW08]. Specifically, we consider a dual-mode commitment scheme where the CRS can be sampled in one of two (computationally indistinguishable) modes: (1) an extractable mode which we use to argue binding; and (2) a statistically hiding mode where the commitments now statistically hide the input. Dual-mode commitments can be constructed from a lossy public-key encryption scheme, which is implied by most number-theoretic intractability assumptions [PW08, HLOV11, AFMP20].

The idea in the hiding proof then is to first switch the dual-mode commitment from binding mode into hiding mode. Observe that this step only changes the public parameters in the scheme. Once the CRS is in hiding mode, the commitments to the PRG seed statistically hide s, regardless of the size of the adversary. In this case, we can appeal to the Gentry-Wichs leakage lemma to argue that the joint distribution (Commit(s), $\operatorname{PRG}(s)$, aux) is computationally indistinguishable from (Commit( $\mathbf{0}), \mathbf{t}$, aux*) assuming only sub-exponential hardness of the PRG. This means the unopened bits in the hidden-bits string are uniformly random and hiding holds. We provide the full details in Section 3 (Theorem 3.4).

Upgrading BARGs to zkBARGs. For completeness, we conclude with a few remarks on using a NIZK for NP to generically upgrade a batch argument to a zero-knowledge batch argument (zkBARG). First, we note that the naïve approach of giving a NIZK proof of knowledge of a BARG proof does not work out of the box. The issue is that the verification algorithm for the BARG needs to read the statements $\left(x_{1}, \ldots, x_{T}\right)$, and thus, the size of the verification circuit scales linear with $T$. Since the size of a NIZK proof can scale polynomially with the size of the verification circuit, the size of the NIZK proof of knowledge of a valid BARG proof for $\left(x_{1}, \ldots, x_{T}\right)$ can scale polynomially with $T$. Nonetheless, we can still apply this general approach in the following settings:

- Index BARGs: An index BARG for an NP language is one where the statements are always fixed to be the integers $1, \ldots, T$ [CJJ21b]. In an index BARG, the verification algorithm only takes the upper bound $T$ as input and is required to run in time that is polylogarithmic in $T$. We can generically compose an index BARG with a NIZK to obtain a zero-knowledge index BARG. We can then apply the index-BARG-to-BARG transformation from [CJJ21b] to the zero-knowledge index BARG for NP to obtain a zkBARG for NP; note here that the [CJJ21b] transformation preserves zero-knowledge.
- BARGs with split verification: A BARG satisfies "split verification" [CJJ21b, WW22, CGJ ${ }^{+} 22$ ] if the verification algorithm decomposes into a (non-succinct) statement-dependent preprocessing step that outputs a short verification key vk and a (succinct) online verification step that takes the preprocessed key vk and the proof $\pi$ and decides whether to accept or reject the proof. Importantly, the online verification step can be implemented by a circuit whose size is polylogarithmic in the number of instances $T$. Given a BARG with a split verification property, it suffices to use a NIZK to prove knowledge of a BARG proof that satisfies the online verification check. This yields a zkBARG with split verification.


## 2 Preliminaries

We write $\lambda$ to denote the security parameter. For a positive integer $n \in \mathbb{N}$, we write $[n]$ to denote the set $\{1, \ldots, n\}$. We use boldface letters (e.g., $\mathbf{x}$ ) to denote vectors. We write poly $(\lambda)$ to denote a fixed function that is $O\left(\lambda^{c}\right)$ for some $c \in \mathbb{N}$ and negl $(\lambda)$ to denote a function that is $o\left(\lambda^{-c}\right)$ for all $c \in \mathbb{N}$. We say an event occurs with overwhelming probability if its complement occurs with negligible probability. We say an algorithm on $\lambda$-bit inputs is efficient if it can be computed by a Boolean circuit of size poly $(\lambda)$, or equivalently, if it can be computed by a Turing machine in $\operatorname{poly}(\lambda)$ time with poly $(\lambda)$ bits of advice.

Let $\mathcal{D}_{1}=\left\{\mathcal{D}_{1, \lambda}\right\}_{\lambda \in \mathbb{N}}$ and $\mathcal{D}_{2}=\left\{\mathcal{D}_{2, \lambda}\right\}_{\lambda \in \mathbb{N}}$ be two ensembles of distributions. For functions $s=s(\lambda)$ and $\varepsilon=\varepsilon(\lambda)$, we say that $\mathcal{D}_{1}$ and $\mathcal{D}_{2}$ are $(s, \varepsilon)$-indistinguishable if for all non-negative polynomials poly $(\cdot)$ and all adversaries $\mathcal{A}$,
modeled as Boolean circuits of size at most $s(\lambda) \cdot \operatorname{poly}(\lambda)$, and all sufficiently large $\lambda \geq \lambda_{\mathcal{A}}$,

$$
\left|\operatorname{Pr}\left[\mathcal{A}(x)=1: x \leftarrow \mathcal{D}_{1, \lambda}\right]-\operatorname{Pr}\left[\mathcal{A}(x)=1: x \leftarrow \mathcal{D}_{2, \lambda}\right]\right| \leq \varepsilon(\lambda) .
$$

We say that $\mathcal{D}_{1}$ and $\mathcal{D}_{2}$ are computationally indistinguishable if there exists a negligible function $\varepsilon(\lambda)=$ negl $(\lambda)$ such that $\mathcal{D}_{1}$ and $\mathcal{D}_{2}$ are $(1, \varepsilon)$-indistinguishable. We say that $\mathcal{D}_{1}$ and $\mathcal{D}_{2}$ are statistically indistinguishable if the statistical distance $\Delta\left(\mathcal{D}_{1}, \mathcal{D}_{2}\right)$ is bounded by a negligible function negl $(\lambda)$.

Min-entropy. We recall some basic definitions on min-entropy. Our definitions are adapted from those in [DRS04]. For a (discrete) random variable $X$, we write $H_{\infty}(X)=-\log \left(\max _{x} \operatorname{Pr}[X=x]\right)$ to denote its min-entropy. For two (possibly correlated) discrete random variables $X$ and $Y$, we define the average min-entropy of $X$ given $Y$ to be $\mathbf{H}_{\infty}(X \mid Y)=-\log \left(\mathbb{E}_{y \leftarrow Y} \max _{x} \operatorname{Pr}[X=x \mid Y=y]\right)$. The optimal probability of an unbounded adversary guessing $X$ given the correlated value $Y$ is $2^{-\mathrm{H}_{\infty}(X \mid Y)}$.

Lemma 2.1 (Conditional Min-Entropy [DRS04, Lemma 2.2]). Let $A, B$ be random variables and suppose there are at most $2^{\lambda}$ elements in the support of $B$. Then $\mathbf{H}_{\infty}(A \mid B) \geq \mathbf{H}_{\infty}(A, B)-\lambda \geq \mathbf{H}_{\infty}(A)-\lambda$.

Gentry-Wichs leakage lemma. Our analysis will rely on the following "leakage lemma" from [GW11]:
Lemma 2.2 (Indistinguishability with Auxiliary Information [GW11, Lemma 3.1]). Let $\lambda$ be a security parameter. There exists a polynomial poly( $\cdot$ ) such that the following property holds. Let $\mathcal{X}=\left\{\mathcal{X}_{\lambda}\right\}_{\lambda \in \mathbb{N}}$ and $\mathcal{Y}=\left\{\mathcal{Y}_{\lambda}\right\}_{\lambda \in \mathbb{N}}$ be arbitrary distributions that are $(s, \varepsilon)$-indistinguishable for some $s=s(\lambda)$ and $\varepsilon=\varepsilon(\lambda)$. Let $X_{\lambda}^{*}=\left\{X_{\lambda}^{*}\right\}_{\lambda \in \mathbb{N}}$ be an augmented distribution where $\mathcal{X}_{\lambda}^{*}$ is a distribution on pairs $\left(x_{\lambda}, \pi_{\lambda}\right)$ where $x_{\lambda} \leftarrow \mathcal{X}_{\lambda}$ and $\pi \in\{0,1\}^{\ell(\lambda)}$ can be arbitrarily correlated with $x_{\lambda}$. Then, there exists a distribution $\boldsymbol{Y}^{*}=\left\{\boldsymbol{y}_{\lambda}^{*}\right\}_{\lambda \in \mathbb{N}}$ with the following properties:

- Each $\mathcal{Y}_{\lambda}^{*}$ is a distribution on tuples $\left(y_{\lambda}, \hat{\pi}_{\lambda}\right)$, where $y_{\lambda} \leftarrow y_{\lambda}$ and $\pi_{\lambda} \in\{0,1\}^{\ell(\lambda)}$.
- The distributions $\mathcal{X}^{*}$ and $\boldsymbol{y}^{*}$ are $\left(s^{\prime}, \varepsilon^{\prime}\right)$-indistinguishable, where $s^{\prime}(\lambda)=s(\lambda) \cdot \operatorname{poly}\left(\varepsilon(\lambda) / 2^{\ell(\lambda)}\right)$ and $\varepsilon^{\prime}(\lambda)=2 \varepsilon(\lambda)$.

Leftover hash lemma. Our construction will also rely on the generalized leftover hash lemma (LHL) from [BDK ${ }^{+}$11]:
Theorem 2.3 (LHL with Conditional Min-Entropy [ $\mathrm{BDK}^{+} 11$, Theorem 3.2, adapted]). Let ( $X, Z$ ) be random variables sampled from some joint distribution $\mathcal{D}$ over $\mathcal{X} \times \mathcal{Z}$. Let $\mathcal{H}=\left\{h: \mathcal{X} \rightarrow\{0,1\}^{v}\right\}$ be a family of universal hash functions, and let $L=\mathrm{H}_{\infty}(X \mid Z)-v$ be the entropy loss. Let $\mathcal{A}(r, h, z)$ be a (possibly probabilistic) distinguisher where

$$
\operatorname{Pr}\left[\mathcal{A}(r, h, z)=1: r \stackrel{\mathrm{R}}{\leftarrow}\{0,1\}^{v}, h \stackrel{\mathrm{R}}{\leftarrow} \mathcal{H},(x, z) \leftarrow \mathcal{D}\right] \leq \varepsilon .
$$

Then, the distinguishing advantage of $\mathcal{A}$ on the following distributions is at most $\sqrt{\varepsilon / 2^{L}}$ :

$$
\left\{(h(x), h, z): \begin{array}{c}
(x, z) \leftarrow \mathcal{D} \\
h \stackrel{\mathrm{R}}{\gtrless} \mathcal{H}
\end{array}\right\} \text { and }\left\{(r, h, z): \begin{array}{c}
(x, z) \leftarrow \mathcal{D} \\
r \stackrel{\mathrm{R}}{\leftarrow}\{0,1\}^{v}, h \stackrel{\mathrm{R}}{ }_{\leftarrow}^{\mathcal{H}}
\end{array}\right\}
$$

Corollary 2.4 (LHL with Conditional Min-Entropy). Let $(X, Z)$ be random variables sampled from some joint distribution $\mathcal{D}$ over $\mathcal{X} \times \mathcal{Z}$. Let $\mathcal{H}=\left\{h: \mathcal{X} \rightarrow\{0,1\}^{v}\right\}$ be a family of universal hash functions. Let $L=\mathbf{H}_{\infty}(X \mid Z)-v$ be the entropy loss. Then the statistical distance between the following distributions is at most $2^{-L / 2}$ :

$$
\left\{(h(x), h, z): \begin{array}{c}
(x, z) \leftarrow \mathcal{D} \\
h \stackrel{\mathrm{R}}{\mathrm{R}} \mathcal{H}
\end{array}\right\} \text { and }\left\{(r, h, z): \begin{array}{c}
(x, z) \leftarrow \mathcal{D} \\
r \stackrel{\mathrm{R}}{\leftarrow}\{0,1\}^{v}, h \mathrm{R}^{\mathrm{R}} \mathcal{H}
\end{array}\right\}
$$

Proof. Follows by setting $\varepsilon=1$ in Theorem 2.3 (which captures all distinguishers).

Pseudorandom generators. We recall the definition of a pseudorandom generator.
Definition 2.5 (Pseudorandom Generator). Let $\lambda$ be a security parameter. A pseudorandom generator with output length $m=m(\lambda)$ is an efficiently-computable function family $\operatorname{PRG}=\left\{\operatorname{PRG}_{\lambda}\right\}_{\lambda \in \mathbb{N}}$ where $\operatorname{PRG}_{\lambda}:\{0,1\}^{\lambda} \rightarrow\{0,1\}^{m(\lambda)}$. For functions $s=s(\lambda)$ and $\varepsilon=\varepsilon(\lambda)$, we say that PRG is $(s, \varepsilon)$-secure if the following two distributions are $(s, \varepsilon)$ indistinguishable:

$$
\left\{\operatorname{PRG}_{\lambda}(x): x \stackrel{\mathrm{R}}{\leftarrow}\{0,1\}^{\lambda}\right\} \quad \text { and } \quad\left\{y \stackrel{\mathrm{R}}{\leftarrow}\{0,1\}^{m(\lambda)}\right\} .
$$

We say that PRG is sub-exponentially secure if there exists a constant $\alpha>0$ and a negligible function $\varepsilon=\operatorname{negl}(\lambda)$ such that PRG is $\left(2^{\lambda^{\alpha}}, \varepsilon\right)$-secure.
Definition 2.6 (Locality of a PRG). We say a PRG: $\{0,1\}^{\lambda} \rightarrow\{0,1\}^{m}$ has locality $k=k(\lambda)$ if each output bit of $\operatorname{PRG}(x)$ is a function of at most $k$ bits of the seed $x$. We say that PRG is computable in $\mathrm{NC}^{0}$ if PRG has constant locality $k=O(1)$.

Local PRGs constructions. Goldreich [Gol00, MST03] gave the first candidate local PRG construction (with constant locality) based on constraint-satisfiability problems over expander graphs. A long line of subsequent works have studied variants of Goldreich's construction [CM01, MST03, CEMT09, App12, ABR12, OW14, AL16, AK19]; we refer to [App15] for an excellent survey of the state of the art. Notably, PRGs with constant locality and super-linear stretch have featured prominently in constructions of indistinguishability obfuscation [Lin17, LT17, JLS21, JLS22].

There has also been an extensive line of works studying attacks and ruling out certain instantiations of local PRGs [MST03, CEMT09, BQ09, OW14, App15, AL16, CDM ${ }^{+}$18, Üna23]. For local PRGs with super-linear stretch $\lambda^{1+\delta}$, the most recent attacks [BQ09, Üna23] run in time roughly $\lambda^{O\left(\lambda^{1-\delta / k}\right)}$ where $k$ is the locality.

