Fully-succinct Publicly Verifiable Delegation from Constant-Size Assumptions

Alonso González* and Alexandros Zacharakis**

[1] Toposware Inc. alonso.gonzalez@toposware.com
[2] Universitat Pompeu Fabra, Barcelona, Spain alexandros.zacharakis@upf.edu

Abstract. We construct a publicly verifiable, non-interactive delegation scheme for any polynomial size arithmetic circuit with proof-size and verification complexity comparable to those of pairing based zk-SNARKS. Concretely, the proof consists of $O(1)$ group elements and verification requires $O(1)$ pairings and n group exponentiations, where n is the size of the input. While known SNARK-based constructions rely on non-falsifiable assumptions, our construction can be proven sound under any constant size ($k \geq 2$) k-Matrix Diffie-Hellman (k-MDDH) assumption. However, the size of the reference string as well as the prover's complexity are quadratic in the size of the circuit. This result demonstrates that we can construct delegation from very simple and well-understood assumptions. We consider this work a first step towards achieving practical delegation from standard, falsifiable assumptions.

Our main technical contributions are first, the introduction and construction of what we call "no-signaling, somewhere statistically binding commitment schemes". These commitments are extractable for any small part $\boldsymbol{x}_S$ of an opening $\boldsymbol{x}$, where $S \subseteq [n]$ is of size at most K. Here n is the dimension of $\boldsymbol{x}$ and $\boldsymbol{x}_S = (x_i)_{i \in S}$. Importantly, for any $S' \subseteq S$, extracting $\boldsymbol{x}_{S'}$ can be done independently of $S \setminus S'$. Second, we use these commitments to construct more efficient "quasi-arguments" with no-signaling extraction, introduced by Paneth and Rothblum (TCC 17). These arguments allow extracting parts of the witness of a statement and checking it against some local constraints *without revealing which part is checked*. We construct pairing-based quasi arguments for linear and quadratic constraints and combine them with the low-depth delegation result of González et. al. (Asiacrypt 19) to construct the final delegation scheme.

Keywords: Delegation · Succinct Arguments · Non Interactive Zero Knowledge.

---

* This work was done while the author was part of LIP laboratory at the ENS de Lyon, France

** Research Supported by fellowships from "la Caixa" Foundation (ID 100010434). The fellowship code is LCF/BQ/DI18/11660053. Funding is also from the European Union's Horizon 2020 research and innovation program under the Marie Skłodowska-Curie grant agreement No. 713673.

1 Introduction

In a delegation scheme, a verifier with limited computational resources (a mobile device for example) wishes to delegate a heavy but still polynomial computation to an untrusted prover. The prover, with more computational power but still of polynomial time, computes a proof which the verifier accepts or rejects. Given the limitations of the verifier, the proof should be as short as possible and the verification process should consume as few computational resources as possible. Additionally, the construction of the proof should not be much costlier than performing the computation itself.

A delegation scheme can be easily constructed from a zero-knowledge Succinct Non-Interactive Argument of Knowledge (zk-SNARK) for NP. Schemes like [19, 25] are very appealing in practice because a proof consists of only a constant number of group elements and verification requires the evaluation of a constant number of pairings.[3] The downside is that these zk-SNARKs are based on strong and controversial assumptions such as the knowledge of exponent assumption or the generic group model.

Such assumptions are called non-falsifiable because there is no way of efficiently deciding whether an adversary breaks the assumption or not. In such assumptions, the adversary is treated in a non black box way and the assumption argues about *how* an adversary performs a computation instead of *what* computation it cannot perform. Since zk-SNARKs can handle even NP computations, soundness becomes an essentially non-falsifiable property where one needs to decide whether an adversary produces a true or false statement without any witness but only with a very short proof. Gentry and Wichs [20] proved that zk-SNARKs for NP are (in a broad sense) impossible to construct without resorting to non-falsifiable assumptions.

While this impossibility result justifies the use of such assumptions for non-deterministic computation, this is not the case for delegation of computation which only considers deterministic computation. Indeed, in this case, soundness becomes an efficiently falsifiable statement: determining whether the adversary breaks soundness simply requires to evaluate the delegated polynomial computation on some input x and check whether it is accepting or rejecting. Actually, getting delegation from falsifiable assumptions is easy in general: let Π be a SNARK for NP. For a binary relation R, the assumption "Π is sound for R" is in general non-falsifiable since checking membership in the corresponding language is hard and the SNARK proof does not help as shown by [20]. On the contrary, for a relation R in P, the assumption becomes falsifiable since one can efficiently compute $R(x)$. Nevertheless, the important issue is to consider the *quality* of the assumption in place since the assumption "the proof system is sound" is tautological. Ideally, we should rely on simple and well understood assumptions *without* sacrificing other desirable properties.

Almost all known constructions that base their soundness on falsifiable assumptions (or even no assumptions at all) come with some compromises: they

[3] Note that zero-knowledge is not necessary.

(1) are not expressive enough to capture all polynomial time computation [32, 24, 11, 29] (2) are interactive [21, 45], (3) are designated verifier [35, 36, 33, 8, 5] or (4) rely on strong (yet falsifiable) assumptions related to obfuscation [12, 40, 6, 3, 13] or multi-linear maps [44].

An exception to this is a construction of Kalai et al. [34] of a delegation scheme for any poly-time computation based on a newly introduced q-size assumption in bilinear groups. The size of the assumption is $q = \log T$ and T is the time needed to perform the computation. As for efficiency, the size of the proof is $\mathsf{polylog}(T)$ group elements which becomes $\mathsf{poly}(\kappa)$ if $T \leq 2^{\kappa}$.

However, in spite of the recent progress, there's still a gap in the proof size and verification with respect to the most efficient known constructions, namely those based on paring based zk-SNARKs.

1.1 Our results

In this work we consider the question *"what are the simplest assumptions that imply publicly verifiable, non-interactive delegation of computation"?* Here "*simple*" should be interpreted as falsifiable and well understood. Having practicality in mind as well, we would also want a delegation scheme that competes in efficiency with the most efficient constructions to date, namely those that are based on non-falsifiable assumptions.