Dual-mode commitments. Next, we recall the notion of a "dual-mode" commitment (also called a "mixed commitment") [DN02]. At a high-level, these are non-interactive commitment schemes in the common reference string (CRS) model where the CRS can be sampled from one of two computationally indistinguishable distributions. In one distribution (or mode), the commitment scheme is extractable (i.e., given trapdoor information, one can efficiently extract the committed value from a commitment), and in the other distribution (or mode), the commitment scheme is statistically hiding. ${ }^{4}$ We give the formal definition below:

Definition 2.7 (Dual-Mode Bit Commitment [DN02]). A dual-mode bit commitment scheme is a tuple of efficient algorithms $\Pi_{B C}=($ Setup, Commit, Verify $)$ with the following syntax:

- $\operatorname{Setup}\left(1^{\lambda}\right.$, mode $) \rightarrow($ crs, td $)$ : On input the security parameter $\lambda$ and mode $\in\{$ bind, hide $\}$, the setup algorithm outputs a common reference string crs and a trapdoor td (possibly empty).
- Commit $(\mathrm{crs}, b) \rightarrow(c, \sigma)$ : On input the common reference string crs and a bit $b \in\{0,1\}$, the commit algorithm outputs a commitment $c$ and an opening $\sigma$.
- Verify (crs, $c, b, \sigma) \rightarrow\{0,1\}$ : On input the common reference string crs, a commitment $c$, a bit $b \in\{0,1\}$, and an opening $\sigma$, the verification algorithm outputs a bit $b^{\prime} \in\{0,1\}$.
Moreover, $\Pi_{B C}$ should satisfy the following properties:
- Correctness: For all security parameters $\lambda \in \mathbb{N}$, all bits $b \in\{0,1\}$, all modes mode $\in\{$ bind, hide $\}$,

$$
\operatorname{Pr}\left[\operatorname{Verify}(\operatorname{crs}, c, b, \sigma)=1: \quad(\mathrm{crs}, \mathrm{td}) \leftarrow \operatorname{Setup}\left(1^{\lambda}, \text { mode }\right) ;\right]=1 .
$$

- Mode indistinguishability: For all efficient adversaries $\mathcal{A}$, and sampling (crs ${ }_{\text {bind }}$, td$) \leftarrow \operatorname{Setup}\left(1^{\lambda}\right.$, bind $)$, ( crs $\left._{\text {hide }}, \operatorname{td}^{\prime}\right) \leftarrow \operatorname{Setup}\left(1^{\lambda}\right.$, hide), we have that

$$
\left|\operatorname{Pr}\left[\mathcal{A}\left(1^{\lambda}, \operatorname{crs}_{\text {bind }}\right)=1\right]-\operatorname{Pr}\left[\mathcal{A}\left(1^{\lambda}, \operatorname{crs}_{\text {hide }}\right)=1\right]\right|=\operatorname{neg} \mid(\lambda) .
$$

[^2]- Extractable in binding mode: There exists an efficient algorithm Extract that takes as input a trapdoor td and a string $c \in\{0,1\}^{*}$, and outputs a bit $b \in\{0,1\}$. Then, for all adversaries $\mathcal{A}$,

$$
\operatorname{Pr}\left[\operatorname{Verify}(\mathrm{crs}, c, b, \sigma)=1 \wedge b \neq b^{\prime}: \quad \begin{array}{l}
(\mathrm{crs}, \mathrm{td}) \leftarrow \operatorname{Setup}\left(1^{\lambda}, \mathrm{bind}\right) ; \\
\\
\\
(c, \sigma, b) \leftarrow \mathcal{A}(\mathrm{crs}) ; \\
b^{\prime} \leftarrow \operatorname{Extract}(\mathrm{td}, c)
\end{array}\right]=\operatorname{neg}(\lambda)
$$

- Statistically hiding in hiding mode: For a security parameter $\lambda$ and a bit $\beta \in\{0,1\}$, we define the hiding game between an adversary $\mathcal{A}$ and a challenger as follows:

1. The challenger starts by sampling $(\operatorname{crs}, \mathrm{td}) \leftarrow \operatorname{Setup}\left(1^{\lambda}\right.$, hide $)$ and gives crs to $\mathcal{A}$.
2. Algorithm $\mathcal{A}$ outputs two messages $b_{0}, b_{1} \in\{0,1\}$.
3. The challenger computes $(c, \sigma) \leftarrow \operatorname{Commit}\left(\operatorname{crs}, b_{\beta}\right)$ and replies to $\mathcal{A}$ with $c$.
4. Algorithm $\mathcal{A}$ outputs a bit $b^{\prime} \in\{0,1\}$, which is the output of the experiment.

Then $\Pi_{B C}$ is statistically hiding in hiding mode if there exists a negligible function negl $(\cdot)$ such that for all adversaries $\mathcal{A}$ in the above hiding experiment,

$$
\left|\operatorname{Pr}\left[b^{\prime}=1 \mid \beta=0\right]-\operatorname{Pr}\left[b^{\prime}=1 \mid \beta=1\right]\right|=\operatorname{neg} \mid(\lambda)
$$

Constructions of dual-mode commitments. Dual-mode commitments (with extraction) can be built from any lossy public-key encryption scheme [BHY09], which can in turn be constructed from most standard algebraic assumptions [PW08, HLOV11, AFMP20]. In particular, a commitment to an input $x$ is just a public-key encryption of $x$ and the opening is the corresponding encryption randomness. In extracting mode, the extraction trapdoor is the decryption key.

### 2.1 Non-Interactive Zero-Knowledge Arguments for NP

We recall the notion of a non-interactive zero-knowledge argument for NP [GMR85, BFM88]. We specifically consider the NP-complete language of Boolean circuit satisfiability. Namely, for a Boolean circuit $C:\{0,1\}^{n} \times\{0,1\}^{h} \rightarrow\{0,1\}$, we say that a statement $\mathbf{x} \in\{0,1\}^{n}$ is a yes instance if there exists a witness $\mathbf{w} \in\{0,1\}^{h}$ such that $C(\mathbf{x}, \mathbf{w})=1$.
Definition 2.8 (NIZK Argument for NP). A non-interactive zero-knowledge argument for Boolean circuit satisfiability is a tuple of efficient algorithms $\Pi_{\text {NIZK }}=($ Setup, Prove, Verify $)$ with the following syntax:

- Setup $\left(1^{\lambda}\right) \rightarrow$ crs: On input the security parameter $\lambda \in \mathbb{N}$, the setup algorithm outputs a common reference string crs.
- Prove (crs, $C, \mathbf{x}, \mathbf{w}) \rightarrow \pi$ : On input the common reference string crs, a Boolean circuit $C:\{0,1\}^{n} \times\{0,1\}^{h} \rightarrow$ $\{0,1\}$, a statement $\mathbf{x} \in\{0,1\}^{n}$, and a witness $\mathbf{w} \in\{0,1\}^{h}$, the prove algorithm outputs a proof $\pi$.
- Verify (crs, $C, \mathbf{x}, \pi) \rightarrow b$ : On input the common reference string crs, the Boolean circuit $C:\{0,1\}^{n} \times\{0,1\}^{h} \rightarrow$ $\{0,1\}$, a statement $\mathbf{x} \in\{0,1\}^{n}$, and a proof $\pi$, the verification algorithm outputs a bit $b \in\{0,1\}$.

Moreover, $\Pi_{\text {NIZK }}$ should satisfy the following properties:

- Completeness: For all $\lambda \in \mathbb{N}$, all Boolean circuits $C:\{0,1\}^{n} \times\{0,1\}^{h} \rightarrow\{0,1\}$, all statements $\mathbf{x} \in\{0,1\}^{n}$, and all witnesses $\mathbf{w} \in\{0,1\}^{h}$ where $C(\mathbf{x}, \mathbf{w})=1$,

$$
\operatorname{Pr}\left[\operatorname{Verify}(\operatorname{crs}, C, \mathbf{x}, \pi)=1: \begin{array}{c}
\operatorname{crs} \leftarrow \operatorname{Setup}\left(1^{\lambda}\right) ; \\
\pi \leftarrow \operatorname{Prove}(\operatorname{crs}, C, \mathbf{x}, \mathbf{w})
\end{array}\right]=1
$$

- Computational soundness: For all efficient adversaries $\mathcal{A}$,

$$
\operatorname{Pr}\left[\mathrm{x} \notin \mathcal{L}_{C} \wedge \operatorname{Verify}(\mathrm{crs}, C, \mathrm{x}, \pi)=1: \begin{array}{c}
\operatorname{crs} \leftarrow \operatorname{Setup}\left(1^{\lambda}\right) \\
(C, \mathbf{x}, \pi) \leftarrow \mathcal{A}\left(1^{\lambda}, \mathrm{crs}\right)
\end{array}\right]=\operatorname{negl}(\lambda),
$$

where for a circuit $C:\{0,1\}^{n} \times\{0,1\}^{h} \rightarrow\{0,1\}$, we define $\mathcal{L}_{C}$ to be the language of Boolean circuit satisfiability: $\mathcal{L}_{C}:=\left\{\mathbf{x} \in\{0,1\}^{n}: \exists \mathbf{w} \in\{0,1\}^{h}, C(\mathbf{x}, \mathbf{w})=1\right\}$.

- Computational zero-knowledge: For every efficient adversary $\mathcal{A}$, there exists an efficient simulator $\mathcal{S}=$ $\left(\mathcal{S}_{1}, \mathcal{S}_{2}\right)$ such that for $\operatorname{crs} \leftarrow \operatorname{Setup}\left(1^{\lambda}\right)$ and $(\widetilde{\mathrm{crs}}$, st $\mathcal{S}) \leftarrow \mathcal{S}_{1}\left(1^{\lambda}\right)$, we have that

$$
\left|\operatorname{Pr}\left[\mathcal{A}^{O_{0}(\mathrm{crs}, \cdot, \cdot,)}\left(1^{\lambda}, \mathrm{crs}\right)=1\right]-\operatorname{Pr}\left[\mathcal{A}^{O_{1}(\mathrm{st} \mathcal{S}, \cdot, \cdot)}\left(1^{\lambda}, \widetilde{\mathrm{crs}}\right)=1\right]\right|=\operatorname{neg}(\lambda)
$$

where the oracles $O_{0}$ and $O_{1}$ are defined as follows:

- $O_{0}($ crs, $C, \mathbf{x}, \mathbf{w}):$ On input crs, a circuit $C:\{0,1\}^{n} \times\{0,1\}^{h} \rightarrow\{0,1\}$, a statement $\mathbf{x} \in\{0,1\}^{n}$, and a witness
$\mathbf{w} \in\{0,1\}^{h}$, the oracle outputs $\perp$ if $C(\mathbf{x}, \mathbf{w})=0$. If $C(\mathbf{x}, \mathbf{w})=1$, it outputs Prove (crs, $\left.C, \mathbf{x}, \mathbf{w}\right)$.
- $O_{1}\left(\operatorname{st}_{\mathcal{S}}, C, \mathbf{x}, \mathbf{w}\right)$ : On input the simulator state st $\mathcal{S}_{\mathcal{S}}$, a circuit $C:\{0,1\}^{n} \times\{0,1\}^{h} \rightarrow\{0,1\}$, a statement
$\mathbf{x} \in\{0,1\}^{n}$, and a witness $\mathbf{w} \in\{0,1\}^{h}$, the oracle outputs $\perp$ if $C(\mathbf{x}, \mathbf{w})=0$. If $C(\mathbf{x}, \mathbf{w})=1$, it outputs
$\mathcal{S}_{2}\left(\mathrm{st}_{\mathcal{S}}, C, \mathrm{x}\right)$.


### 2.2 Non-Interactive Batch Arguments for NP

The main cryptographic primitive we consider in this work is a non-interactive batch argument for NP. As before, we consider the NP-complete language of Boolean circuit satisfiability. We now recall the definition of a non-interactive batch argument for NP from [KPY19, CJJ21a]. Our construction relies on the notion of semi-adaptive soundness used in [CJJ21b, WW22, DGKV22, KLVW23, CGJ ${ }^{+} 22$ ].

Definition 2.9 (Batch Argument for NP [CJJ21b, adapted]). A non-interactive batch argument (BARG) for Boolean circuit satisfiability is a tuple of three efficient algorithms $\Pi_{\text {BARG }}=$ (Setup, Prove, Verify) with the following syntax:

- Setup $\left(1^{\lambda}, 1^{T}, 1^{s}\right) \rightarrow$ crs: On input the security parameter $\lambda \in \mathbb{N}$, the number of instances $T \in \mathbb{N}$, and a bound on the circuit size $s \in \mathbb{N}$, the setup algorithm outputs a common reference string crs.
- Prove $\left(\operatorname{crs}, C,\left(\mathbf{x}_{1}, \ldots, \mathbf{x}_{T}\right),\left(\mathbf{w}_{1}, \ldots, \mathbf{w}_{T}\right)\right) \rightarrow \pi$ : On input the common reference string crs, a Boolean circuit $C:\{0,1\}^{n} \times\{0,1\}^{h} \rightarrow\{0,1\}$, statements $\mathbf{x}_{1}, \ldots, \mathbf{x}_{T} \in\{0,1\}^{n}$, and witnesses $\mathbf{w}_{1}, \ldots, \mathbf{w}_{T} \in\{0,1\}^{h}$, the prove algorithm outputs a proof $\pi$.
- Verify $\left(\mathrm{crs}, C,\left(\mathbf{x}_{1}, \ldots, \mathbf{x}_{T}\right), \pi\right) \rightarrow b$ : On input the common reference string crs, the Boolean circuit $C:\{0,1\}^{n} \times$ $\{0,1\}^{h} \rightarrow\{0,1\}$, statements $\mathbf{x}_{1}, \ldots, \mathbf{x}_{T} \in\{0,1\}^{n}$ and a proof $\pi$, the verification algorithm outputs a bit $b \in\{0,1\}$.
Moreover, $\Pi_{B A R G}$ should satisfy the following properties:
- Completeness: For all $\lambda, T, s \in \mathbb{N}$, all circuits $C:\{0,1\}^{n} \times\{0,1\}^{h} \rightarrow\{0,1\}$ of size at most $s$, all statements $\mathbf{x}_{1}, \ldots, \mathbf{x}_{T} \in\{0,1\}^{n}$, and all witnesses $\mathbf{w}_{1}, \ldots, \mathbf{w}_{T} \in\{0,1\}^{h}$ where $C\left(\mathbf{x}_{i}, \mathbf{w}_{i}\right)=1$ for all $i \in[T]$,

$$
\operatorname{Pr}\left[\operatorname{Verify}(\operatorname{crs}, C, \mathbf{x}, \pi)=1: \begin{array}{c}
\operatorname{crs} \leftarrow \operatorname{Setup}\left(1^{\lambda}, 1^{T}, 1^{s}\right) ; \\
\pi \leftarrow \operatorname{Prove}(\operatorname{crs}, C, \mathbf{x}, \mathbf{w})
\end{array}\right]=1,
$$

where $\mathbf{x}=\left(\mathbf{x}_{1}, \ldots, \mathbf{x}_{T}\right)$ and $\mathbf{w}=\left(\mathbf{w}_{1}, \ldots, \mathbf{w}_{T}\right)$.

- Succinct proof size: ${ }^{5}$ There exists a polynomial $\operatorname{poly}(\cdot, \cdot, \cdot)$ such that for all $\lambda, T, s \in \mathbb{N}$, all crs in the support of $\operatorname{Setup}\left(1^{\lambda}, 1^{T}, 1^{s}\right)$, and all Boolean circuits $C:\{0,1\}^{n} \times\{0,1\}^{h} \rightarrow\{0,1\}$ of size at most $s$, the size of the proof $\pi$ output by Prove (crs, $C, \cdot, \cdot)$ satisfies $|\pi| \leq \operatorname{poly}(\lambda, \log T, s)$.

[^3]- Semi-adaptive soundness: For a security parameter $\lambda$, we define the semi-adaptive soundness game between an adversary $\mathcal{A}$ and a challenger as follows:

1. Algorithm $\mathcal{A}$ starts by outputting the number of instances $1^{T}$, the bound on the circuit size $1^{s}$, and an index $i \in[T]$
2. The challenger samples a common reference string crs $\leftarrow \operatorname{Setup}\left(1^{\lambda}, 1^{T}, 1^{s}\right)$ and gives 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 $s$, statements $\left(\mathbf{x}_{1}, \ldots, \mathbf{x}_{T}\right)$ where each $\mathbf{x}_{i} \in\{0,1\}^{n}$, and a proof $\pi$.
4. The output of the experiment is $b=1$ if $\operatorname{Verify}\left(\operatorname{crs}, C,\left(\mathbf{x}_{1}, \ldots, \mathbf{x}_{T}\right), \pi\right)=1$ and for all $\mathbf{w}_{i} \in\{0,1\}^{h}$, $C\left(\mathbf{x}_{i}, \mathbf{w}_{i}\right)=0$. Otherwise, the output is $b=0$.

Then $\Pi_{\text {BARG }}$ satisfies semi-adaptive soundness if for all efficient adversaries $\mathcal{A}, \operatorname{Pr}[b=1]=\operatorname{negl}(\lambda)$ in the semi-adaptive soundness game.

Constructions of batch arguments for NP. Batch arguments for NP have recently been realized from a broad range of standard assumptions including lattice-based assumptions [CJJ21b, DGKV22] as well as assumptions over pairing groups [KVZ21, WW22] and pairing-free groups [CGJ ${ }^{+} 22$ ].

### 2.3 Hidden-Bits Generator

We recall the notion of a hidden-bits generator with subset-dependent proofs from [KMY20]. For a bitstring $\mathbf{r} \in\{0,1\}^{n}$ and a set of indices $I \subseteq[n]$, we write $\mathbf{r}_{I} \in\{0,1\}^{|I|}$ to denote the substring corresponding to the bits of $\mathbf{r}$ indexed by $I$.
Definition 2.10 (Hidden-Bits Generator [KMY20, Definition 11]). A hidden-bits generator with subset-dependent proofs is a tuple of efficient algorithms $\Pi_{\mathrm{HBG}}=$ (Setup, GenBits, Prove, Verify) with the following syntax:

- Setup $\left(1^{\lambda}, 1^{m}\right) \rightarrow$ crs: On input the security parameter $\lambda$, and the output length $m$, the setup algorithm outputs a common reference string crs.
- GenBits(crs) $\rightarrow(\mathbf{r}, \mathrm{st})$ : On input the the common reference string crs, the generator algorithm outputs a string $\mathbf{r} \in\{0,1\}^{m}$ and a state st.
- Prove $(\mathrm{st}, I) \rightarrow \pi$ : On input the state st and a subset $I \subseteq[m]$, the prove algorithm outputs a proof $\pi$.
- Verify $\left(\mathrm{crs}, I, \mathbf{r}_{I}, \pi\right) \rightarrow b$ : On input a common reference string crs, a subset $I \subseteq[m]$, a string $\mathbf{r}_{I} \in\{0,1\}^{|I|}$, and a proof $\pi$, the verification algorithm outputs a bit $b \in\{0,1\}$.
We require $\Pi_{H B G}$ to satisfy the following properties:
- Correctness: For all $m, \lambda \in \mathbb{N}$ and all subsets $I \subseteq[m]$, we have

$$
\operatorname{Pr}\left[\begin{array}{cc}
\operatorname{crs} \leftarrow \operatorname{Setup}\left(1^{\lambda}, 1^{m}\right) ; \\
\operatorname{Verify}\left(\operatorname{crs}, I, \mathbf{r}_{I}, \pi\right)=1: & (\mathbf{r}, \mathrm{st}) \leftarrow \operatorname{GenBits}(\mathrm{crs}) ; \\
\pi \leftarrow \operatorname{Prove}(\mathrm{st}, I)
\end{array}\right]=1
$$

- Somewhat computational binding: For every crs in the support of the algorithm $\operatorname{Setup}\left(1^{\lambda}, 1^{m}\right)$, there exists a set $\mathcal{V}^{\text {crs }}$ with the following properties:
(i) Output sparsity. There exists a universal constant $\gamma<1$ and a fixed polynomial $p(\cdot)$ such that for every polynomial $m=m(\lambda)$, and every crs in the support of $\operatorname{Setup}\left(1^{\lambda}, 1^{m}\right),\left|\mathcal{V}^{\text {crs }}\right| \leq 2^{m^{\gamma} \cdot p(\lambda)}$
(ii) Computational binding. For every efficient and stateful adversary $\mathcal{A}$,

$$
\operatorname{Pr}\left[\mathbf{r}_{I} \notin \mathcal{V}_{I}^{\text {crs }} \wedge \operatorname{Verify}\left(\operatorname{crs}, I, \mathbf{r}_{I}, \pi\right)=1: \quad \begin{array}{c}
1^{m} \leftarrow \mathcal{A}\left(1^{\lambda}\right) ; \\
\\
\left(I, \mathbf{r}_{I}, \pi\right) \leftarrow \mathcal{A}(\mathrm{crs})
\end{array}\right]=\operatorname{setup}\left(1^{\lambda}, 1^{m}\right) ; \quad \operatorname{negl}(\lambda),
$$

where $\mathcal{V}_{I}^{\text {crs }}:=\left\{\mathbf{r}_{I}: \mathbf{r} \in \mathcal{V}^{\text {crs }}\right\}$.

- Computationally hiding: For every polynomial $m=m(\lambda)$, every subset $I \subseteq[m]$, and all efficient adversaries $\mathcal{A}$, we have

$$
\left|\operatorname{Pr}\left[\mathcal{A}\left(\operatorname{crs}, I, \mathbf{r}_{I}, \pi, \mathbf{r}_{\bar{I}}\right)=1\right]-\operatorname{Pr}\left[\mathcal{A}\left(\operatorname{crs}, I, \mathbf{r}_{I}, \pi, \mathbf{r}_{\bar{I}}^{\prime}\right)=1\right]\right|=\operatorname{negl}(\lambda),
$$

where crs $\leftarrow \operatorname{Setup}\left(1^{\lambda}, 1^{m}\right),(\mathbf{r}, \mathrm{st}) \leftarrow \operatorname{GenBits}(\mathrm{crs}), \pi \leftarrow \operatorname{Prove}(\mathrm{st}, I), \mathbf{r}^{\prime} \stackrel{\mathrm{R}}{\leftarrow}\{0,1\}^{m}$, and $\bar{I}=[m] \backslash I$.
Theorem 2.11 (NIZK from Hidden-Bits Generator [KMY20, Theorem 5]). If there exists a hidden-bits generator with subset-dependent proofs, then there exists a computational NIZK argument for NP.

## 3 Hidden-Bits Generator from Batch Arguments

In this section, we show how to construct a hidden-bits generator with subset-dependent proofs using a batch argument for NP together with a dual-mode commitment and a low-complexity PRG. Then combined with Theorem 2.11, we obtain a NIZK for NP from the same underlying set of assumptions.

Construction 3.1 (Hidden-Bits Generator from Batch Arguments). Let $\lambda \in \mathbb{N}$ be a security parameter and $m$ be an output length parameter. Let $n=n(\lambda, m)$ be a PRG seed length parameter and let $B=B(\lambda, m)$ be a block length parameter. These parameters will be determined in the security analysis. Our construction relies on the following primitives:

- Let $G_{\lambda}:\{0,1\}^{\lambda} \rightarrow\{0,1\}^{\ell(\lambda)}$ be a family of PRGs. Let $k=k(\lambda)$ be the locality of the PRG (i.e., each output bit of $G_{\lambda}$ depends on at most $k$ input bits). In the following description, we require that $\ell(n) \geq m B$.
- Let $\Pi_{B C}=(B C . S e t u p, B C . C o m m i t, B C . V e r i f y)$ be a dual-mode bit commitment scheme.
- Let $\Pi_{\text {BARG }}=($ BARG.Setup, BARG.Prove, BARG.Verify $)$ be a batch argument for NP with proof length $\ell_{\mathrm{BARG}}=$ $\ell_{\text {BARG }}(\lambda, T, s)$, where $s$ denotes the size of the underlying NP relation and $T$ denotes the number of instances.
- For an index $i \in[\ell]$ where $\ell=\ell(n)$, let $i_{1}, \ldots, i_{k} \in[n]$ be the indices of the $k$ seed bits on which the $i^{\text {th }}$ output bit of $G_{n}(\cdot)$ depends. Let $s \in\{0,1\}^{n}$ be a seed for the PRG, and let $G_{n}^{(i)}:\{0,1\}^{k} \rightarrow\{0,1\}$ be the circuit that takes as input the seed bits $s_{i_{1}}, \ldots, s_{i_{k}} \in\{0,1\}$ and outputs the $i^{\text {th }}$ bit of $G_{n}(\mathbf{s})$. Then, for a common reference string $\operatorname{crs}_{B C}$ for the bit commitment scheme, define the NP relation $\mathcal{R}\left[n, \mathrm{crs}_{\mathrm{BC}}\right]$ as follows:

Hard-wired: PRG seed length $n$, common reference string $\operatorname{crs}_{B C}$
Statement: circuit $G_{n}^{(i)}:\{0,1\}^{k} \rightarrow\{0,1\}$, commitments $c_{1}, \ldots, c_{k}$, output $t \in\{0,1\}$
Witness: bits $s_{1}, \ldots, s_{k} \in\{0,1\}$, openings $\sigma_{1}, \ldots, \sigma_{k}$
Output 1 if all of the following conditions hold:

- For each $i \in[k], \mathrm{BC} . \operatorname{Verify}\left(\operatorname{crs}_{\mathrm{BC}}, c_{i}, s_{i}, \sigma_{i}\right)=1$;
- $t=G_{n}^{(i)}\left(s_{1}, \ldots, s_{k}\right)$.

Otherwise, output 0 .

Figure 1: Relation $\mathcal{R}\left[n, \operatorname{crs}_{\mathrm{BC}}\right]\left(\left(G_{n}^{(i)}, c_{1}, \ldots, c_{k}, t\right),\left(s_{1}, \ldots, s_{k}, \sigma_{1}, \ldots, \sigma_{k}\right)\right)$.
We now construct our hidden-bits generator $\Pi_{\text {HBG }}=($ Setup, GenBits, Prove, Verify $)$ as follows:

- Setup $\left(1^{\lambda}, 1^{m}\right)$ : On input the security parameter $\lambda$ and the output length $m$, the setup algorithm proceeds as follows:

1. Sample a CRS for the dual-mode commitment scheme: $\left(\operatorname{crs}_{B C}, t d\right) \leftarrow B C \cdot \operatorname{Setup}\left(1^{\lambda}\right.$, bind $)$.
2. Let $n=n(\lambda, m)$ be the PRG seed length. Let $C$ be the circuit that computes the NP relation $\mathcal{R}\left[n, \operatorname{crs}_{\mathrm{BC}}\right]$. Sample a common reference string $\operatorname{crs}_{\text {BARG }} \leftarrow \operatorname{BARG} \cdot \operatorname{Setup}\left(1^{\lambda}, 1^{m B}, 1^{|C|}\right)$.
3. Let $B=B(\lambda, m)$ be the block size and sample $\mathbf{v}_{1}, \ldots, \mathbf{v}_{m} \stackrel{R}{\leftarrow}\{0,1\}^{B}$.

Output the common reference string crs $=\left(n\right.$, crs $\left._{\text {BARG }}, \mathrm{crs}_{\mathrm{BC}}, \mathrm{v}_{1}, \ldots, \mathbf{v}_{m}\right)$.

- GenBits(crs): On input the common reference string crs $=\left(n, \operatorname{crs}_{\mathrm{BARG}}, \mathrm{crs}_{\mathrm{BC}}, \mathbf{v}_{1}, \ldots, \mathbf{v}_{m}\right)$, the generator algorithm proceeds as follows:

1. Sample a PRG seed $\mathbf{s} \stackrel{\mathrm{R}}{\leftarrow}\{0,1\}^{n}$, and compute $\mathbf{t}=\{0,1\}^{m B} \leftarrow G_{n}(\mathbf{s}) .{ }^{6}$ Then, for each $i \in[n]$, compute a commitment $\left(c_{i}, \sigma_{i}\right) \leftarrow \mathrm{BC}$.Commit $\left(\operatorname{crs}_{\mathrm{BC}}, s_{i}\right)$ to the bits of the seed.
2. Split $\mathbf{t}=\mathbf{t}_{1}\left\|\mathbf{t}_{2}\right\| \cdots \| \mathbf{t}_{m}$ into blocks where each $\mathbf{t}_{i} \in\{0,1\}^{B}$ for each $i \in[m]$. Next, for each $i \in[m]$, compute $r_{i} \leftarrow \mathbf{v}_{i}^{\top} \mathbf{t}_{i}$ (where the vectors $\mathbf{v}_{i}$ and $\mathbf{t}_{i}$ are interpreted as vectors in $\mathbb{Z}_{2}^{B}$ ).

The algorithm outputs the hidden-bits string $\mathbf{r}=r_{1}\left\|r_{2}\right\| \cdots \| r_{m} \in\{0,1\}^{m}$ together with the generator state st $=\left(n, \operatorname{crs}_{\text {BARG }}, \operatorname{crs}_{\mathrm{BC}}, \mathrm{s}, c_{1}, \ldots, c_{n}, \sigma_{1}, \ldots, \sigma_{n}\right)$.

- Prove(st, $I$ ): On input the state $s t=\left(n, \operatorname{crs}_{\mathrm{BARG}}, \operatorname{crs}_{\mathrm{BC}}, \mathrm{s}, c_{1}, \ldots, c_{n}, \sigma_{1}, \ldots, \sigma_{n}\right)$ and a set of indices $I \subseteq[m]$, the prove algorithm proceeds as follows:

1. Let $\mathbf{t}=G_{n}(\mathbf{s})$ and parse $\mathbf{t}=\mathbf{t}_{1}\left\|\mathbf{t}_{2}\right\| \cdots \| \mathbf{t}_{m}$ where each $\mathbf{t}_{i} \in\{0,1\}^{B}$. We will also use the notation $\mathbf{t}[i]:=\mathbf{t}_{i}$ to refer to the $i^{\text {th }}$ block of t and $G_{n}[i]:=G_{n}^{(i)}$ to refer to the circuit computing the $i^{\text {th }}$ bit of $G_{n}(\mathrm{~s})$. In the analysis, we will often associate an index $i \in[m B]$ with a pair $(j, \beta) \in[m] \times[B]$ and vice versa (where $i=(j-1) B+\beta)$.
2. Let $I=\left\{i^{(1)}, \ldots, i^{(L)}\right\}$, where the indices $i^{(1)}, \ldots, i^{(L)} \in[m]$ are in sorted order. For each $j \in[L], \beta \in[B]$, let $\mathbf{t}\left[i^{(j)}, \beta\right] \in\{0,1\}$ denote the $\beta^{\text {th }}$ bit of $\mathbf{t}\left[i^{(j)}\right]$.
3. By construction, the value of $\mathrm{t}\left[i^{(j)}, \beta\right]$ depends on at most $k$ bits of $s$. We define $G_{n}\left[i^{(j)}, \beta\right]$ to denote the circuit that reads up to $k$ bits of $\boldsymbol{s}$ and outputs $\mathrm{t}\left[i^{(j)}, \beta\right]$. Next, we define the function idx: $[L] \times[B] \times[k] \rightarrow$ $[n]$ where $\operatorname{idx}\left(i^{(j)}, \beta, \gamma\right)$ outputs the $\gamma^{\text {th }}$ input bit of $s$ on which the output bit $\mathbf{t}\left[i^{(j)}, \beta\right]$ depends. In particular, the inputs to the circuit $G_{n}\left[i^{(j)}, \beta\right]$ consist of bits idx $\left(i^{(j)}, \beta, 1\right), \ldots, \operatorname{idx}\left(i^{(j)}, \beta, k\right)$ of $\mathbf{s}$.
4. For each $j \in[L]$ and $\beta \in[B]$, define the statement $x_{j, \beta}$ and associated witness $w_{j, \beta}$ as follows:

$$
\begin{align*}
x_{j, \beta} & =\left(G_{n}\left[i^{(j)}, \beta\right], c_{\mathrm{idx}\left(i^{(j)}, \beta, 1\right)}, \ldots, c_{\mathrm{idx}\left(i^{(j)}, \beta, k\right)}, \mathrm{t}\left[i^{(j)}, \beta\right]\right)  \tag{3.1}\\
w_{j, \beta} & =\left(s_{\mathrm{idx}\left(i^{(j)}, \beta, 1\right)}, \ldots, s_{\mathrm{idx}\left(i^{(j)}, \beta, k\right)}, \sigma_{\mathrm{idx}\left(i^{(j)}, \beta, 1\right)}, \ldots, \sigma_{\mathrm{idx}\left(i^{(j)}, \beta, k\right)}\right) \tag{3.2}
\end{align*}
$$

Let $C$ be the circuit that computes the NP relation in Fig. 1. Then, compute the proof
$\pi_{\text {BARG }} \leftarrow$ BARG.Prove $\left(\operatorname{crs}_{\text {BARG }}, C,\left(x_{1,1}, \ldots, x_{1, B}, \ldots, x_{L, 1}, \ldots, x_{L, B}\right),\left(w_{1,1}, \ldots, w_{1, B}, \ldots, w_{L, 1}, \ldots, w_{L, B}\right)\right)$.
5. Output $\pi=\left(\pi_{\mathrm{BARG}},\left(c_{1}, \ldots, c_{n}\right),\left(\mathbf{t}_{i^{(1)}}, \ldots, \mathbf{t}_{i^{(L)}}\right)\right)$.

- Verify $\left(\mathrm{crs}, I, \mathbf{r}_{I}, \pi\right)$ : On input crs $=\left(n, \operatorname{crs}_{\mathrm{BARG}}, \operatorname{crs}_{\mathrm{BC}}, \mathbf{v}_{1}, \ldots, \mathbf{v}_{m}\right)$, a set of indices $I=\left\{i^{(1)}, \ldots, i^{(L)}\right\} \subseteq[m]$ (in sorted order), a string $\mathbf{r}_{I} \in\{0,1\}^{L}$, and a proof $\pi=\left(\pi_{\mathrm{BARG}},\left(c_{1}, \ldots, c_{n}\right),\left(\mathbf{t}_{i^{(1)}}, \ldots, \mathbf{t}_{i^{(L)}}\right)\right)$, the verification algorithm proceeds as follows:

1. For each $j \in[L]$, let $r_{i}(j) \in\{0,1\}$ be the bit of $\mathbf{r}_{I}$ associated with index $i^{(j)}$. Then, for each $j \in[L]$, check that $r_{i^{(j)}}=\mathbf{v}_{i^{(j)}}^{\top} \mathbf{t}_{i^{(j)}}$. Output 0 if any check fails.
2. Using the commitments $c_{1}, \ldots, c_{n}$ and the bits of $\mathbf{t}_{i^{(1)}}, \ldots, \mathbf{t}_{i^{(L)}}$, construct the statements $x_{j, \beta}$ for each $j \in[L]$ and $\beta \in[B]$ according to Eq. (3.1). Let $C$ be the circuit that computes the NP relation $\mathcal{R}\left[n, \mathrm{crs}_{\mathrm{BC}}\right]$ in Fig. 1.
3. Output BARG.Verify $\left(\operatorname{crs}_{\mathrm{BARG}}, C,\left(x_{1,1}, \ldots, x_{1, B}, \ldots, x_{L, 1}, \ldots, x_{L, B}\right), \pi_{\mathrm{BARG}}\right)$.
[^4]Theorem 3.2 (Correctness). If $\Pi_{\mathrm{BARG}}$ is complete and $\Pi_{\mathrm{BC}}$ is correct, then Construction 3.1 is correct.
Proof. Take any security parameter $\lambda$, output length $m$, and set of indices $I \subseteq[m]$. Let crs $\leftarrow \operatorname{Setup}\left(1^{\lambda}, 1^{m}\right)$ where crs $=\left(n, \operatorname{crs}_{\text {BARG }}, \operatorname{crs}_{B C}, \mathbf{v}_{1}, \ldots, \mathbf{v}_{m}\right)$. Let $(\mathbf{r}, \mathrm{st}) \leftarrow \operatorname{GenBits}(\mathrm{crs})$ and $\pi \leftarrow \operatorname{Prove}(\mathrm{st}, I)$. Consider the output of Verify (crs, $\left.I, \mathbf{r}_{I}, \pi\right)$.