The main contribution of this work is the construction of a fully-succinct, non-interactive, publicly verifiable delegation scheme from any k-Matrix Diffie-Hellman assumption (k-MDDH) for $k \geq 2$, as for example the decisional linear assumption (DLin) [7]. In the more efficient setting of asymmetric groups, soundness can be based on the natural translation of symmetric DLin where the challenge is encoded in both groups (the SDlin assumption of [22]). Here by fully-succinct we mean that the proof size is linear in the security parameter and verification requires a linear number of operations (whose complexity depends only on the security parameter) in the size of the input of the computation. We achieve these goals but with the drawback that the prover computation and the size of the crs are quadratic in the size of the circuit. Our main contribution is summarized in the next (informal) theorem.

Theorem 1. *(*Informal*). There exists a non-interactive, publicly verifiable delegation scheme for any polynomial size circuit C with n-size input that is adaptively sound under any k-MDDH assumption for $k \geq 2$ with the following efficiency properties: the crs size is $\mathsf{poly}(\kappa)|C|^2$, prover complexity is $\mathsf{poly}(\kappa)|C|^2$, proof size is $\mathsf{poly}(\kappa)$ and verification complexity is $\mathsf{poly}(\kappa)n$.*

Our construction is also concretely efficient as far as proof size and verification complexity are concerned. The proof comprises of 10+8 group elements of an asymmetric bilinear group and verification requires n exponentiations plus 36 evaluations of the pairing function, where n is the size of the input. The attractive concrete efficiency is achieved due to the structure-preserving nature [1] of our construction. This notion captures that all algorithms solely perform group

operations, namely they are *algebraic*, and there is no need to encode cryptographic primitives such as hash functions or pairings as arithmetic circuits, a process that is very inefficient in practice.

This result demonstrates two things. First, delegation of computation can be based on very simple, standard assumptions. Second, its structure preserving nature hints to the plausibility of practically efficient delegation schemes comparable in efficiency with the ones based on SNARKs, but under simple, standard assumptions. In table 1 we present a comparison of our delegation of computation construction with other pairing based schemes.

Table 1: Comparison between different pairing based delegation schemes and our results.

	Language	Verification	Proof size	CRS size	Assumption
[19][25]	AC	$n\mathsf{e} + O(1)\mathsf{p}$	$O(\kappa)$	$O(\|C\|\kappa)$	Non Falsifiable
[34] (base case)	RM	$n\mathsf{e} + \mathsf{poly}(\log d)\mathsf{p}$	$O(\kappa \log d)$	$O((n+d)\kappa)$	$\log d$-Assumption
[24]	AC	$n\mathsf{e} + O(d)\mathsf{p}$	$O(d\kappa)$	$O(\|C\|\kappa)$	s-Assumption
This work	AC	$n\mathsf{e} + O(1)\mathsf{p}$	$O(\kappa)$	$O(\|C\|^2\kappa)$	DLin/SDLin

Verification is given in number exponentiations (e) and pairings (p). d is the circuit depth/number of steps of a computation, n the number of inputs, s the circuit width/computation space and $|C|$ the circuit size. AC stands for "Arithmetic Circuit" and RM for "RAM Machine". For [34] we only consider the "base case" and not the "bootstrapped" constructions, because bootstrapping adds a considerable overhead and is thus incomparable in terms of group operations. We stress out, however, that the crs size of the bootstrapped construction is sublinear in the time of the computation.

No-Signaling SSB Commitments and Succinct Pairing-based Quasi-Arguments. We follow and extend the ideas of Paneth and Rothblum [44] and Kalai et al. [34] for constructing delegation schemes for poly-time computations from what they called quasi-arguments of knowledge with no-signaling extractors. First, we formalize a similar notion for commitment schemes and show that the somewhere statistically binding (SSB) commitments of [22, 18] are no-signaling when they also have what we call an "oblivious trapdoor generator". Second, we use the no-signaling SSB commitments to construct more efficient constant-sized quasi-arguments of knowledge for linear and quadratic relations. We achieve this by combining SSB commitments with the very efficient quasi-adaptive non-interactive zero-knowledge arguments for linear [30, 41, 31, 39] and quadratic relations [22, 16]. To this aim, we also show that the QA-NIZK arguments can be easily modified to have no-signaling extractors under standard assumptions.

Applications to NIZK. Our construction can be turned into a NIZK argument for NP of size $n+O(1)$ group elements -namely $O(n\kappa)$ proof size- under the same

assumptions where n is the number of public an secret inputs of the circuit. In table 2 we provide a comparison of our NIZK construction and the literature. Using standard techniques, the argument implies compact NIZK for NP with proof size $O(n) + \mathsf{poly}(\kappa)$. That is, the size of the proof is proportional to the size of the input and the security parameter only gives an additive overhead. In comparison, the state of the art is $O(|C|) + \mathsf{poly}(\kappa)$ for poly-sized boolean circuits and $O(n) + \mathsf{poly}(\kappa)$ for log-depth boolean circuits [38, 37]. We note that a similar result can be obtained by [34], albeit with a stronger assumption.

Table 2: Comparison between different pairing based NIZK schemes and our results.

	Language	Verification	Proof size	CRS size	Assumption
[26]	AC	$O(\|C\|)\mathtt{p}$	$O(\|C\|\kappa)$	$O(\kappa)$	SXDH
[19][25]	AC	$O(1)\mathtt{p}$	$O(\kappa)$	$O(\|C\|\kappa)$	Non Falsifiable
[24]	BC	$O(n+d)\mathtt{p}$	$O((n+d)\kappa)$	$O(\|C\|\kappa)$	s-Assumption
[37]	NC[1]	$O(\|C\|)\mathsf{poly}(\kappa)$	$n\mathsf{poly}(\kappa)$	$\mathsf{poly}(\|C\|, \kappa, 2^d)$	DLin
This work	BC	$O(n)\mathtt{p}$	$nO(\kappa)$	$O(\|C\|^2\kappa)$	DLin/SDLin

Verification is given in number of pairings $\mathtt{p}$. d is the circuit depth, n the number of (public and secret) inputs, s the circuit width and $|C|$ the circuit size. AC stands for "Arithmetic Circuit" and BC for "Boolean Circuit".