- By construction of GenBits, $r_{i}=\mathbf{v}_{i}^{\top} \mathbf{t}_{i}$ for all $i \in[L]$. Thus, the first set of checks in Verify pass.
- Next, for each $j \in[L]$ and $\beta \in[B]$, let $x_{j, \beta}$ and $w_{j, \beta}$ be the statement and witness defined as in Eqs. (3.1) and (3.2). By correctness of $\Pi_{\mathrm{BC}}$, it follows that $\left(x_{j, \beta}, w_{j, \beta}\right) \in \mathcal{R}\left[n, \mathrm{crs}_{\mathrm{BC}}\right]$.
- Let $\pi=\left(\pi_{\mathrm{BARG}},\left(c_{1}, \ldots, c_{n}\right),\left(\mathbf{t}_{i^{(1)}}, \ldots, \mathbf{t}_{i^{(L)}}\right)\right)$. Since $\left(x_{j, \beta}, w_{j, \beta}\right) \in \mathcal{R}\left[n, \operatorname{crs}_{\mathrm{BC}}\right]$ for all $j \in[L]$ and $\beta \in[B]$, completeness of $\Pi_{\text {BARG }}$ implies that

$$
\text { BARG.Verify }\left(\operatorname{crs}_{\mathrm{BARG}}, C,\left(x_{1,1}, \ldots, x_{1, B}, \ldots, x_{L, 1}, \ldots, x_{L, B}\right), \pi_{\mathrm{BARG}}\right)=1
$$

Correspondingly, Verify (crs, $\left.I, \mathbf{r}_{I}, \pi\right)=1$, as required.
Theorem 3.3 (Somewhat Computational Binding). Let $\lambda$ be a security parameter. Suppose there exists a universal constant $\delta<1$ and a fixed polynomial $p(\cdot)$ such that for every polynomial $m=m(\lambda)$, it follows that $n=n(\lambda, m) \leq$ $m^{\delta} \cdot p(\lambda)$. Suppose also that $\Pi_{\mathrm{BARG}}$ satisfies semi-adaptive soundness, $\Pi_{\mathrm{BC}}$ is extractable in binding mode, and that $B=B(\lambda, m)$ is polynomially bounded. Then, Construction 3.1 satisfies somewhat computational binding.
Proof. Let crs $=\left(n, \operatorname{crs}_{\text {BARG }}, \operatorname{crs}_{\mathrm{BC}}, \mathbf{v}_{1}, \ldots, \mathbf{v}_{m}\right)$ be a common reference string in the support of $\operatorname{Setup}\left(1^{\lambda}, 1^{m}\right)$. We define the set $\mathcal{V}^{\text {crs }} \subset\{0,1\}^{m}$ as follows:

$$
\mathcal{V}^{\mathrm{crs}}:=\left\{\left(\mathbf{v}_{1}^{\top} \mathbf{t}_{1}, \ldots, \mathbf{v}_{m}^{\top} \mathbf{t}_{m}\right) \mid \exists \mathbf{s} \in\{0,1\}^{n}: \mathbf{t}_{1}\|\cdots\| \mathbf{t}_{m}=G_{n}(\mathbf{s})\right\}
$$

We now show that each of the requirements in Definition 2.10 is satisfied:
Output sparsity. This is immediate from the construction: $\left|\mathcal{V}^{\text {crs }}\right| \leq 2^{n} \leq 2^{m^{\delta}} \cdot p(\lambda)$.

Computational binding. To argue computational binding, we appeal to the fact that $\Pi_{\mathrm{BC}}$ is extractable in binding mode and to semi-adaptive soundness of $\Pi_{\text {BARG }}$. Formally, suppose there is an efficient adversary $\mathcal{A}$ that breaks computational binding of Construction 3.1 with non-negligible advantage $\varepsilon$. We construct an adversary $\mathcal{B}$ that breaks semi-adaptive soundness of the BARG as follows:

1. Algorithm $\mathcal{B}$ starts running $\mathcal{A}$ on input the security parameter $1^{\lambda}$. Algorithm $\mathcal{A}$ chooses the output length $1^{m}$.
2. Algorithm $\mathcal{B}$ then samples $\left(\right.$ crs $\left._{B C}, \mathrm{td}\right) \leftarrow B C \cdot \operatorname{Setup}\left(1^{\lambda}\right.$, bind) as well as an index $i^{*}{ }^{*}{ }^{\mathrm{R}}[m B]$. It outputs $1^{m B}$ as the number of instances, $1^{s}$ as the size of the circuit (for computing the relation $\mathcal{R}\left[n, \operatorname{crs}_{\mathrm{BC}}\right]$ in Fig. 1), and the chosen index $i^{*}$.
3. Algorithm $\mathcal{B}$ receives $\operatorname{crs}_{\text {BARG }}$ from its challenger. Then, it samples the strings $\mathbf{v}_{1}, \ldots, \mathbf{v}_{m} \stackrel{R}{\leftarrow}\{0,1\}^{B}$. It gives $\mathrm{crs}=\left(n, \mathrm{crs}_{\mathrm{BARG}}, \mathrm{crs}_{\mathrm{BC}}, \mathrm{v}_{1}, \ldots, \mathrm{v}_{m}\right)$ to $\mathcal{A}$.
4. Algorithm $\mathcal{A}$ outputs an opening $\left(I, \mathbf{r}_{I}, \pi\right)$.
5. Algorithm $\mathcal{B}$ parses $I=\left\{i^{(1)}, \ldots, i^{(L)}\right\}$ and $\pi=\left(\pi_{\mathrm{BARG}},\left(c_{1}, \ldots, c_{n}\right),\left(\mathbf{t}_{i^{(1)}}, \ldots, \mathbf{t}_{i^{(L)}}\right)\right)$. It constructs the statement $\mathbf{x}=\left(x_{1,1}, \ldots, x_{1, B}, \ldots, x_{L, 1}, \ldots, x_{L, B}\right)$ from $c_{1}, \ldots, c_{m}$ and $\mathbf{t}_{i^{(1)}}, \ldots, \mathbf{t}_{i^{(L)}}$ according to Eq. (3.1) and defines $C$ to be the circuit that computes the NP relation $\mathcal{R}\left[n, \mathrm{crs}_{B C}\right]$ from Fig. 1 .
6. Now, for each $i \in[n]$, algorithm $\mathcal{B}$ runs $s_{i} \leftarrow \operatorname{BC}$. Extract $\left(\operatorname{td}, c_{i}\right)$. Let $s=s_{1}\|\cdots\| s_{n} \in\{0,1\}^{n}$ be the extracted seed. Algorithm $\mathcal{B}$ now outputs ( $C, \mathbf{x}, \pi_{\mathrm{BARG}}$ ) if the index $i^{*}$ satisfies $i^{*} \in I$ and $t_{i^{*}} \neq t_{i^{*}}^{\prime}$, where $\mathbf{t}^{\prime}=G_{n}(\mathbf{s})$. Otherwise algorithm $\mathcal{B}$ outputs $\perp$.

First, we argue that algorithm $\mathcal{B}$ is admissible.

- Suppose that $t_{i^{*}} \neq t_{i^{*}}^{\prime}$ where $\mathbf{t}^{\prime}=G_{n}(\mathbf{s})$. Write $i^{*}=\left(i^{(j)}, \beta\right) \in[m] \times[B]$. Then,

$$
\mathbf{t}\left[i^{(j)}, \beta\right]=t_{i^{*}} \neq G_{n}^{\left(i^{*}\right)}(\mathbf{s})=G_{n}\left[i^{(j)}, \beta\right]\left(s_{\mathrm{idx}\left(i^{(j)}, \beta, 1\right)}, \ldots, s_{\mathrm{idx}\left(i^{(j)}, \beta, k\right)}\right) .
$$

- Consider the instance

$$
x_{i}(j), \beta=\left(G_{n}\left[i^{(j)}, \beta\right], c_{\mathrm{idx}\left(i^{(j)}, \beta, 1\right)}, \ldots, c_{\mathrm{idx}\left(i^{(j)}, \beta, k\right)}, \mathbf{t}\left[i^{(j)}, \beta\right]\right),
$$

and any candidate witness

$$
w_{i^{(j)}, \beta}=\left(s_{\mathrm{idx}\left(i^{(j)}, \beta, 1\right)}^{\prime}, \ldots, s_{\mathrm{idx}\left(i^{(j)}, \beta, k\right)}^{\prime}, \sigma_{\mathrm{idx}\left(i^{(j)}, \beta, 1\right)}, \ldots, \sigma_{\mathrm{idx}\left(i^{(j)}, \beta, k\right)}\right) .
$$

We consider two possibilities:

- Suppose there exists $\gamma \in[k]$ where $s_{\mathrm{idx}\left(i^{(j)}, \beta, \gamma\right)}^{\prime} \neq s_{\mathrm{idx}\left(i^{(j)}, \beta, \gamma\right)}$. By extractability of $\Pi_{\mathrm{BC}}$, with overwhelming probability over the choice of $\operatorname{crs}_{\mathrm{BC}}$,

$$
\text { BC.Verify }\left(\operatorname{crs}_{\mathrm{BC}}, c_{\mathrm{idx}\left(i^{(j)}, \beta, \gamma\right)}, s_{\mathrm{idx}\left(i^{(j)}, \beta, \gamma\right)}^{\prime}, \sigma_{\mathrm{idx}\left(i^{(j)}, \beta, \gamma\right)}\right)=0 .
$$

Correspondingly, $\mathcal{R}\left[\operatorname{crs}_{\mathrm{BC}}\right]\left(x_{i^{(j)}, \beta}, w_{i^{(j)}, \beta}\right)=0$.

- Suppose that for all $\gamma \in[k], s_{\mathrm{idx}\left(i^{(j)}, \beta, \gamma\right)}^{\prime}=s_{\mathrm{idx}\left(i^{(j)}, \beta, \gamma\right)}$. In this case,

$$
G_{n}\left[i^{(j)}, \beta\right]\left(s_{\mathrm{idx}\left(i^{(j)}, \beta, 1\right)}^{\prime}, \ldots, s_{\mathrm{idx}\left(i^{(j)}, \beta, k\right)}^{\prime}\right)=G_{n}\left[i^{(j)}, \beta\right]\left(s_{\mathrm{idx}\left(i^{(j)}, \beta, 1\right)}, \ldots, s_{\mathrm{idx}\left(i^{(j)}, \beta, k\right)}\right) \neq \mathbf{t}\left[i^{(j)}, \beta\right]
$$

Once again, $\mathcal{R}\left[\operatorname{crs}_{\mathrm{BC}}\right]\left(x_{i(j), \beta}, w_{i(j), \beta}\right)=0$.
Thus, we conclude that if $t_{i^{*}} \neq t_{i^{*}}^{\prime}$, then instance $x_{i^{(j)}, \beta}=x_{i^{*}}$ is false with all but negligible probability over the choice of $\operatorname{crs}_{\mathrm{BC}}$. Algorithm $\mathcal{B}$ only produces an output when $t_{i^{*}} \neq t_{i^{*}}^{\prime}$ (i.e., when $x_{i^{*}}$ is false), so algorithm $\mathcal{B}$ is admissible for the semi-adaptive soundness game. To conclude the proof, we compute the advantage of $\mathcal{B}$. In the semi-adaptive soundness game, the challenger constructs crs $_{\text {BARG }}$ using $\operatorname{BARG}$. $\operatorname{Setup}\left(1^{\lambda}, 1^{m B}, 1^{|C|}\right)$, which is identical to the distribution in computational binding game. Thus, algorithm $\mathcal{B}$ perfectly simulates an execution of the binding game for $\mathcal{A}$. This means that with probability $\varepsilon$, algorithm $\mathcal{A}$ outputs $\left(I, r_{I}, \pi\right)$ where $\pi=\left(\pi_{\mathrm{BARG}},\left(c_{1}, \ldots, c_{n}\right),\left(\mathbf{t}_{i^{(1)}}, \ldots, \mathbf{t}_{i^{(L)}}\right)\right)$ with the following two properties:

- Let $\mathbf{x}=\left(x_{1,1}, \ldots, x_{1, B}, \ldots, x_{L, 1}, \ldots, x_{L, B}\right)$ be the statement constructed from $c_{1}, \ldots, c_{n}$ and $\mathbf{t}_{i^{(1)}}, \ldots, \mathbf{t}_{i^{(L)}}$ according to Eq. (3.1). Then, we have BARG.Verify $\left(\operatorname{crs}_{\mathrm{BARG}}, C, \mathrm{x}, \pi_{\mathrm{BARG}}\right)=1$, where $C$ is the circuit computing the NP relation $\mathcal{R}\left[\mathrm{crs}_{\mathrm{BC}}\right]$ from Fig. 1.
- The bits $\mathbf{r}_{I}$ satisfy $\mathbf{r}_{I} \notin \mathcal{V}_{I}^{\text {crs. }}$. This means that for every seed $\mathbf{s} \in\{0,1\}^{n}$, there must exist some output index $i \in[m B]$ such that $t_{i} \neq G_{n}^{(i)}(\mathbf{s})$.

Thus, with probability $\varepsilon$, both of the above conditions hold. In particular, this means that $\mathcal{A}$ outputs $\left(I, \mathbf{r}_{I}, \pi\right)$ such that there exists some index $\hat{i} \in[m B]$ where $t_{\hat{i}} \neq t_{\hat{i}}^{\prime}=G_{n}^{(\hat{i})}(\mathrm{s})$. Now, algorithm $\mathcal{B}$ samples $i^{*} \stackrel{\mathrm{R}}{\leftarrow}[m B]$ and moreover $i^{*}$ is independent of $\mathcal{A}$ 's view. Thus, $\hat{i}=i^{*}$ with probability at least $1 / m B$, in which case, algorithm $\mathcal{B}$ outputs the instance $(C, \mathbf{x})$ with the proof $\pi_{\mathrm{BARG}}$. Again from the above conditions, BARG.Verify $\left(\operatorname{crs}_{\mathrm{BARG}}, C, \mathbf{x}, \pi_{\mathrm{BARG}}\right)=1$, and algorithm $\mathcal{B}$ succeeds in breaking semi-adaptive soundness of $\Pi_{\text {BARG }}$. We conclude that algorithm $\mathcal{B}$ breaks semi-adaptive soundness of $\Pi_{\text {BARG }}$ with advantage $\varepsilon / m B-\operatorname{negl}(\lambda)$, and the claim follows.

Theorem 3.4 (Computational Hiding). Suppose the following conditions hold:

- The $P R G G_{\lambda}$ is sub-exponentially secure (i.e., there exists a constant $\alpha>0$ and a negligible function $\varepsilon_{\text {PRG }}=\operatorname{negl}(\lambda)$ such that $G_{\lambda}$ is $\left(2^{\lambda^{\alpha}}, \varepsilon_{\mathrm{PRG}}\right)$-secure $)$.
- The bit commitment scheme $\Pi_{\mathrm{BC}}$ satisfies mode indistinguishability and is statistically hiding in hiding mode.
- The length of the $\operatorname{BARG}$ proof $\ell_{\mathrm{BARG}}=\operatorname{poly}(\lambda, m)$ is polynomially-bounded.
- The length of the PRG seed satisfies $n=n(\lambda, m) \geq \max \left(\lambda, \ell_{\mathrm{BARG}}^{c}\right)$ for some constant $c>1 / \alpha$, and the block size satisfies $B=B(\lambda, m) \geq \omega(\log \lambda)+\ell_{\text {BARG }}$.