Our argument can be also used to construct zk-SNARKs from quantitatively weaker assumptions than the state of the art. Indeed, the strongest assumption used in zk-SNARKs such as [19, 25] is a knowledge assumption which states that an adversary computing some elements of a bilinear group, satisfying a particular relation, must know their discrete logarithms.[4] Such assumption is used to extract an assignment to each of the circuit wires. The "size" of such assumption is proportional to the number of extracted values, which in this case is the size of the circuit. Since our argument only requires the reduction to know the input of the circuit, we can rely on a knowledge assumption only for extracting the input. As a consequence the size of the assumption is drastically shortened. Since these assumptions are stronger as the size of the assumption increases and given that we lack good understanding of them, it is always safer to rely on shorter assumptions. Also, weaker assumptions translates to better concrete efficiency by using smaller security parameters.[5]

[4] Actually, the adversary must know a representation of these values as a linear combination of a set of group elements that she receives as input.

[5] We note, however, that in the case of non-falsifiable assumptions it not clear how an appropriate security parameter should be chosen.

2 Technical Overview

To construct the delegation scheme we follow a commit-and-prove approach, which means that we first commit to the witness (the satisfying assignment of wires in a circuit) and then show that this witness satisfies some relation. We use somewhere statistically binding (SSB) commitments as those used in [22, 23, 18] and show that they satisfy a *no-signaling extraction* property. Then, we do the same for the so called quasi-adaptive NIZK arguments for linear spaces [30, 41, 31, 39] and for quadratic relations [22, 16]. From these primitives we can construct delegation for bounded-space computations/bounded width circuits with proof-size independent of the depth of the computation by following the techniques of [44, 34]. To get a succinct proof-size, in addition to the "depth compression" we must also perform a "width compression". To this end, we use ideas from the delegation scheme for bounded depth computations of González and Ràfols [24] and remove the necessity of a q-assumption to rely solely on constant size assumptions. To combine both "compressions" efficiently we exploit the fact that [24] is structure preserving and the verifier is a bounded width circuit. In the next sections we present these techniques.

2.1 No-Signaling Somewhere Statistically Binding Commitments/Hashing

Somewhere statistically binding (SSB) hashing/commitments[6] were introduced by Hubacek and Wichs [28] and then improved by [43], and have been used for constructing efficient NIZK proofs [22, 23] as well as ring signatures [4].

An SSB commitment scheme is a generalization of dual mode commitments [27] where the commitment key can be sampled from many computationally indistinguishable distributions, each of which is making the commitments statistically binding for a number of K coordinates of the commited value. That is, when commiting to a vector $\boldsymbol{m} = (m_1, \ldots, m_n)$ with a commitment key ck_S associated with a set $S \subseteq [n]$ of size at most K, no (even computationally unbounded) adversary can compute a commitment c and two valid openings $\boldsymbol{m}, \boldsymbol{m}'$ such that for some $i \in S$ it holds that $m_i \neq m'_i$, except with negligible probability. Importantly, the size of the commitment c should be independent of n but may depend on the value K.

Known SSB commitments constructions are also extractable[7], that is, there exists an efficient algorithm that has some trapdoor information associated with

[6] Through this paper we will refer to "commitments" while technically they are "hashes". We do so because in the context of NIZK proofs is traditional to commit to the witness and then prove that the committed value satisfy some relation. However, since we are less interested in zero-knowledge, the randomness of such commitments is 0 (or fixed/inexistent) and we end up with hashes.

[7] In the context of bilinear groups, we can consider f-extraction where one only extracts f applied to the witness. In particular, it is usual to consider f the (one-way) function that maps elements in $\mathbb{Z}_p$ to one of the base groups $\mathbb{G}_1$ or $\mathbb{G}_2$. This is the notion of extractability we use in this work and is enough to obtain our results.

ck_S and can efficiently extract from a commitment c a valid opening $(m_i)_{i\in S}$. Note that the notion of a "valid opening" is well-defined due to the statistical binding property on the set S.

We argue that the SSB extractor has many similarities with the no-signaling extractors of [44, 34]. First, we briefly recall what a no-signaling extractor is in the context of quasi arguments of knowledge. A quasi argument is a proof system for a relation that defines some local constraints on the statement/witness pair. The requirement is that there exists a *no signaling extractor* that allows extracting a part of the witness from a verifying proof that is locally correct. Furthermore, each part of the extracted local witness can be in a sense extracted independently. This is formalized by requiring that extracting local witness w_S for a set S and restricting it to the variables $S' \subseteq S$ is computationally indistinguishable from extracting $w_{S'}$ for the set S'. As we shall see shortly, this property is extremely useful when constructing delegation schemes.

In the case of SSB commitments, extractability of the local opening is just a local soundness guarantee. Additionally, indistinguishability of the commitment keys is a weaker form of the no-signaling property. Indeed, a no-signaling extractor must produce commitment keys which are indistinguishable for the various possible extractable sets. Otherwise a distinguisher for sets S, S' can be used for wining in the no-signaling game even without the extracted value. Nevertheless, this alone does not satisfy the no-signaling property: some information about the positions where the crs is programmed to extract might be revealed by (parts of) the extracted local openings.

We strengthen the indistinguishability property of the distributions of the commitment keys of SSB commitments to give them a no-signaling flavour. Roughly speaking, we require that the distributions of the commitment keys are computationally indistinguishable *even if the adversary has access to local openings associated with a set S' of committed values*. These local openings trivially reveal information about the set S' but we require that they do not leak information about the values outside of S'. That is, for any sets $S' \subseteq S$ of size at most K, the commitment keys $ck_S, ck_{S'}$ are computationally indistinguishable even if we allow the distinguisher access to local openings of S'.

Remark 1 (Connection with PIR). Somewhere statistically binding commitments/hashing is closely related with single server Private Information Retrieval Schemes (PIR) when the SSB commitment is also extractable. Indeed, we can think of the commitment key for an index i of the SSB as a PIR query and the commitment/hash as the PIR answer. Then, one can decode the PIR query using the trapdoor associated with the commitment key. In our work, the SSB commitments we use are different from PIRs in three ways: (1) we do not extract the PIR answers, but we f-extract, specifically we extract encodings of messages in a group but not their discrete logarithms, (2) we directly use SSBs with locality greater than one instead of making parallel PIR queries to improve concrete efficiency and (3) the size of the commitment key is proportional to the size of the commited values, while in PIRs the query should be small compared to the database size. Furthermore, we exploit in a non-black box way the properties

as well as the algebraic structure of the SSB commitments to compose them with other protocols, such as group based quasi-adaptive non-interactive zero knowledge arguments.

SSB Commitments with Oblivious Trapdoor Generation. We define a stronger notion for SSB commitment schemes, *oblivious trapdoor generation*, which implies the no-signaling property. This notion is easier to work with in our particular constructions.

Intuitively, this notion captures that there exists a different, *oblivious* key generation algorithm that can generate the commitment key for S and a trapdoor for a subset $S' \subseteq S$ obliviously of $S \setminus S'$ for any subset S' of the larger set S of binding coordinates. More concretely, the oblivious key generation algorithm takes as input a commitment key ck_S binding at S and the description of a subset $S' \subseteq S$ and outputs an *identically distributed key* together with a trapdoor for extracting values in the small set S'. We emphasize that this algorithm does not take as input neither the description of S nor the trapdoor associated with it. Intuitively, the key generation algorithm is oblivious of $S \setminus S'$ (it might even be that $S \setminus S' = \emptyset$) due to the indistinguishability of commitment keys associated with different sets, in this case S and S'.

This property implies no-signaling commitments. Indeed, this follows easily since (1) by the index set hiding property the commitment key itself does not reveal any information about $S \setminus S'$ and (2) we can use the oblivious key generation algorithm to create a trapdoor for extracting the smaller set *without skewing the distribution of the commitment key*. The latter property means essentially that we are given an oracle to extract the smaller set (by computing the trapdoor for an identically distributed key) which is exactly what the no-signaling property captures.

Constructing Oblivious SSB Commitments. We next describe how to construct efficient SSB commitments with oblivious trapdoor generator. A natural way to construct oblivious SSB commitment with locality parameter K is to concatenate K SSB commitments with locality parameter 1. Consider a set $S = \{s_1, \ldots, s_t\}$ for some $t \leq K$. We can construct a commitment key associated with S by computing t commitment keys/trapdoor pairs $(ck_1, \tau_1), \ldots, (ck_t, \tau_t)$ for sets $\{s_1\}, \ldots, \{s_t\}$, complementing with $K - t$ keys for $\emptyset$ if necessary. To commit to some $\boldsymbol{x} \in \mathcal{M}^n$, where $\mathcal{M}$ is the message space of the commitment, one simply computes $c_1 = \mathsf{Com}_{ck_1}(\boldsymbol{x}), \ldots, c_K = \mathsf{Com}_{ck_K}(\boldsymbol{x})$. Extraction of each x_{s_i} is done using c_{s_i} and the trapdoor τ_{s_i}, independently of the others. The oblivious extractor on input the commitment keys for some unknown S and the description of $S' \subseteq S$ just re-samples the commitment keys for S'.[8] Since it doesn't matter if the trapdoors for positions $i \notin S'$ are not known, this trivial extractor can obliviously generate the trapdoor $\{\tau_i : i \in S'\}$.

[8] Actually, the oblivious key generation needs to know which of the commitments keys $ck_1, \ldots, ck_K$ are perfectly binding for $s' \in S'$. Nevertheless, it should be still oblivious of whether the rest of commitment keys are binding or not.

While this generic construction is enough, we can construct more efficient ones if we consider specific instantiations. More specifically, as we present next, we can have more efficient instantiations (roughly half commitment size compared to the generic one) in the case of commitments derived from the Pedersen commitment scheme.

Notation. We first need to introduce some notation. When $S \subseteq [n]$ we denote with $\overline{S}$ the set $[n] \setminus S$. For a vector $\boldsymbol{x}$ (resp. matrix $\mathbf{G}$) we denote $\boldsymbol{x}_S = (x_i)_{i \in S}$ (resp. $\mathbf{G}_S = (\boldsymbol{g}_i)_{i \in S}$ where $\boldsymbol{g}_i$ is the i-th column of $\mathbf{G}$). Finally, we use implicit notation for groups. That is, given a group $\mathbb{G}$ and a fixed generator $\mathcal{P}$ we denote with $[r]$ the element $r\mathcal{P}$. For vectors and matrices $\boldsymbol{a}, \mathbf{A}$ respectively, we denote with $[\boldsymbol{a}], [\mathbf{A}]$ the natural embeddings of $\boldsymbol{a}, \mathbf{A}$ to $\mathbb{G}$.

For vectors $\boldsymbol{a}, \boldsymbol{b}$, we denote $\boldsymbol{a} \circ \boldsymbol{b} = (a_i b_i)_i$ the Hadamard product of them, and for matrices $\mathbf{A} = (a_{i,j})_{i,j}$, $\mathbf{B}$ we denote $\mathbf{A} \otimes \mathbf{B} = (a_{i,j}\mathbf{B})_{i,j}$ their Kronecker product. We will be using the mixed-product property of kronecker products, which says that $(\mathbf{A} \otimes \mathbf{B})(\mathbf{C} \otimes \mathbf{D}) = (\mathbf{AC}) \otimes (\mathbf{BD})$ whenever $\mathbf{A}, \mathbf{B}, \mathbf{C}, \mathbf{D}$ have the appropriate dimensions.

Efficient SSB Commitments. We next present an oblivious SSB construction based on the Pedersen commitment scheme. This construction was implicit in [22] and later generalized in [18]. Later we will see that it also satisfies the stronger notion of oblivious trapdoor generation.