Then, for all polynomially-bounded $m=m(\lambda)$, Construction 3.1 is computationally hiding.
Proof. Let $I \subseteq[m]$ be an arbitrary subset. We start by defining two distributions $\mathcal{D}_{\text {real }}$ and $\mathcal{D}_{\text {ideal }}$ that will be helpful for our analysis:

- $\mathcal{D}_{\text {real }}\left(1^{\lambda}\right)$ : On input the security parameter $\lambda \in \mathbb{N}$, the real distribution constructs the output as follows:
- Sample $\left(\right.$ crs $\left._{B C}, \mathrm{td}\right) \leftarrow \operatorname{BC} . \operatorname{Setup}\left(1^{\lambda}\right.$, hide).
- Sample $\operatorname{crs}_{\text {BARG }} \leftarrow \operatorname{BARG} . \operatorname{Setup}\left(1^{\lambda}, 1^{m B}, 1^{|C|}\right)$.
- Let $n=n(\lambda, m)$ and sample $\mathbf{s} \stackrel{\mathrm{R}}{\leftarrow}\{0,1\}^{n}$ and compute $\left(c_{i}, \sigma_{i}\right) \leftarrow$ BC.Commit $\left(\operatorname{crs}_{\mathrm{BC}}, s_{i}\right)$ for each $i \in[n]$.
- Compute $\mathbf{t} \leftarrow G_{n}(\mathbf{s})$ and output $\left(I, \operatorname{crs}_{B A R G}, \operatorname{crs}_{B C}, c_{1}, \ldots, c_{n}, \mathbf{t}\right)$.
- $\mathcal{D}_{\text {ideal }}\left(1^{\lambda}\right)$ : Same as $\mathcal{D}_{\text {real }}\left(1^{\lambda}\right)$ except we replace each $c_{i}$ with a commitment to 0 and $\mathbf{t}$ with a uniformly random string: $\left(c_{i}, \sigma_{i}\right) \leftarrow$ BC.Commit $\left(\operatorname{crs}_{B C}, 0\right)$ for each $i \in[n]$ and $\mathrm{t} \stackrel{\mathrm{R}}{\leftarrow}\{0,1\}^{m B}$.

We now show that if $G_{n}$ is sub-exponentially secure (and the commitment scheme is statistically hiding), then $\mathcal{D}_{\text {real }}$ and $\mathcal{D}_{\text {ideal }}$ are also indistinguishable to a sub-exponential time algorithm.
Lemma 3.5. Suppose $G_{\lambda}$ is $\left(2^{\lambda^{\alpha}}, \varepsilon_{\text {PRG }}\right)$-secure for some constant $\alpha>0$ and negligible function $\varepsilon_{\text {PRG }}=$ negl $(\lambda)$ and that $\Pi_{B C}$ is statistically hiding in hiding mode. Suppose also that $\ell_{\mathrm{BARG}}=\operatorname{poly}(\lambda, m)$, and $B \geq \omega(\log \lambda)+\ell_{\mathrm{BARG}}$, $n \geq \max \left(\lambda, \ell_{\mathrm{BARG}}^{c}\right)$ for some constant $c>1 / \alpha$. Then, there exists a negligible function $\varepsilon_{i d e a l}=\operatorname{negl}(\lambda)$ such that for all subsets $I \subseteq[m], \mathcal{D}_{\text {real }}$ and $\mathcal{D}_{\text {ideal }}$ are $\left(2^{n^{\alpha}}, \varepsilon_{\text {ideal }}\right)$-indistinguishable

Proof. We start by defining a sequence of hybrid experiments:

- $\mathrm{Hyb}_{0}$ : This is the real distribution $\mathcal{D}_{\text {real }}$.
- $\mathrm{Hyb}_{i}$ : Same as $\mathrm{Hyb}_{0}$ except for all $j \leq i$, we now sample commitments $\left(c_{j}, \sigma_{j}\right) \leftarrow$ BC.Commit $\left(\mathrm{crs}_{\mathrm{BC}}, 0\right)$. The commitments for $j>i$ are sampled as in $\mathrm{Hyb}_{0}$.
- Hyb ${ }_{n+1}$ : Same as $\operatorname{Hyb}_{n}$ except $\mathrm{t} \stackrel{\mathrm{R}}{\leftarrow}\{0,1\}^{m B}$. This is the ideal distribution $\mathcal{D}_{\text {ideal }}$.

We now show that each adjacent pair of experiments are indistinguishable.
Claim 3.6. Suppose $\Pi_{B C}$ is statistically hiding in hiding mode. Then, there exists a negligible function $\varepsilon_{0}=\operatorname{negl}(\lambda)$ such that for all (possibly super-polynomial) functions $s_{0}=s_{0}(\lambda)$ and all $i \in[n]$, the distributions $\mathrm{Hyb}_{i-1}$ and $\mathrm{Hyb}_{i}$ are ( $s_{0}, \varepsilon_{0}$ )-indistinguishable.

Proof. Suppose there exists an adversary $\mathcal{A}$ of size $s_{0}$ that can distinguish $\mathrm{Hyb}_{i-1}$ and $\mathrm{Hyb}_{i}$ with non-negligible advantage $\delta$. We use $\mathcal{A}$ to construct an adversary $\mathcal{B}$ that breaks hiding of $\Pi_{\mathrm{BC}}$ as follows:

1. Algorithm $\mathcal{B}$ receives $\operatorname{crs}_{B C}$ from its challenger. It samples $\operatorname{crs}_{\text {BARG }} \leftarrow \operatorname{BARG} \cdot \operatorname{Setup}\left(1^{\lambda}, 1^{m B}, 1^{|C|}\right), s \stackrel{R}{\leftarrow}\{0,1\}^{n}$, and computes $\mathrm{t} \leftarrow G_{n}(\mathbf{s})$.
2. Then, for $j<i$, algorithm $\mathcal{B}$ computes $\left(c_{j}, \sigma_{j}\right) \leftarrow \operatorname{BC}$.Commit $\left(\operatorname{crs}_{\mathrm{BC}}, 0\right)$ and for $j>i$, it computes $\left(c_{j}, \sigma_{j}\right) \leftarrow$ BC.Commit $\left(\operatorname{crs}_{\mathrm{BC}}, s_{j}\right)$. Algorithm $\mathcal{B}$ submits $\left(s_{i}, 0\right)$ as its challenge and sets $c_{i}$ to be the challenger's response.
3. Algorithm $\mathcal{B}$ gives $\left(I, \operatorname{crs}_{\mathrm{BARG}}, \operatorname{crs}_{\mathrm{BC}}, c_{1}, \ldots, c_{n}, \mathrm{t}\right)$ to $\mathcal{A}$ and outputs whatever $\mathcal{A}$ outputs.

If $c_{i}$ is a commitment to $s_{i}$, then algorithm $\mathcal{B}$ perfectly simulates distribution $\mathrm{Hyb}_{i-1}$ and if $c_{i}$ is a commitment to 0 , then algorithm $\mathcal{B}$ perfectly simulates distribution $\mathrm{Hyb}_{i}$. Thus, algorithm $\mathcal{B}$ also succeeds with advantage $\delta$, and the claim follows.

Claim 3.7. Suppose $G_{\lambda}$ is $\left(2^{\lambda^{\alpha}}, \varepsilon_{\text {PRG }}\right)$-secure for some constant $\alpha>0$ and negligible function $\varepsilon_{\text {PRG }}=\varepsilon_{\text {PRG }}(\lambda)=\operatorname{neg}(\lambda)$. Then, $\mathrm{Hyb}_{n}$ and $\mathrm{Hyb}_{n+1}$ are $\left(2^{n^{\alpha}}, \varepsilon_{\mathrm{PRG}}^{\prime}\right)$-indistinguishable, where $\varepsilon_{\mathrm{PRG}}^{\prime}=\varepsilon_{\mathrm{PRG}}(n)$.

Proof. Suppose there exists an adversary $\mathcal{A}$ of size $s_{\mathcal{A}} \leq 2^{n^{\alpha}}$ that can distinguish $\operatorname{Hyb}_{n}\left(1^{\lambda}\right)$ and $\operatorname{Hyb}_{n+1}\left(1^{\lambda}\right)$ with advantage $\delta>\varepsilon_{\text {PRG }}^{\prime}$. We use $\mathcal{A}$ to construct an adversary $\mathcal{B}$ that breaks PRG security with seed length $n$ :

1. At the beginning of the experiment, algorithm $\mathcal{B}$ receives a challenge $\mathbf{t} \in\{0,1\}^{m B}$.
2. $\mathcal{B}$ samples $\operatorname{crs}_{\mathrm{BC}} \leftarrow \operatorname{BC} \cdot \operatorname{Setup}\left(1^{\lambda}\right.$, hide $), \operatorname{crs}_{\text {BARG }} \leftarrow \operatorname{BARG} . \operatorname{Setup}\left(1^{\lambda}, 1^{m B}, 1^{|C|}\right)$. For each $i \in[n]$, it computes $\left(c_{i}, \sigma_{i}\right) \leftarrow$ BC.Commit $\left(\mathrm{crs}_{\mathrm{BC}}, 0\right)$.
3. Algorithm $\mathcal{B}$ gives $\left(I, \operatorname{crs}_{B A R G}, \operatorname{crs}_{B C}, c_{1}, \ldots, c_{n}, t\right)$ to $\mathcal{A}$ and outputs whatever $\mathcal{A}$ outputs.

By construction, the size of algorithm $\mathcal{B}$ is $s_{\mathcal{A}}+\operatorname{poly}(\lambda, m, B,|C|) \leq s_{\mathcal{A}} \cdot \operatorname{poly}(n)$, where the inequality holds since $m, B,|C|$ are all polynomially-bounded (in both $\lambda$ and $n$ ). If $\mathbf{t}=G_{n}(\mathbf{s})$ for some $\mathbf{s} \stackrel{\mathrm{R}}{\leftarrow}\{0,1\}^{n}$, then $\mathcal{B}$ perfectly simulates $\operatorname{Hyb}_{n}\left(1^{\lambda}\right)$ for $\mathcal{A}$. Otherwise, if $\mathrm{t} \stackrel{\mathrm{R}}{\leftarrow}\{0,1\}^{m B}$, then $\mathcal{B}$ perfectly simulates $\operatorname{Hyb}_{n+1}\left(1^{\lambda}\right)$ for $\mathcal{A}$. Thus, algorithm $\mathcal{B}$ breaks security of $G_{n}$ with the same advantage $\delta>\varepsilon_{\mathrm{PRG}}^{\prime}=\varepsilon_{\mathrm{PRG}}(n)$.

By Claims 3.6 and 3.7, we can set $\varepsilon_{\text {ideal }}=n \cdot \varepsilon_{0}+\varepsilon_{\operatorname{PRG}}(n(\lambda, m))=\operatorname{negl}(\lambda)$. The latter equality follows since $n(\lambda, m) \geq \lambda$. The lemma now follows by a hybrid argument.

To complete the proof, we start by appealing to the Gentry-Wichs leakage simulation lemma (Lemma 2.2). Take any subset $I \subseteq[m]$. We start by defining the augmented distribution $\mathcal{D}_{\text {real }}^{*}=\mathcal{D}_{\text {real }}^{*}\left(1^{\lambda}\right)$ :

- Sample $\left(I, \operatorname{crs}_{\text {BARG }}, \operatorname{crs}_{B C}, c_{1}, \ldots, c_{n}, \mathbf{t}\right) \leftarrow \mathcal{D}_{\text {real }}\left(1^{\lambda}\right)$ according to the real distribution. Each commitment $c_{i}$ is computed as $\left(c_{i}, \sigma_{i}\right) \leftarrow \mathrm{BC}$.Commit $\left(\operatorname{crs}_{\mathrm{BC}}, s_{i}\right)$ where $\mathrm{s} \stackrel{\mathrm{R}}{\leftarrow}\{0,1\}^{n}$ and $\mathbf{t}=G_{n}(\mathrm{~s})$.
- Let st $=\left(n, \operatorname{crs}_{\mathrm{BARG}}, \mathrm{crs}_{\mathrm{BC}}, \mathbf{s}, c_{1}, \ldots, c_{n}, \sigma_{1}, \ldots, \sigma_{n}\right)$, and compute $\pi=\left(\pi_{\mathrm{BARG}},\left(c_{1}, \ldots, c_{n}\right),\left(\mathbf{t}_{i^{(1)}}, \ldots, \mathbf{t}_{i^{(L)}}\right)\right) \leftarrow$ Prove( $\mathrm{st}, I$ ).
- Output $\left(I, \mathrm{crs}_{\mathrm{BARG}}, \mathrm{crs}_{\mathrm{BC}}, c_{1}, \ldots, c_{n}, \mathbf{t}\right)$ and the auxiliary information aux $=\pi_{\mathrm{BARG}}$. By definition, $|\mathrm{aux}|=$ $\left|\pi_{\mathrm{BARG}}\right|=\ell_{\mathrm{BARG}}$.

By Lemma 3.5, the distributions $\mathcal{D}_{\text {real }}$ and $\mathcal{D}_{\text {ideal }}$ are $\left(s_{\text {ideal }}, \varepsilon_{\text {ideal }}\right)$-indistinguishable for $s_{\text {ideal }}=2^{n^{\alpha}}$ and a negligible function $\varepsilon_{\text {ideal }}=\operatorname{neg}(\lambda)$. Without loss of generality, we can assume that $\varepsilon_{\text {ideal }} \geq 2^{-n^{\alpha / 2}}$ (e.g., we can set $\varepsilon_{\text {ideal }}=$ $\max \left(\varepsilon_{\text {ideal }}^{\prime}, 2^{-n^{\alpha / 2}}\right)$, where $\varepsilon_{\text {ideal }}^{\prime}$ is the negligible function from Lemma 3.5). By Lemma 2.2, there exists an augmented distribution $\mathcal{D}_{\text {ideal }}^{*}=\mathcal{D}_{\text {ideal }}^{*}\left(1^{\lambda}\right)$ over tuples $\left(\left(I, \operatorname{crs}_{\text {BARG }}, \operatorname{crs}_{B C}, c_{1}, \ldots, c_{n}, \mathbf{t}\right)\right.$, aux $)$ where $\left(I, \operatorname{crs}_{\text {BARG }}, \operatorname{crs}_{B C}, c_{1}, \ldots, c_{n}, \mathbf{t}\right) \leftarrow$ $\mathcal{D}_{\text {ideal }}\left(1^{\lambda}\right)$ and aux $x^{\prime} \in\{0,1\}^{\ell_{\text {BARG }}}$. Moreover, the distributions $\mathcal{D}_{\text {real }}^{*}$ and $\mathcal{D}_{\text {ideal }}^{*}$ are $\left(s_{\text {aug }}, \varepsilon_{\text {aug }}\right)$-indistinguishable where

$$
s_{\text {aug }}=s_{\text {ideal }} \cdot \operatorname{poly}\left(\varepsilon_{\text {ideal }} / 2^{\ell_{\text {BARG }}}\right)=2^{n^{\alpha}} \cdot \operatorname{poly}\left(\varepsilon_{\text {ideal }} / 2^{\ell_{\text {BARG }}}\right)
$$

and $\varepsilon_{\text {aug }}=2 \cdot \varepsilon_{\text {ideal }}=\operatorname{negl}(\lambda)$. Since $\varepsilon_{\text {ideal }} \geq 2^{-n^{\alpha / 2}}$ and $n \geq \ell_{\text {BARG }}^{c}$ for some constant $c>1 / \alpha$, this means that $s_{\text {aug }}=2^{\Omega\left(n^{\alpha}\right)}$. We summarize this in the following claim:

Claim 3.8. Under the same conditions as in the statement of Lemma 3.5, the distributions $\mathcal{D}_{\text {real }}^{*}$ and $\mathcal{D}_{\text {ideal }}^{*}$ are ( $s_{\text {aug }}, \varepsilon_{\text {aug }}$ )indistinguishable where $s_{\text {aug }}=2^{\Omega\left(n^{\alpha}\right)}$ and $\varepsilon_{\text {aug }}=\operatorname{negl}(\lambda)$.