Let $\mathbb{G}$ be a group of size p. For message space $\mathbb{Z}_p^d$, locality parameter $K \in \mathbb{N}$ and a subset $S \subseteq [d]$ of size $t \leq K$, the commitment key is defined as follows: $\mathbf{G} = (\mathbf{G}_S | \mathbf{G}_{\overline{S}})\mathbf{P}$ and

$$\mathbf{G}_S \leftarrow \mathbb{Z}_p^{(K+1)\times t}, \qquad \mathbf{G}_0 \leftarrow \mathbb{Z}_p^{(K+1)\times(K+1-t)},^{9}$$

$$\mathbf{\Gamma} \leftarrow \mathbb{Z}_p^{(K+1-t)\times(d-t)}, \qquad \mathbf{G}_{\overline{S}} = \mathbf{G}_0\mathbf{\Gamma}.$$

Matrix $\mathbf{P} \in \{0,1\}^{d\times d}$ is a permutation matrix associated to S such that $\mathbf{P}\boldsymbol{e}_{s_i} = \boldsymbol{e}_i$, for $i \leq t$ and $\boldsymbol{e}_i$ the i-th vector of the canonical basis. A commitment to $\boldsymbol{x} \in \mathbb{Z}_p^d$ is computed as $[\boldsymbol{c}] = [\mathbf{G}]\boldsymbol{x} = [\mathbf{G}_S|\mathbf{G}_{\overline{S}}]\mathbf{P}\boldsymbol{x} = [\mathbf{G}_S]\boldsymbol{x}_S + [\mathbf{G}_{\overline{S}}]\boldsymbol{x}_{\overline{S}}$. Note that the columns of $\mathbf{G}_S$ are linearly independent from the columns of $\mathbf{G}_{\overline{S}}$ with overwhelming probability, since $\mathbf{Im}(\mathbf{G}_{\overline{S}}) \subseteq \mathbf{Im}(\mathbf{G}_0)$ and $(\mathbf{G}_S|\mathbf{G}_0)$ is a basis of $\mathbb{Z}_p^{K+1}$ w.o.p. since this corresponds to a uniform matrix of dimensions $K+1 \times K+1$.

This distribution of commitment keys implies that the parts of the input indexed by S go to the space spanned by $\mathbf{G}_S$ of dimension t, while the rest is mapped to the space spanned by $\mathbf{G}_0$ of dimension $K+1-t$. Since $\mathrm{rank}(\mathbf{G}_S) = t$

[9] It is not always the case that this matrix is uniform. The actual property needed is that this matrix satisfies some hardness assumption. Specifically, the index set hiding property reduces to the $\mathcal{G}$-MDDH assumption (see Section 2.2 for an informal definition) where $\mathcal{G}$ is the distributions from which we sample $\mathbf{G}_0$. When working with symmetric groups, we instantiate using the DLIN assumption. For the sake of simplicity we consider the uniform case in the technical overview.

with overwhelming probability, all the information of $\boldsymbol{x}_S \in \mathbb{Z}_p^t$ can be retrieved from $\boldsymbol{c}$. Even more, there exists an efficiently computable trapdoor $\mathbf{T}_S \in \mathbb{Z}_p^{(K+1)\times t}$ such that $\mathbf{T}_S^\top \mathbf{G}_S = \mathbf{I}_{t\times t}$ and $\mathbf{T}_S^\top \mathbf{G}_{\overline{S}} = \mathbf{0}_{t\times(d-t)}$, and hence

$$\mathbf{T}_S^\top[\boldsymbol{c}] = \mathbf{T}_S^\top[\mathbf{G}\boldsymbol{x}] = \mathbf{T}_S^\top[\mathbf{G}_S\boldsymbol{x}_S + \mathbf{G}_{\overline{S}}\boldsymbol{x}_{\overline{S}}] = [\boldsymbol{x}_S].$$

To compute $\mathbf{T}_S$, it is enough to solve the linear system $\mathbf{T}_S^\top(\mathbf{G}_S \mid \mathbf{G}_0) = (\mathbf{I}_S \mid \mathbf{0})$ which admits a solution since $(\mathbf{G}_S \mid \mathbf{G}_0)$ is a basis of $\mathbb{Z}_p^{K+1}$ with overwhelming probability.

Note that this shows also that the commitment is statistically binding in S. The indistinguishability of commitment keys can be shown with a tight reduction to the DDH assumption as in [18].

Oblivious Trapdoor Generation. One of the main technical contributions of this work is an oblivious trapdoor generator for this commitment scheme, which in turns implies that it is no-signaling. Recall that the property requires that there exists an efficient algorithm, called the oblivious key generation algorithm, that receives as input the description of a set S' of size $t' \leq K$ and a commitment key $[\mathbf{G}]$ sampled for being binding at some unknown $S \supseteq S'$. The algorithm computes a new commitment key $[\mathbf{H}]$ with the following guarantees: (1) it is *statistically close* to $[\mathbf{G}]$ and (2) we also obtain a trapdoor $\mathbf{T}_{S'}$ that allows us to extract local openings for the small set S'.

Since we know that columns in S' are uniformly distributed, we could attempt to sample a uniform matrix $\mathbf{H}_{S'} \leftarrow \mathbb{Z}_p^{(K+1)\times t'}$ and solve the equation $\mathbf{T}_{S'}^\top \mathbf{H}_{S'} = \mathbf{I}_{t'\times t'}$ for some $\mathbf{T}_{S'}$. However, since we don't know the distribution of $[\mathbf{G}_{\overline{S}'}]$ the only hope seems to be to define $[\mathbf{H}_{\overline{S}'}] = [\mathbf{G}_{\overline{S}'}]$ and try to find some $\mathbf{T}_{S'}$ such that $\mathbf{T}_{S'}^\top \mathbf{G}_{\overline{S}'} = \mathbf{0}_{t'\times(d-t')}$. Unfortunately, this amounts to finding elements in the kernel of $[\mathbf{G}_{\overline{S}'}]^\top$ which is in general a computationally hard problem [42].

Instead we make the following observation. Regardless of the distribution of the columns in $S \setminus S'$, the t' lower rows of $\mathbf{G}_{\overline{S}}$ can be always written as a random linear combination of the first $K + 1 - t'$ rows. That is

$$\mathbf{G}_{\overline{S}'} = \begin{pmatrix} \mathbf{A} \\ \mathbf{R}\mathbf{A} \end{pmatrix}, \text{ where } \mathbf{A} \in \mathbb{Z}_p^{K+1-t'\times d-t'} \text{ and } \mathbf{R} \leftarrow \mathbb{Z}_p^{t'\times K+1-t'}.$$

In this case, if we know the matrix $\mathbf{R}$ in the field, it is possible to compute elements in the kernel of $\mathbf{G}_{\overline{S}'}$ by setting

$$\mathbf{T}_{S'} = \begin{pmatrix} -\mathbf{R}^\top\mathbf{C} \\ \mathbf{C} \end{pmatrix}, \text{ for any } \mathbf{C} \in \mathbb{Z}_p^{t'\times t'}.$$

If additionally, we choose some $\mathbf{C}$ that satisfies $\mathbf{T}_{S'}^\top \mathbf{H}_{S'} = \mathbf{I}_{t'\times t'}$ we have computed a trapdoor for S'. This yields a way to compute the rest of the columns: discard the lower t' rows of $\mathbf{G}_{\overline{S}}$, sample a uniform matrix $\mathbf{R}$ as above and complete the last rows with the elements $\mathbf{R}[\mathbf{A}]$. Then, using $\mathbf{R}$, $\mathbf{H}_{S'}$ (which are known in the field) find some $\mathbf{C}$ that satisfies the linear equations and use it to define the trapdoor $\mathbf{T}'_S$.

Lets see in more detail why the previous observation holds. Consider the matrix $\mathbf{G}_0 \in \mathbb{Z}_p^{(K+1)\times(K+1-t)}$ and note that the upper part $\overline{\mathbf{G}}_0$ is a uniformly distributed matrix with more rows than columns; hence $\mathbf{R}\overline{\mathbf{G}}_0$, for $\mathbf{R} \leftarrow \mathbb{Z}_p^{t'\times(K+1-t')}$, is uniformly distributed. This is also valid for all non-binding coordinates since $\mathbf{G}_{\overline{S}} = \mathbf{G}_0\mathbf{\Gamma}$ and then the lower rows follow distribution $\mathbf{R}\overline{\mathbf{G}}_{\overline{S}}$. Next, consider the columns corresponding to the (unknown) binding coordinates $S \setminus S'$. The same argument holds: for some uniform $\mathbf{R}'\overline{\mathbf{G}}_{S\setminus S'}$ is uniform when $\mathbf{R}' \leftarrow \mathbb{Z}_p^{t'\times(K+1-t')}$. It remains to show that using the same randomness for both column sets, i.e. setting $\mathbf{R} = \mathbf{R}'$, does not alter the distribution of the commitment key. Indeed, with overwhelming probability, the columns of $\overline{\mathbf{G}}_0 \in \mathbb{Z}_p^{(K+1-t')\times(K+1-t)}$ and of $\overline{\mathbf{G}}_{S\setminus S'} \in \mathbb{Z}_p^{(K+1-t')\times(t-t')}$ form a basis of $\mathbb{Z}_p^{K+1-t'}$, which means that the matrix $\mathbf{R}^\top$ can be decomposed into two independent components: a random element in $\mathrm{Im}(\overline{\mathbf{G}}_{S\setminus S'}^{\perp})$ and another in $\mathrm{Im}(\overline{\mathbf{G}}_0^{\perp})$. This shows that $\mathbf{R}\overline{\mathbf{G}}_0 = \mathbf{R}_2(\mathbf{G}_{\overline{S}\setminus S'}^{\perp})^\top\overline{\mathbf{G}}_0$ and $\mathbf{R}\overline{\mathbf{G}}_{S\setminus S'} = \mathbf{R}_1(\mathbf{G}_0^{\perp})^\top\overline{\mathbf{G}}_{S\setminus S'}$ are independent and then $\begin{pmatrix} \overline{\mathbf{G}}_{S\setminus S'} & \overline{\mathbf{G}}_0\mathbf{\Gamma} \\ \mathbf{R}\overline{\mathbf{G}}_{S\setminus S'} & \mathbf{R}\overline{\mathbf{G}}_0\mathbf{\Gamma} \end{pmatrix}$ is correctly distributed.

2.2 Pairing-based Quasi-Arguments

Paneth and Rothblum [44] and then Kalai et al. [34] used a weakened version of an argument of knowledge called quasi-argument, as an intermediate step for obtaining a delegation scheme. Quasi arguments are defined for languages that can be expressed as a set of *local constraints*. Roughly speaking, this means that a witness $\boldsymbol{w}$ for membership of a statement x in a language can be decomposed in parts, namely $\boldsymbol{w} = (w_1, \ldots, w_n)$, and for each subset $S \subseteq [n]$, the partial witness $\boldsymbol{w}_S$ satisfies some local relations, that is, a predicate $\mathcal{R}(x, \boldsymbol{w}_S)$ holds. For example, in the case of a CNF formula of n variables, the witness is an accepting assignment of the formula and a local constraint with respect to some set S captures that every clause that only has variables w_i, w_j, w_k for $i, j, k \in S$ is satisfied. Note that it can be the case that even unsatisfiable formulas can satisfy all local constraints for families of sets of small size (yet, no global satisfying assignment exists).

Unlike an argument of knowledge, a quasi-argument has only local extraction, meaning that only a small part of the witness of size at most K, the locality parameter, is extracted. This is formalized by means of an extractor which on input a set $S \subseteq [n]$ of size at most K, where n is the size of the witness, programs a crs so that it can later extract positions of the witness defined by S. Central to quasi-arguments is the notion of no-signaling local extraction which is aimed to capture a strong *local soundness* guarantee.

Local soundness requires that the extracted local witness is consistent with the relation and doesn't lead to a local contradiction, that is, it satisfies the local constraints associated to some set S. The *no-signaling* requirement is defined for any two sets S, S' where $S' \subseteq S$ and of size at most K. It states that the result of programming extraction for S and then output only the extracted value for

S', should be indistinguishable from the result of programming extraction for S' and output the extracted value for S'. Intuitively, this strengthens locality by requiring that the small parts of the local witness are extracted independently from rest.

We next outline the construction of pairing-based quasi-arguements for two specific languages of interest, satisfiability of linear and quadratic relations on committed values. For ease of presentation we do so for symmetric bilinear groups but we streess out that we also translate these to the more efficient setting of asymetric bilinear groups. We will later rely on these quasi arguments to construct a delegation scheme for polynomial sized arithmetic circuits but we emphasize that these constructions are of independent interest; they capture a form of "succinct" aggregation of relations and -importantly- they do so under standard falsifiable assumptions. While full knowledge soundness is not achieved, the weakened notion of no-signaling extraction might be enough for some applicaitons. Thus, we choose to present them in full generality.