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

- $\mathrm{Hyb}_{0}$ : This is the real distribution where the challenger samples the bits $\mathbf{r}$ and the proof $\pi$ as in the real scheme:
- The challenger first samples crs $\leftarrow \operatorname{Setup}\left(1^{\lambda}, 1^{m}\right)$. In particular, $\operatorname{crs}=\left(n, \operatorname{crs}_{\mathrm{BARG}}, \operatorname{crs}_{\mathrm{BC}}, \mathbf{v}_{1}, \ldots, \mathbf{v}_{m}\right)$, where $\left(\operatorname{crs}_{B C}, \operatorname{td}\right) \leftarrow \operatorname{BC} . \operatorname{Setup}\left(1^{\lambda}\right.$, bind $), \operatorname{crs}_{\text {BARG }} \leftarrow \operatorname{BARG} . \operatorname{Setup}\left(1^{\lambda}, 1^{m B}, 1^{|C|}\right), C$ is the circuit that computes the NP relation $\mathcal{R}\left[n, \mathrm{crs}_{\mathrm{BC}}\right]$ from Fig. 1 , and $\mathbf{v}_{1}, \ldots, \mathbf{v}_{m} \stackrel{R}{\leftarrow}\{0,1\}^{B}$.
- Next, the challenger samples the bits $\mathbf{r}$ by running $(\mathbf{r}, \mathrm{st}) \leftarrow \operatorname{GenBits}(\mathrm{crs})$ and a proof by computing $\pi \leftarrow \operatorname{Prove}(\mathrm{st}, I)$. In particular, the challenger first samples a seed $\mathbf{s} \stackrel{\mathrm{R}}{\leftarrow}\{0,1\}^{n}$ and computes $\mathbf{t} \leftarrow G_{n}(\mathbf{s})$. It splits $\mathbf{t}=\mathbf{t}_{1}\left\|\mathbf{t}_{2}\right\| \cdots \| \mathbf{t}_{m}$ into blocks where each $\mathbf{t}_{i} \in\{0,1\}^{B}$ for each $i \in[m]$. For each $i \in[m]$, the challenger computes $r_{i} \leftarrow \mathbf{v}_{i}^{\top} \mathbf{t}_{i}$ and sets $\mathbf{r}=r_{1}\|\cdots\| r_{m} \in\{0,1\}^{m}$.
- To construct the proof $\pi$, the challenger computes the commitments $\left(c_{i}, \sigma_{i}\right) \leftarrow$ BC.Commit $\left(\operatorname{crs}_{\mathrm{BC}}, s_{i}\right)$ for each $i \in[n]$. It then parses $I=\left\{i^{(1)}, \ldots, i^{(L)}\right\}$, where the indices $i^{(1)}, \ldots, i^{(L)} \in[m]$ are in sorted order. The challenger then constructs the statement $\left(x_{1,1}, \ldots, x_{1, B}, \ldots, x_{L, 1}, \ldots, x_{L, B}\right)$ and witness $\left(w_{1,1}, \ldots, w_{1, B}, \ldots, w_{L, 1}, \ldots, w_{L, B}\right)$ according to Eq. (3.1) and Eq. (3.2). It constructs the BARG proof as in Prove:

$$
\pi_{\mathrm{BARG}} \leftarrow \text { BARG.Prove }\left(\operatorname{crs}_{\mathrm{BARG}}, C,\left(x_{1,1}, \ldots, x_{1, B}, \ldots, x_{L, 1}, \ldots, x_{L, B}\right),\left(w_{1,1}, \ldots, w_{1, B}, \ldots, w_{L, 1}, \ldots, w_{L, B}\right)\right)
$$

$$
\text { and sets } \pi=\left(\pi_{\mathrm{BARG}},\left(c_{1}, \ldots, c_{n}\right),\left(\mathbf{t}_{i^{(1)}}, \ldots, \mathbf{t}_{i^{(L)}}\right)\right)
$$

- The challenger gives ( $\left.\operatorname{crs}, I, \mathbf{r}_{I}, \pi, \mathbf{r}_{\bar{I}}\right)$ to $\mathcal{A}$. Algorithm $\mathcal{A}$ then outputs a bit $b \in\{0,1\}$ which is the output of the experiment.
- $\mathrm{Hyb}_{1}$ : Same as $\mathrm{Hyb}_{0}$ except the challenger now samples the commitment $\mathrm{CRS}\left(\mathrm{crs}_{\mathrm{BC}}, \mathrm{td}\right) \leftarrow \mathrm{BC}$. $\operatorname{Setup}\left(1^{\lambda}\right.$, hide $)$. In this experiment, the distribution of $\left(\left(\operatorname{crs}_{\mathrm{BARG}}, \mathrm{crs}_{\mathrm{BC}}, c_{1}, \ldots, c_{n}, \mathrm{t}\right), \pi_{\mathrm{BARG}}\right)$ is distributed according to $\mathcal{D}_{\text {real }}^{*}$.
- $\mathrm{Hyb}_{2}$ : Same as $\mathrm{Hyb}_{1}$ except the challenger samples components $\left(\left(\operatorname{crs}_{\mathrm{BARG}}, \mathrm{crs}_{\mathrm{BC}}, c_{1}, \ldots, c_{n}, \mathrm{t}\right), \pi_{\mathrm{BARG}}\right) \leftarrow \mathcal{D}_{\text {ideal }}^{*}$. Specifically, the experiment now proceeds as follows:
- The challenger samples $\left(\left(\operatorname{crs}_{\mathrm{BARG}}, \operatorname{crs}_{\mathrm{BC}}, c_{1}, \ldots, c_{n}, \mathrm{t}\right), \pi_{\mathrm{BARG}}\right) \leftarrow \mathcal{D}_{\text {ideal }}^{*}, \mathbf{v}_{1}, \ldots, \mathbf{v}_{m} \stackrel{\mathrm{R}}{\leftarrow}\{0,1\}^{B}$, and sets crs $=\left(n\right.$, crs $_{\text {BARG }}$, crs $\left._{\mathrm{BC}}, \mathbf{v}_{1}, \ldots, \mathbf{v}_{m}\right)$.
- Next, the challenger splits $\mathbf{t}=\mathbf{t}_{1}\left\|\mathbf{t}_{2}\right\| \cdots \| \mathbf{t}_{m}$ into blocks where each $\mathbf{t}_{i} \in\{0,1\}^{B}$ for each $i \in[m]$. For each $i \in[m]$, the challenger computes $r_{i} \leftarrow \mathbf{v}_{i}^{\top} \mathbf{t}_{i}$ and sets $\mathbf{r}=r_{1}\|\cdots\| r_{m} \in\{0,1\}^{m}$.
- The challenger sets the proof $\pi=\left(\pi_{\mathrm{BARG}},\left(c_{1}, \ldots, c_{n}\right),\left(\mathbf{t}_{i^{(1)}}, \ldots, \mathbf{t}_{i^{(L)}}\right)\right)$ and gives (crs, $\left.I, \mathbf{r}_{I}, \pi, \mathbf{r}_{\bar{I}}\right)$ to $\mathcal{A}$.
- $\mathrm{Hyb}_{3}$ : Same as $\mathrm{Hyb}_{2}$ except the challenger samples $\mathrm{r}_{\bar{I}} \stackrel{\mathrm{R}}{\leftarrow}\{0,1\}^{|\bar{I}|}$.
- $\mathrm{Hyb}_{4}$ : Same as $\mathrm{Hyb}_{3}$ except the challenger samples $\left(\left(\operatorname{crs}_{\mathrm{BARG}}, \mathrm{crs}_{\mathrm{BC}}, c_{1}, \ldots, c_{n}, \mathrm{t}\right), \pi_{\mathrm{BARG}}\right) \leftarrow \mathcal{D}_{\text {real }}^{*}$.
- $\mathrm{Hyb}_{5}$ : Same as $\mathrm{Hyb}_{4}$ except the challenger samples $\operatorname{crs}_{B C} \leftarrow B C \cdot \operatorname{Setup}\left(1^{\lambda}\right.$, bind). Note that this coincides with the ideal distribution.

Lemma 3.9. Suppose $\Pi_{B C}$ satisfies mode indistinguishability. Then, $\mathrm{Hyb}_{0}$ and $\mathrm{Hyb}_{1}$ are computationally indistinguishable.

Proof. Suppose there is an adversary $\mathcal{A}$ of size $s_{0}=\operatorname{poly}(\lambda)$ that distinguishes the outputs of $\mathrm{Hyb}_{0}$ and $\mathrm{Hyb}_{1}$ with non-negligible probability $\delta$. We use $\mathcal{A}$ to construct an adversary $\mathcal{B}$ that breaks mode indistinguishability:

1. At the beginning of the game, algorithm $\mathcal{B}$ receives the security parameter $1^{\lambda}$ and a common reference string $\mathrm{crs}_{B C}$ from the challenger.
2. $\mathcal{B}$ samples $\operatorname{crs}_{\text {BARG }} \leftarrow \operatorname{BARG} . \operatorname{Setup}\left(1^{\lambda}, 1^{m B}, 1^{|C|}\right)$ and $\mathbf{v}_{1}, \ldots, \mathbf{v}_{m} \stackrel{\mathrm{R}}{\leftarrow}\{0,1\}^{B}$. It constructs the common reference string crs $=\left(n, \mathrm{crs}_{\mathrm{BARG}}, \mathrm{crs}_{\mathrm{BC}}, \mathbf{v}_{1}, \ldots, \mathbf{v}_{m}\right)$.
3. Algorithm $\mathcal{B}$ computes $(\mathbf{r}, \mathrm{st}) \leftarrow \operatorname{GenBits}(\mathrm{crs})$ and $\pi \leftarrow \operatorname{Prove}(\mathrm{st}, I)$.
4. Algorithm $\mathcal{B}$ gives ( $\operatorname{crs}, I, \mathbf{r}_{I}, \pi, \mathbf{r}_{\bar{I}}$ ) to $\mathcal{A}$ and outputs whatever $\mathcal{A}$ outputs.

By construction, algorithm $\mathcal{B}$ has size $s_{0}+\operatorname{poly}(\lambda, m, B,|C|) \leq s_{0} \cdot \operatorname{poly}(\lambda)$ which holds due to $m, B,|C|$ all being $\operatorname{poly}(\lambda)$. When $\operatorname{crs}_{B C}$ is sampled in binding mode, then algorithm $\mathcal{B}$ perfectly simulates $\mathrm{Hyb}_{0}$ for $\mathcal{A}$. Alternatively, if $\operatorname{crs}_{B C}$ is sampled in hiding mode, then algorithm $\mathcal{B}$ perfectly simulates $\mathrm{Hyb}_{1}$ for $\mathcal{A}$. Critically, neither the GenBits nor the Prove algorithms require knowledge of the trapdoor td for the bit commitment scheme. Thus, algorithm $\mathcal{B}$ succeeds with the same advantage $\delta$.

Lemma 3.10. Under the same conditions as in the statement of Claim 3.8, $\mathrm{Hyb}_{1}$ and $\mathrm{Hyb}_{2}$ are ( $s_{\text {aug }}, \varepsilon_{\mathrm{aug}}$ )-indistinguishable for $s_{\text {aug }}=2^{\Omega\left(n^{\alpha}\right)}$ and $\varepsilon_{\text {aug }}=\operatorname{negl}(\lambda)$.

Proof. Suppose there is an adversary $\mathcal{A}$ with size $s_{\text {aug }}$ that distinguishes $\mathrm{Hyb}_{1}$ and $\mathrm{Hyb}_{2}$ with advantage $\delta>\varepsilon_{\text {aug. We }}$. We construct algorithm $\mathcal{B}$ that distinguishes the distributions $\mathcal{D}_{\text {real }}^{*}\left(1^{\lambda}\right)$ and $\mathcal{D}_{\text {ideal }}^{*}\left(1^{\lambda}\right)$ as follows:

1. Algorithm $\mathcal{B}$ receives $\left(I, \operatorname{crs}_{\mathrm{BARG}}, \operatorname{crs}_{\mathrm{BC}}, c_{1}, \ldots, c_{n}, \mathbf{t}, \pi_{\mathrm{BARG}}\right.$ ) from the challenger. It parses $\mathbf{t}=\mathbf{t}_{1}\|\cdots\| \mathbf{t}_{m} \in$ $\{0,1\}^{m B}$ where each $\mathbf{t}_{i} \in\{0,1\}^{B}$. In addition, algorithm $\mathcal{B}$ samples $\mathbf{v}_{1}, \ldots, \mathbf{v}_{m} \stackrel{R}{R}_{\leftarrow}\{0,1\}^{B}$.
2. Algorithm $\mathcal{B}$ computes $n=n(\lambda, m)$ and sets crs $=\left(n, \operatorname{crs}_{B A R G}, \operatorname{crs}_{B C}, \mathbf{v}_{1}, \ldots, \mathbf{v}_{m}\right)$. For each $i \in[m]$, it computes $r_{i} \leftarrow \mathbf{v}_{i}^{\top} \mathbf{t}_{i}$ and sets $\mathbf{r}=r_{1}\|\cdots\| r_{m}$. Finally, it sets $\pi=\left(\pi_{\mathrm{BARG}},\left(c_{1}, \ldots, c_{n}\right),\left(\mathbf{t}_{i^{(1)}}, \ldots, \mathbf{t}_{i^{(L)}}\right)\right)$, where $I=\left\{i^{(1)}, \ldots, i^{(L)}\right\}$.
3. Algorithm $\mathcal{B}$ gives (crs, $I, \mathbf{r}_{I}, \pi, \mathbf{r}_{\bar{I}}$ ) to $\mathcal{A}$ and outputs whatever $\mathcal{A}$ outputs.

Since algorithm $\mathcal{A}$ has size $s_{\text {aug }}$, the size of algorithm $\mathcal{B}$ is bounded by $s_{\text {aug }}+\operatorname{poly}(\lambda, m, B) \leq s_{\text {aug }} \cdot \operatorname{poly}(\lambda)$ since $m$ and $B$ are both polynomially-bounded. By construction, when the challenge is sampled from the real distribution $\mathcal{D}_{\text {real }}^{*}$, algorithm $\mathcal{B}$ perfectly simulates the distribution in $\mathrm{Hyb}_{1}$. Alternatively, if the challenge is sampled from the ideal distribution $\mathcal{D}_{\text {ideal }}^{*}$, algorithm $\mathcal{B}$ perfectly simulates the distribution in $\mathrm{Hyb}_{2}$. Correspondingly, algorithm $\mathcal{B}$ is able to distinguish $\mathcal{D}_{\text {real }}^{*}\left(1^{\lambda}\right)$ and $\mathcal{D}_{\text {ideal }}^{*}\left(1^{\lambda}\right)$ with advantage $\delta>\varepsilon_{\text {aug }}$ which contradicts Claim 3.8.
Lemma 3.11. Suppose $B \geq \omega(\log \lambda)+\ell_{\text {BARG }}$. Then, $\mathrm{Hyb}_{2}$ and $\mathrm{Hyb}_{3}$ are statistically indistinguishable.
Proof. Let $\bar{I}=\left\{i^{(1)}, i^{(2)}, \ldots, i^{(m-L)}\right\} \subseteq[m]$. We define a sequence of intermediate experiments $\mathrm{Hyb}_{2, j}$ for each $j \in\{0, \ldots, m-L\}$ as follows:

- $\mathrm{Hyb}_{2,0}$ : Same as $\mathrm{Hyb}_{2}$. In particular, the challenger samples $\left(\left(I, \operatorname{crs}_{\mathrm{BARG}}, \mathrm{crs}_{\mathrm{BC}}, c_{1}, \ldots, c_{n}, \mathbf{t}\right), \pi_{\mathrm{BARG}}\right) \leftarrow \mathcal{D}_{\text {ideal }}^{*}$, $\mathbf{v}_{1}, \ldots, \mathbf{v}_{m} \stackrel{\mathrm{R}}{\leftarrow}\{0,1\}^{B}$ and sets crs $=\left(n, \operatorname{crs}_{\text {BARG }}\right.$, crs $\left._{\mathrm{BC}}, \mathbf{v}_{1}, \ldots, \mathbf{v}_{m}\right)$. It parses $\mathbf{t}=\mathbf{t}_{1}\|\cdots\| \mathbf{t}_{m}$ where $\mathbf{t}_{i} \in\{0,1\}^{B}$ and computes $\mathbf{r} \leftarrow\left(\mathbf{v}_{1}^{\top} \mathbf{t}_{1}\|\cdots\| \mathbf{v}_{m}^{\top} \mathbf{t}_{m}\right)$. Finally, it sets $\pi=\left(\pi_{\mathrm{BARG}},\left(c_{1}, \ldots, c_{n}\right),\left(\mathbf{t}_{i^{(1)}}, \ldots, \mathbf{t}_{i^{(L)}}\right)\right)$ and gives (crs, $I, \mathbf{r}_{I}, \pi, \mathbf{r}_{\bar{I}}$ ) to the adversary.
- $\mathrm{Hyb}_{2, j}$ : Same as $\mathrm{Hyb}_{2, j-1}$ except $r_{i(j)} \stackrel{\mathrm{R}}{\leftarrow}\{0,1\}$. Note that $\mathrm{Hyb}_{2, m-L}$ is identical to $\mathrm{Hyb}_{3}$.

We now appeal to the leftover hash lemma to show that for all $j \in[m-L]$, the statistical distance between $\operatorname{Hyb}_{2, j-1}\left(1^{\lambda}\right)$ and $\operatorname{Hyb}_{2, j}\left(1^{\lambda}\right)$ is negligible.

Claim 3.12. Suppose $B \geq \omega(\log \lambda)+\ell_{\text {BARG }}$. Then, for all $j \in[m-L]$, the statistical distance between $\operatorname{Hyb}_{2, j-1}\left(1^{\lambda}\right)$ and $\operatorname{Hyb}_{2, j}\left(1^{\lambda}\right)$ is negligible.
Proof. The only difference between the two distributions is that in $\mathrm{Hyb}_{2, j-1}$, the challenger samples $r_{i^{(j)}} \leftarrow \mathbf{v}_{i^{(j)}}^{\top} \mathbf{t}_{i^{(j)}}$, whereas in $\mathrm{Hyb}_{2, j}$, the challenger samples $r_{i(j)} \stackrel{\mathrm{R}}{\leftarrow}\{0,1\}$. First, define the random variable $Z$ to be

$$
Z=\left(n, \operatorname{crs}_{\mathrm{BARG}}, \operatorname{crs}_{\mathrm{BC}},\left\{\mathbf{v}_{i}\right\}_{i \neq i(j)}, I, \mathbf{r}_{I},\left(\pi_{\mathrm{BARG}},\left(c_{1}, \ldots, c_{n}\right),\left\{\mathbf{t}_{i}\right\}_{i \in I}\right), \mathbf{r}_{\bar{I} \backslash\left\{i^{(j)}\right\}}\right) .
$$

Observe that the adversary's view in the two experiments then consists of the tuple $\left(r_{i^{(j)}}, \mathbf{v}_{i^{(j)}}, Z\right)$. In both $\operatorname{Hyb}_{2, j-1}$ and $\mathrm{Hyb}_{2, j}$, the challenger samples $\mathbf{t} \stackrel{\mathrm{R}}{\leftarrow}\{0,1\}^{m B}$. By construction, $\mathbf{t}_{i(j)}$ is independent of all of the components in $Z$ other than $\pi_{\text {BARG }}$. In conjunction with Lemma 2.1, we can now write

$$
\mathbf{H}_{\infty}\left(\mathbf{t}_{i}(j) \mid Z\right)=\mathbf{H}_{\infty}\left(\mathbf{t}_{i^{(j)}} \mid \pi_{\mathrm{BARG}}\right) \geq \mathbf{H}_{\infty}\left(\mathbf{t}_{i^{(j)}}\right)-\left|\pi_{\mathrm{BARG}}\right|=B-\ell_{\mathrm{BARG}} \geq \omega(\log \lambda)
$$

since $B \geq \omega(\log \lambda)+\ell_{\text {BARG }}$. Then, by the (generalized) leftover hash lemma (Corollary 2.4), we can conclude that the statistical distance between the distributions

$$
\left(\mathbf{v}_{i}^{\top}{ }^{\top} \mathbf{t}_{i^{(j)}}, \mathbf{v}_{i^{(j)}}, Z\right) \quad \text { and } \quad\left(r_{i(j)}, \mathbf{v}_{i}(j), Z\right)
$$

where $\mathbf{v}_{i^{(j)}} \stackrel{\mathrm{R}}{\leftarrow}\{0,1\}^{B}$ and $r_{i^{(j)}} \stackrel{\mathrm{R}}{\leftarrow}\{0,1\}$ is at most $2^{-(\omega(\log \lambda)-1) / 2}=\operatorname{negl}(\lambda)$. Since the statistical distance between the two experiments is negligible, the claim holds.

The lemma now follows from Claim 3.12 and a standard hybrid argument (since $m=\operatorname{poly}(\lambda)$ ).
Lemma 3.13. Under the same conditions as in the statement of Claim 3.8, $\mathrm{Hyb}_{3}$ and $\mathrm{Hyb}_{4}$ are ( $s_{\text {aug }}, \varepsilon_{\mathrm{aug}}$ )-indistinguishable for $s_{\text {aug }}=2^{\Omega\left(n^{\alpha}\right)}$ and $\varepsilon_{\text {aug }}=\operatorname{negl}(\lambda)$.

Proof. Follows by an analogous argument as the proof of Lemma 3.10.
Lemma 3.14. Suppose $\Pi_{B C}$ satisfies mode indistinguishability. Then, $\mathrm{Hyb}_{4}$ and $\mathrm{Hyb}_{5}$ are computationally indistinguishable.

Proof. Follows by an analogous argument as the proof of Lemma 3.9.
Combining Lemmas 3.9 to 3.11, 3.13 and 3.14 yields the theorem.

Parameter selection. We now describe one candidate approach for instantiating the parameters in Construction 3.1:
Corollary 3.15 (Hidden-Bits Generator from Batch Arguments). Let $k=k(\lambda)$ be a locality parameter and suppose that $G_{\lambda}:\{0,1\}^{\lambda} \rightarrow\{0,1\}^{\ell(\lambda)}$ is a $k$-local PRG. Suppose $\Pi_{B A R G}$ is a non-interactive batch argument satisfying semi-adaptive soundness, $\Pi_{\mathrm{BC}}$ is a dual-mode commitment scheme, and that these underlying schemes satisfy the following conditions:

- PRG parameters: Suppose there exists a constant $\alpha \in(0,1)$ and a negligible function $\varepsilon_{\text {PRG }}=$ negl $(\lambda)$ such that $G_{\lambda}$ is $\left(2^{\lambda^{\alpha}}, \varepsilon_{\text {PRG }}\right)$ secure. Moreover, suppose there exists a constant $\delta_{1} \in(0,1)$ such that $k \leq \delta_{1} \log \lambda$ and a constant $s>1$ such that $\ell(\lambda) \geq \lambda^{s}$. In words, we assume that $G_{\lambda}$ has super-linear stretch, logarithmic locality, and sub-exponential security.
- BARG succinctness: Suppose there exists constants $d>0, \delta_{2} \in(0,1)$ and a polynomial $q=q(\lambda)$ such that the proof length $\ell_{\mathrm{BARG}}=\ell_{\mathrm{BARG}}(\lambda, T, s)$ for $\Pi_{\mathrm{BARG}}$ satisfies

$$
\ell_{\mathrm{BARG}}(\lambda, T, s) \leq s^{d} \cdot T^{\delta_{2}} \cdot q(\lambda),
$$

where $T$ denotes the number of instances ands denotes a bound on the size of the circuit.

- Block size: Suppose $B=\lambda+\ell_{\text {BARG }}$.

Let $\delta_{1}^{\prime}=\left(d \cdot \delta_{1}+\varepsilon\right) /\left(1-\delta_{2}\right)$ for an arbitrarily small constant $\varepsilon>0, \delta_{2}^{\prime}=\delta_{2} /\left(1-\delta_{2}\right)$, and $q^{\prime}(\lambda)=q(\lambda)^{1 /\left(1-\delta_{2}\right)}$. Suppose moreover that the parameters satisfy the following properties:

- Hardness parameter: $\alpha>\delta_{1}^{\prime}+\delta_{2}^{\prime}$.
- Seed length: $n=\max \left(\lambda,\left(m^{\delta_{2}^{\prime}} \cdot q^{\prime}(\lambda) \cdot O\left(\lambda^{\delta_{2}^{\prime}}\right)\right)^{1 /\left(\alpha-\delta_{1}^{\prime}-\varepsilon^{\prime}\right)}\right)$ for an arbitrary constant $0<\varepsilon^{\prime}<\alpha-\delta_{1}^{\prime}-\delta_{2}^{\prime}$.
- Stretch: $s \geq\left(\alpha-\delta_{1}^{\prime}-\varepsilon^{\prime}\right)\left(1+\delta_{2}^{\prime}\right) / \delta_{2}^{\prime}+\delta_{1}^{\prime}$.

Then Construction 3.1 is a hidden-bits generator with subset-dependent proofs.
Proof. Take any input length $m$. Let crs $=\left(n, \operatorname{crs}_{\text {BARG }}, \operatorname{crs}_{B C}, \mathbf{v}_{1}, \ldots, \mathbf{v}_{m}\right) \leftarrow \operatorname{Setup}\left(1^{\lambda}, 1^{m}\right)$. We first bound the size of the circuit $C$ that computes the relation $\mathcal{R}\left[n, \mathrm{crs}_{B C}\right]$ :

- By construction, $\left|\operatorname{crs}_{B C}\right|=\operatorname{poly}(\lambda)$. Correspondingly, the size of the circuit computing BC.Verify is poly $(\lambda)$.
- Next, $G_{n}^{(i)}$ is a function on $k$-bit inputs, so it can be computed by a circuit of size $2^{k} \cdot \operatorname{poly}(k)$. Since $k \leq \delta_{1} \log n$, we can bound

$$
|C| \leq n^{\delta_{1}} \cdot \operatorname{poly}(\log n)=O\left(n^{\delta_{1}+\varepsilon / d}\right)
$$

For this choice of parameters, the length $\ell_{\text {BARG }}$ of the BARG proof satisfies

$$
\ell_{\mathrm{BARG}}=\ell_{\mathrm{BARG}}(\lambda, m B,|C|) \leq|C|^{d} \cdot(m B)^{\delta_{2}} \cdot q(\lambda)=n^{d \delta_{1}+\varepsilon} \cdot m^{\delta_{2}} \cdot \ell_{\mathrm{BARG}}^{\delta_{2}} \cdot q(\lambda) \cdot O\left(\lambda^{\delta_{2}}\right) .
$$

Equivalently, this means

$$
\ell_{\mathrm{BARG}} \leq\left(n^{d \delta_{1}+\varepsilon} \cdot m^{\delta_{2}}\right)^{1 /\left(1-\delta_{2}\right)} \cdot q(\lambda)^{1 /\left(1-\delta_{2}\right)} \cdot O\left(\lambda^{\delta_{2} /\left(1-\delta_{2}\right)}\right)=n^{\delta_{1}^{\prime}} \cdot m^{\delta_{2}^{\prime}} \cdot q^{\prime}(\lambda) \cdot O\left(\lambda^{\delta_{2}^{\prime}}\right),
$$

We now consider the requirements of Theorem 3.3, Theorem 3.4 and the requirement on the PRG stretch:

- Theorem 3.4 requires that $n \geq \max \left(\lambda, \ell_{\mathrm{BARG}}^{c}\right)$ for some constant $c>1 / \alpha$. Let $c=1 /\left(\alpha-\varepsilon^{\prime}\right)>1 / \alpha$. By assumption, we now have

$$
n^{\alpha-\delta_{1}^{\prime}-\varepsilon^{\prime}} \geq m^{\delta_{2}^{\prime}} \cdot q^{\prime}(\lambda) \cdot O\left(\lambda^{\delta_{2}^{\prime}}\right)
$$

In particular, this means that

$$
n^{\alpha-\varepsilon^{\prime}} \geq n^{\delta_{1}^{\prime}} \cdot m^{\delta_{2}^{\prime}} \cdot q^{\prime}(\lambda) \cdot O\left(\lambda^{\delta_{2}^{\prime}}\right) \geq \ell_{\mathrm{BARG}}
$$

Correspondingly, we have $\left(n^{\alpha-\varepsilon^{\prime}}\right)^{c}=n \geq \ell_{\mathrm{BARG}}^{c}$, as required.

- Theorem 3.3 requires that $n \leq m^{\delta} \cdot \operatorname{poly}(\lambda)$ for some (universal) constant $\delta \in(0,1)$. Since $q=\operatorname{poly}(\lambda)$ and $\alpha, \delta_{1}^{\prime}, \varepsilon^{\prime}$ are all constants, we currently have that $n \leq m^{\delta_{2}^{\prime} /\left(\alpha-\delta_{1}^{\prime}-\varepsilon^{\prime}\right)} \cdot \operatorname{poly}(\lambda)$. By construction, we have that $0<\alpha-\delta_{1}^{\prime}-\delta_{2}^{\prime}-\varepsilon^{\prime}$, so $\delta_{2}^{\prime}<\alpha-\delta_{1}^{\prime}-\varepsilon$. Thus, setting $\delta=\delta_{2}^{\prime} /\left(\alpha-\delta_{1}^{\prime}-\varepsilon^{\prime}\right)<1$ satisfies the requirement.
- Finally, we require that $\ell(n) \geq m B$, or equivalently, $n^{s} \geq m B$. By construction,

$$
\begin{aligned}
n^{s}=n^{\delta_{1}^{\prime}} n^{s-\delta_{1}^{\prime}} & \geq n^{\delta_{1}^{\prime}} \cdot\left(m^{\delta_{2}^{\prime}} \cdot q^{\prime}(\lambda) \cdot O\left(\lambda^{\delta_{2}^{\prime}}\right)\right)^{\left(s-\delta_{1}^{\prime}\right) /\left(\alpha-\delta_{1}^{\prime}-\varepsilon^{\prime}\right)} \\
& \geq n^{\delta_{1}^{\prime}} \cdot\left(m^{\delta_{2}^{\prime}} \cdot q^{\prime}(\lambda) \cdot O\left(\lambda^{\delta_{2}^{\prime}}\right)\right)^{\left(1+\delta_{2}^{\prime}\right) / \delta_{2}^{\prime}} \\
& \geq n^{\delta_{1}^{\prime}} \cdot m^{1+\delta_{2}^{\prime}} \cdot\left(q^{\prime}(\lambda)\right)^{\left(1+\delta_{2}^{\prime}\right) / \delta_{2}^{\prime}} \cdot O\left(\lambda^{1+\delta_{2}^{\prime}}\right) .
\end{aligned}
$$

Finally, we have

$$
m B=m \lambda+m \ell_{\mathrm{BARG}} \leq m \lambda+n^{\delta_{1}^{\prime}} \cdot m^{1+\delta_{2}^{\prime}} \cdot q^{\prime}(\lambda) \cdot O\left(\lambda^{\delta_{2}^{\prime}}\right) \leq n^{s},
$$

as required.

Candidate instantiations. For illustrative purposes, we now describe some instantiations of Corollary 3.15.

- Suppose we instantiate Construction 3.1 and Corollary 3.15 with a batch argument where the proof size scales polylogarithmically with the number of instances:

$$
\ell_{\mathrm{BARG}}(\lambda, T, s) \leq s^{d} \cdot \operatorname{polylog}(T) \cdot q(\lambda)
$$

for some constant $d>0$. This is satisfied by most existing BARG constructions [CJJ21b, WW22, DGKV22, KLVW23, $\mathrm{CGJ}^{+}$22]. In this case, the constant $\delta_{2}$ in Corollary 3.15 can be made arbitrarily small. Then we can instantiate Corollary 3.15 with any $k$-local PRG that is secure against $2^{\lambda^{\alpha}}$-size adversaries with locality $k \leq \delta_{1} \log \lambda$ and stretch $s>1+d \delta_{1}$, provided that $\alpha / \delta_{1}>d$. For example, we can rely on sub-exponential hardness of Goldreich's local PRG [Gol00] with logarithmic locality.

- We can also instantiate Construction 3.1 and Corollary 3.15 with a "mildly-succinct" batch argument where the BARG proof size scales polynomially with the number of instances: ${ }^{7}$

$$
\ell_{\mathrm{BARG}}(\lambda, T, s) \leq s^{d} \cdot T^{\delta_{2}} \cdot q(\lambda)
$$

for constants $\delta_{2} \in(0,1 / 2)$ and $d>0$. In this case, we can instantiate Corollary 3.15 with a $k$-local PRG that is secure against $2^{\lambda^{\alpha}}$-size adversaries with locality $k \leq \delta_{1} \log \lambda$ and stretch $s>1+\delta_{1}^{\prime}+\delta_{2}^{\prime}$, as long as $\delta_{1}^{\prime}<\alpha-\delta_{2}^{\prime}$ (for $\delta_{1}^{\prime}, \delta_{2}^{\prime}$ as in Corollary 3.15). In particular, we can still rely on sub-exponential hardness of Goldreich's PRG with logarithmic locality, but the sub-exponential hardness parameter $\alpha$ increases as $\delta_{2}$ increases.

[^5]NIZK from batch arguments. Combining Theorem 2.11 and Corollary 3.15, we now obtain a NIZK for NP from a batch argument for NP:

Corollary 3.16 (NIZK from Batch Arguments). Suppose there exists a semi-adaptively-sound BARG, a dual-mode commitment scheme, and a sub-exponentially secure PRG with super-linear stretch and locality at most $k=c \log n$ with $c<1$ and n-bit inputs. Then there exists a computational NIZK argument for NP.