Preliminaries In this section we introduce some necessary preliminaries for the construction of the quasi arguments for linear and quadratic relations. First, we introduce the Matrix and Kernel Diffie-Hellman [17, 42] assumption families. Then we introduce Quasi-Adaptive NIZK [30] and sketch the QA-NIZK construction for membership in linear spaces of [39] and finally the knowledge transfer arguments introduced in [24] which allow to construct QA-NIZK under falsifiable assumptions in some more restricted setting.

Cryptographic assumptions. We introduce informally the Matrix and Kernel Diffie-Hellman assumptions [17, 42]. These are natural generalizations of assumptions used in group based cryptography (either with pairings or not). Both assumption families are parametrized by distributions over matrices in $\mathbb{Z}_p$, that is, we consider distribution ensembles $\mathcal{D}_{\ell,k}$ that output matrices in $\mathbb{Z}_p^{\ell\times k}$. When $\ell = k+1$ we simply write $\mathcal{D}_k$.

The $\mathcal{D}_{\ell,k}$-Matrix Diffie-Hellman Assumption ($\mathcal{D}_{\ell,k}$-MDDH) states that elements in the image of a matrix $\mathbf{A}$ sampled from $\mathcal{D}_{\ell,k}$ are computationally indistinguishable from uniformly random elements.

Assumption. (Informal) *$\mathcal{D}_{\ell,k}$-MDDH holds in $\mathbb{G}$ if the distributions $\{[\mathbf{A}], [\mathbf{A}\boldsymbol{w}]\}$ and $\{[\mathbf{A}], [\boldsymbol{z}]\}$ are computationally indisthinguishable, where $\boldsymbol{w}, \boldsymbol{z}$ are random elements of $\mathbb{Z}_p^k$ and $\mathbb{Z}_p^\ell$ respectively, and $\mathbf{A} \leftarrow \mathcal{D}_{\ell,k}$.*

Consider the uniform distribution $\mathcal{U}_{2,1}$ that outputs random elements in $\mathbb{Z}_p^{2\times 1}$. It is easy to assert that the $\mathcal{U}_{2,1}$-MDDH assumption is equivalent to the Decisional Diffie-Hellman assumption in $\mathbb{G}$.[10] In the setting of symmetric bilinear groups -where the DDH assumption does not hold- we consider a slightly stronger assumption, namely the Decisional Linear assumption (DLIN) [7]. This assumption

[10] In fact, the assumption is weaker since we implicitly assume a uniformly distributed generator of $\mathbb{G}$, which need not be the case for DDH. To show that it is weaker, it is enough to note that one can randomize a DDH instance.

can be stated as the $\mathcal{L}_{3,2}$-MDDH assumption, where $\mathcal{L}_{3,2}$ is the distribution

$$\mathcal{L}_{3,2} = \left\{ \begin{pmatrix} a_1 & 0 \\ 0 & a_2 \\ 1 & 1 \end{pmatrix} \;\middle|\; a_1, a_2 \leftarrow \mathbb{Z}_p \right\}$$

The $\mathcal{D}_{\ell,k}$-Kernel Diffie-Hellman Assumption is a natural computational analogue of the $\mathcal{D}_{\ell,k}$-MDDH for bilinear groups. The assumption states that it is infeasible to find non-trivial elements of the co-kernel of $\mathbf{A} \leftarrow \mathcal{D}_{\ell,k}$ given $[\mathbf{A}]$.

Assumption. (Informal) *$\mathcal{D}_{\ell,k}$-MDDH holds in $\mathbb{G}$ if it is computationally hard to find a non-zero element $[\boldsymbol{z}] \in \mathbb{G}^{\ell}$ such that $[\boldsymbol{z}^\top \mathbf{A}]_T = [\mathbf{0}]_T$ given $[\mathbf{A}]$, where $\mathbf{A} \leftarrow \mathcal{D}_{\ell,k}$.*

Note that the assumption is efficiently falsifiable since we can check the winning condition by employing the pairing operation, that is check if $e([\boldsymbol{z}]^\top, [\mathbf{A}]) = [\mathbf{0}]_T$. This assumption family abstracts and generalizes various computational assumptions in bilinear group, such as the Simultaneous Double Pairing Assumption [2].

It is well known that $\mathcal{D}_{\ell,k}$-MDDH implies $\mathcal{D}_{\ell,k}$-Kernel Diffie-Hellman assumption. Intuitively, this holds since if we can sample an element $\boldsymbol{r}$ in the co-kernel of $\mathbf{A}$, it always holds that $\boldsymbol{r}^\top \mathbf{A} \boldsymbol{w} = \mathbf{0}$ while for a uniformly distributed vector $\boldsymbol{z}$, with overwhelming probability $\boldsymbol{r}^\top \boldsymbol{z} \neq 0$, which translates to an efficient distinguisher for the two distributions defined by $\mathcal{D}_{\ell,k}$-MDDH assumption.

Quasi-Adaptive NIZK for membership in linear spaces. Quasi-Adaptive NIZK (QA-NIZK)[11] arguments are NIZK arguments where the CRS is allowed to depend on the specific language for which proofs have to be generated [30]. We are interested in the specific language of membership in linear spaces. Specifically, given a matrix $\mathbf{M}$ and a description of a group gk, we consider the language of vectors of group elements that lie in the image of $\mathbf{M}$, that is,

$$\mathcal{L}_{gk,\mathbf{M}} = \{[\boldsymbol{x}] \mid \exists \boldsymbol{w} \text{ s.t. } \boldsymbol{x} = \mathbf{M}\boldsymbol{w}\}$$