Remark 3.17 (Using Non-Local PRGs). We note that a local PRG is not strictly necessary for Construction 3.1. It is sufficient to construct a PRG where each output bit of the PRG can be verified by a circuit of size $n^{\delta}$ where $n$ is the seed length and $\delta<1$ is a constant. Any PRG with this local verification property suffices for our main transformation.

## Acknowledgments

We thank Brent Waters for helpful feedback on this work. D. J. Wu is supported by NSF CNS-2151131, CNS-2140975, a Microsoft Research Faculty Fellowship, and a Google Research Scholar award.

## References

[ABR12] Benny Applebaum, Andrej Bogdanov, and Alon Rosen. A dichotomy for local small-bias generators. In TCC, 2012.
[ACL ${ }^{+} 22$ ] Martin R. Albrecht, Valerio Cini, Russell W. F. Lai, Giulio Malavolta, and Sri AravindaKrishnan Thyagarajan. Lattice-based SNARKs: Publicly verifiable, preprocessing, and recursively composable. In CRYPTO, 2022.
[AFMP20] Navid Alamati, Luca De Feo, Hart Montgomery, and Sikhar Patranabis. Cryptographic group actions and applications. In ASIACRYPT, 2020.
[AK19] Benny Applebaum and Eliran Kachlon. Sampling graphs without forbidden subgraphs and unbalanced expanders with negligible error. In FOCS, 2019.
[AL16] Benny Applebaum and Shachar Lovett. Algebraic attacks against random local functions and their countermeasures. In STOC, 2016.
[App12] Benny Applebaum. Pseudorandom generators with long stretch and low locality from random local one-way functions. In STOC, 2012.
[App15] Benny Applebaum. The cryptographic hardness of random local functions - survey. IACR Cryptol. ePrint Arch., 2015.
[BBHR18] Eli Ben-Sasson, Iddo Bentov, Yinon Horesh, and Michael Riabzev. Scalable, transparent, and post-quantum secure computational integrity. IACR Cryptol. ePrint Arch., 2018, 2018.
[BCC88] Gilles Brassard, David Chaum, and Claude Crépeau. Minimum disclosure proofs of knowledge. 7. Comput. Syst. Sci., 37(2), 1988.
[ $\left.B C C^{+} 17\right]$ Nir Bitansky, Ran Canetti, Alessandro Chiesa, Shafi Goldwasser, Huijia Lin, Aviad Rubinstein, and Eran Tromer. The hunting of the SNARK. 7. Cryptol., 30(4), 2017.
[BCCT12] Nir Bitansky, Ran Canetti, Alessandro Chiesa, and Eran Tromer. From extractable collision resistance to succinct non-interactive arguments of knowledge, and back again. In ITCS, 2012.
[ $\mathrm{BCI}^{+}$13] Nir Bitansky, Alessandro Chiesa, Yuval Ishai, Rafail Ostrovsky, and Omer Paneth. Succinct non-interactive arguments via linear interactive proofs. In TCC, 2013.
[BCPR14] Nir Bitansky, Ran Canetti, Omer Paneth, and Alon Rosen. On the existence of extractable one-way functions. In STOC, 2014.
[ $\left.\mathrm{BDK}^{+} 11\right]$ Boaz Barak, Yevgeniy Dodis, Hugo Krawczyk, Olivier Pereira, Krzysztof Pietrzak, François-Xavier Standaert, and Yu Yu. Leftover hash lemma, revisited. In CRYPTO, 2011.
[BFM88] Manuel Blum, Paul Feldman, and Silvio Micali. Non-interactive zero-knowledge and its applications (extended abstract). In STOC, 1988.
[BHY09] Mihir Bellare, Dennis Hofheinz, and Scott Yilek. Possibility and impossibility results for encryption and commitment secure under selective opening. In EUROCRYPT, 2009.
[BISW17] Dan Boneh, Yuval Ishai, Amit Sahai, and David J. Wu. Lattice-based SNARGs and their application to more efficient obfuscation. In EUROCRYPT, 2017.
[BISW18] Dan Boneh, Yuval Ishai, Amit Sahai, and David J. Wu. Quasi-optimal SNARGs via linear multi-prover interactive proofs. In EUROCRYPT, 2018.
[BKM20] Zvika Brakerski, Venkata Koppula, and Tamer Mour. NIZK from LPN and trapdoor hash via correlation intractability for approximable relations. In CRYPTO, 2020.
[ $\left.\mathrm{BKP}^{+} 23 \mathrm{a}\right]$ Nir Bitansky, Chethan Kamath, Omer Paneth, Ron Rothblum, and Prashant Nalini Vasudevan. Batch proofs are statistically hiding. IACR Cryptol. ePrint Arch., 2023. https://eprint.iacr.org/archive/ 2023/754/20230525:044715.
[ $\left.\mathrm{BKP}^{+} 23 \mathrm{~b}\right]$ Nir Bitansky, Chethan Kamath, Omer Paneth, Ron Rothblum, and Prashant Nalini Vasudevan. Batch proofs are statistically hiding. IACR Cryptol. ePrint Arch., 2023. https://eprint.iacr.org/archive/ 2023/754/20230725:080608.
[BQ09] Andrej Bogdanov and Youming Qiao. On the security of goldreich's one-way function. In APPROXRANDOM, 2009.
[BS23] Ward Beullens and Gregor Seiler. Labrador: Compact proofs for R1CS from module-sis. In EUROCRYPT, 2023.
[BY92] Mihir Bellare and Moti Yung. Certifying cryptographic tools: The case of trapdoor permutations. In CRYPTO, 1992.
[CBBZ23] Binyi Chen, Benedikt Bünz, Dan Boneh, and Zhenfei Zhang. Hyperplonk: Plonk with linear-time prover and high-degree custom gates. In EUROCRYPT, 2023.
[CCH $\left.{ }^{+} 19\right]$ Ran Canetti, Yilei Chen, Justin Holmgren, Alex Lombardi, Guy N. Rothblum, Ron D. Rothblum, and Daniel Wichs. Fiat-shamir: from practice to theory. In STOC, 2019.
$\left[\mathrm{CDM}^{+} 18\right]$ Geoffroy Couteau, Aurélien Dupin, Pierrick Méaux, Mélissa Rossi, and Yann Rotella. On the concrete security of goldreich's pseudorandom generator. In ASIACRYPT, 2018.
[CEMT09] James Cook, Omid Etesami, Rachel Miller, and Luca Trevisan. Goldreich's one-way function candidate and myopic backtracking algorithms. In TCC, 2009.
[CGJ ${ }^{+} 22$ Arka Rai Choudhuri, Sanjam Garg, Abhishek Jain, Zhengzhong Jin, and Jiaheng Zhang. Correlation intractability and SNARGs from sub-exponential DDH. IACR Cryptol. ePrint Arch., 2022.
[CHK03] Ran Canetti, Shai Halevi, and Jonathan Katz. A forward-secure public-key encryption scheme. In EUROCRYPT, 2003.
$\left[\mathrm{CHM}^{+} 20\right]$ Alessandro Chiesa, Yuncong Hu, Mary Maller, Pratyush Mishra, Noah Vesely, and Nicholas P. Ward. Marlin: Preprocessing zkSNARKs with universal and updatable SRS. In EUROCRYPT, 2020.
[CJJ21a] Arka Rai Choudhuri, Abhishek Jain, and Zhengzhong Jin. Non-interactive batch arguments for NP from standard assumptions. In CRYPTO, 2021.
[CJJ21b] Arka Rai Choudhuri, Abhishek Jain, and Zhengzhong Jin. SNARGs for P from LWE. In FOCS, 2021.
[CL18] Ran Canetti and Amit Lichtenberg. Certifying trapdoor permutations, revisited. In TCC, 2018.
[CM01] Mary Cryan and Peter Bro Miltersen. On pseudorandom generators in NC. In FOCS, 2001.
[COS20] Alessandro Chiesa, Dev Ojha, and Nicholas Spooner. Fractal: Post-quantum and transparent recursive proofs from holography. In EUROCRYPT, 2020.
[CPW23] Suvradip Chakraborty, Manoj Prabhakaran, and Daniel Wichs. A map of witness maps: New definitions and connections. In PKC, 2023.
[DFH12] Ivan Damgård, Sebastian Faust, and Carmit Hazay. Secure two-party computation with low communication. In TCC, 2012.
[DGI ${ }^{+}$19] Nico Döttling, Sanjam Garg, Yuval Ishai, Giulio Malavolta, Tamer Mour, and Rafail Ostrovsky. Trapdoor hash functions and their applications. In CRYPTO, pages 3-32, 2019.
[DGKV22] Lalita Devadas, Rishab Goyal, Yael Kalai, and Vinod Vaikuntanathan. Rate-1 non-interactive arguments for batch-NP and applications. In FOCS, 2022.
[DN02] Ivan Damgård and Jesper Buus Nielsen. Perfect hiding and perfect binding universally composable commitment schemes with constant expansion factor. In CRYPTO, 2002.
[DRS04] Yevgeniy Dodis, Leonid Reyzin, and Adam D. Smith. Fuzzy extractors: How to generate strong keys from biometrics and other noisy data. In EUROCRYPT, 2004.
[FLS90] Uriel Feige, Dror Lapidot, and Adi Shamir. Multiple non-interactive zero knowledge proofs based on a single random string (extended abstract). In FOCS, 1990.
[GGPR13] Rosario Gennaro, Craig Gentry, Bryan Parno, and Mariana Raykova. Quadratic span programs and succinct NIZKs without PCPs. In EUROCRYPT, 2013.
[GMR85] Shafi Goldwasser, Silvio Micali, and Charles Rackoff. The knowledge complexity of interactive proofsystems (extended abstract). In STOC, 1985.
[Gol00] Oded Goldreich. Candidate one-way functions based on expander graphs. IACR Cryptol. ePrint Arch., 2000.
[GOS06] Jens Groth, Rafail Ostrovsky, and Amit Sahai. Perfect non-interactive zero knowledge for NP. In EUROCRYPT, 2006.
[GOS12] Jens Groth, Rafail Ostrovsky, and Amit Sahai. New techniques for non-interactive zero-knowledge. 7 . ACM, 59(3), 2012.
[GR13] Oded Goldreich and Ron D. Rothblum. Enhancements of trapdoor permutations. 7. Cryptol., 26(3), 2013.
[Gro10] Jens Groth. Short pairing-based non-interactive zero-knowledge arguments. In ASIACRYPT, 2010.
[Gro16] Jens Groth. On the size of pairing-based non-interactive arguments. In EUROCRYPT, 2016.
[GSWW22] Rachit Garg, Kristin Sheridan, Brent Waters, and David J. Wu. Fully succinct batch arguments for np from indistinguishability obfuscation. In TCC, 2022.
[GW11] Craig Gentry and Daniel Wichs. Separating succinct non-interactive arguments from all falsifiable assumptions. In STOC, 2011.
[HJKS22] James Hulett, Ruta Jawale, Dakshita Khurana, and Akshayaram Srinivasan. Snargs for P from subexponential DDH and QR. In EUROCRYPT, 2022.
[HLOV11] Brett Hemenway, Benoît Libert, Rafail Ostrovsky, and Damien Vergnaud. Lossy encryption: Constructions from general assumptions and efficient selective opening chosen ciphertext security. In ASIACRYPT, 2011.
[JJ21] Abhishek Jain and Zhengzhong Jin. Non-interactive zero knowledge from sub-exponential DDH. In EUROCRYPT, 2021.
[JLS21] Aayush Jain, Huijia Lin, and Amit Sahai. Indistinguishability obfuscation from well-founded assumptions. In STOC, 2021.
[JLS22] Aayush Jain, Huijia Lin, and Amit Sahai. Indistinguishability obfuscation from LPN over $\mathbb{F}_{p}$, DLIN, and PRGs in $\mathrm{NC}^{0}$. In EUROCRYPT, 2022.
[KLVW23] Yael Tauman Kalai, Alex Lombardi, Vinod Vaikuntanathan, and Daniel Wichs. Boosting batch arguments and RAM delegation. In STOC, 2023.
[KMY20] Fuyuki Kitagawa, Takahiro Matsuda, and Takashi Yamakawa. NIZK from SNARG. In TCC, 2020.
[KPY19] Yael Tauman Kalai, Omer Paneth, and Lisa Yang. How to delegate computations publicly. In STOC, 2019.
[KVZ21] Yael Tauman Kalai, Vinod Vaikuntanathan, and Rachel Yun Zhang. Somewhere statistical soundness, post-quantum security, and SNARGs. In TCC, 2021.
[Lin17] Huijia Lin. Indistinguishability obfuscation from SXDH on 5-linear maps and locality-5 PRGs. In CRYPTO, 2017.
[Lip13] Helger Lipmaa. Succinct non-interactive zero knowledge arguments from span programs and linear error-correcting codes. In ASIACRYPT, 2013.
[LPWW20] Benoît Libert, Alain Passelègue, Hoeteck Wee, and David J. Wu. New constructions of statistical NIZKs: Dual-mode DV-NIZKs and more. In EUROCRYPT, 2020.
[LT17] Huijia Lin and Stefano Tessaro. Indistinguishability obfuscation from trilinear maps and block-wise local PRGs. In CRYPTO, 2017.
[Mic95] Silvio Micali. Computationally-sound proofs. In Proceedings of the Annual European Summer Meeting of the Association of Symbolic Logic, 1995.
[MST03] Elchanan Mossel, Amir Shpilka, and Luca Trevisan. On $\varepsilon$-biased generators in $\mathrm{NC}^{0}$. In FOCS, 2003.
[OW14] Ryan O'Donnell and David Witmer. Goldreich's PRG: evidence for near-optimal polynomial stretch. In CCC, 2014.
[PHGR13] Bryan Parno, Jon Howell, Craig Gentry, and Mariana Raykova. Pinocchio: Nearly practical verifiable computation. In IEEE Symposium on Security and Privacy, 2013.
[PS19] Chris Peikert and Sina Shiehian. Noninteractive zero knowledge for NP from (plain) learning with errors. In CRYPTO, 2019.
[PW08] Chris Peikert and Brent Waters. Lossy trapdoor functions and their applications. In STOC, 2008.
[QRW19] Willy Quach, Ron D. Rothblum, and Daniel Wichs. Reusable designated-verifier NIZKs for all NP from CDH. In EUROCRYPT, 2019.
[Set20] Srinath T. V. Setty. Spartan: Efficient and general-purpose zkSNARKs without trusted setup. In CRYPTO, 2020.
[Üna23] Akin Ünal. Worst-case subexponential attacks on PRGs of constant degree or constant locality. In EUROCRYPT, 2023.
[WW22] Brent Waters and David J. Wu. Batch arguments for NP and more from standard bilinear group assumptions. In CRYPTO, 2022.


[^0]:    ${ }^{1}$ More recently, Chakraborty et al. [CPW23] showed a similar implication holds starting from a mildly-compact computational witness map (a simpler primitive that is implied by a SNARG for NP).
    ${ }^{2}$ For instance, this is satisfied if the length of the proof scales polylogarithmically in the number of instances. More generally, we can instantiate the theorem even if the proof size scales with $T^{1 / 2-\varepsilon}$, where $T$ is the number of instances and $\varepsilon>0$ is a constant. We refer to Corollary 3.15 for the precise characterization.

[^1]:    ${ }^{3}$ Semi-adaptive soundness for a batch argument says that the adversary must first commit to the index $i$ of the false statement in the soundness game. It can adaptively choose the statements $x_{1}, \ldots, x_{T}$ after seeing the CRS, with the restriction that instance $x_{i}$ must be false.

[^2]:    ${ }^{4}$ In some settings, we can require a stronger "equivocation" property in hiding mode where given trapdoor information, one can sample a commitment $c$ and openings for $c$ to any value. Our constructions do not require equivocation.

[^3]:    ${ }^{5}$ Previous works [KPY19, CJJ21a, CJJ21b, WW22, DGKV22, CGJ ${ }^{+} 22$ ] also impose requirements on the size of the CRS and the running time of the verifier. These additional properties are not needed in our work.

[^4]:    ${ }^{6}$ As noted above, we require that $\ell(n) \geq m B$. If $\ell(n)>m B$, we truncate the output of $G_{n}$ to output a string of length exactly $m B$.

[^5]:    ${ }^{7}$ We note here that additionally assuming a rate-1 string oblivious transfer protocol [DGI $\left.{ }^{+} 19\right]$, such a BARG can be transformed into a BARG where the proof size scales polylogarithmically with the number of instances [KLVW23]. In this case, we would be able to appeal to our previous instantiation.