In the quasi-adaptive case, we allow the common reference string to depend on $gk, \mathbf{M}$ but an adversary can choose the statement $[\boldsymbol{x}]$ adaptively. There are very efficient constructions in this setting. We briefly describe the construction of Kiltz and Wee [39]. First we consider the designated verifier case. Let $\mathbf{M}$ be an $\ell \times n$ matrix. The construction is essentially a hash proof system [14]. The crs contains the projection $[\mathbf{B}] = [\mathbf{M}^\top \mathbf{K}]$ for a random secret key $\mathbf{K} \in \mathbb{Z}_p^{\ell \times k}$. To prove a statement $[\boldsymbol{x}] = [\mathbf{M}]\boldsymbol{w}$, the prover sends $[\boldsymbol{\pi}] = \boldsymbol{w}^\top [\mathbf{B}]$ and the verifier asserts that $[\boldsymbol{\pi}] = [\boldsymbol{x}]^\top \mathbf{K}$. Now it is easy to see that this simple protocol is complete. Indeed

$$\boldsymbol{\pi} = \boldsymbol{w}^\top [\mathbf{B}] = \boldsymbol{w}^\top \mathbf{M}^\top \mathbf{K} = \boldsymbol{x}^\top \mathbf{K}$$

For soundness, roughly speaking, the value $\boldsymbol{x}^\top \mathbf{K}$ is random for $\boldsymbol{x}$ that does not belong to the image of $\mathbf{M}$ conditioned on $\mathbf{B}$. Thus, a cheating (even unbounded)

[11] In this work we do not need the zero knowledge property so we omit it from the discussion.

prover has only negligible probability of producing a verifying proof for elements not in the image of $\mathbf{M}$.

To make the scheme publicly verifiable, groups equipped with a bilinear map are employed. To enable the verifier to perform the test without knowing the secret $\mathbf{K}$, we also add to the crs the value $[\mathbf{C}] = [\mathbf{KA}]$, where $\mathbf{A}$ is a matrix that satisfies some hardness condition. Now, the verifier can test $e([\boldsymbol{\pi}], [\mathbf{A}]) = e([\boldsymbol{x}^\top], [\mathbf{C}])$. Note that this corresponds to multiplying the verification equation of the designated verifier case from the right with $\mathbf{A}$. Now, if

(1) the designated verifier relation does not hold, namely, $\boldsymbol{\pi} \neq \boldsymbol{x}^\top \mathbf{K}$ and
(2) the proof verifies, namely $\boldsymbol{\pi}\mathbf{A} = \boldsymbol{x}^\top \mathbf{KA}$,

then $[\boldsymbol{\pi}] - [\boldsymbol{x}^\top]\mathbf{K}$ is a non-trivial element in the co-kernel of $[\mathbf{A}]$. Thus, the publicly verifiable scheme is sound if we additionally assume that $\mathbf{A}$ is sampled by a distributions $\mathcal{D}$ such that the $\mathcal{D}$-Kernel Diffie-Hellman assumption holds.

Note that if $\mathbf{M}$ spans the entire linear space, then the language is trivial. In this case, only knowledge soundness is a meaningful property. However, we do not whether knowledge soundness of this construction can be proven under falsifiable assumptions or not.

Knowledge Transfer Arguments. To achieve succinct arguments, in principle, one needs to use shrinking commitments. When trying to use such commitments with QA-NIZK such as [39], the aforementioned "triviality" problem arises and it seems like one has to resort to non-falsifiable assumptions or the generic group model. Motivated by the problem of constructing delegation schemes under falsifiable assumptions and in order to overcome the above issue, [24] relax the knowledge soundness property.

When considering delegation using the natural approach of (deterministically) committing to the wires of the circuit, one can observe that full knowledge soundness seems to be an unnecessarily strong requirement. Indeed, given the input $\boldsymbol{x}$ of the circuit, one can compute (or verify) these commitments efficiently by evaluating the circuit. This means intuitively, that we already know how a "correct" opening of the commitments looks like in the soundness security reduction. [24] exploits this fact and manages to relax the knowledge soundness requirement by considering statements of the form "if commitment $[\boldsymbol{c}]$ opens to $\boldsymbol{w}$, then commitment $[\boldsymbol{d}]$ opens to $f(\boldsymbol{w})$" for publicly known function f. As we shall see later, they show that this notion of soundness is enough to construct delegation for low-depth circuits. They also construct two knowledge transfer arguments for linear and quadratic relations under falsifiable assumptions. More concretely, they consider statements of the form

- "if $[\boldsymbol{c}]$ opens to $\mathbf{M}\boldsymbol{w}$, then $[\boldsymbol{d}]$ opens to $\mathbf{N}\boldsymbol{w}$ for some publicly known $\mathbf{M}, \mathbf{N}$, and
- "if $[\boldsymbol{c}_1]$ opens to $\boldsymbol{w}_1$ and $[\boldsymbol{c}_2]$ opens to $\boldsymbol{w}_2$, then $[\boldsymbol{d}]$ opens to $\boldsymbol{w}_1 \circ \boldsymbol{w}_2$ where $\circ$ denotes the pairwise product of vectors.

In the soundness definition, the adversary is required to output the valid opening along with the statement proof-pair. We emphasize that this is only part

of the soundness definition and in the protocol execution the prover does not have to output the valid opening. Consider for example the first case for linear relations. An adversary wins if it manages to output a statement $[\boldsymbol{c}], [\boldsymbol{d}]$ with an accepting proof *and* a $\boldsymbol{w}$ such that $[\boldsymbol{c}] = [\mathbf{M}]\boldsymbol{w}$ *but* $[\boldsymbol{d}] \neq [\mathbf{N}]\boldsymbol{w}$. Such statements essentially give the guarantee that some a priori knowledge about a commitment is "correctly" transferred to another commitment.

For the former construction, namely linear relations, they use the [39] construction where they define $\mathbf{M}$ as a two block matrix where the upper part corresponds to $[\boldsymbol{c}]$ and the lower to $[\boldsymbol{d}]$. Now, using [39], the prover simply needs to convince the verifier that $\begin{bmatrix}\boldsymbol{c}\\ \boldsymbol{d}\end{bmatrix} = \begin{bmatrix}\mathbf{M}\\ \mathbf{N}\end{bmatrix}\boldsymbol{w}$. They show that this construction is knowledge transfer sound if the upper matrix $\mathbf{M}$ is sampled from a distribution $\mathcal{D}$ for which the $\mathcal{D}$-MDDH assumption holds.

For proving the quadratic relations, they do a different analysis of standard techniques used for the construction of pairing-based succinct arguments that exploit the properties of the Lagrange basis.

They also modify these constructions to be compatible with the more efficient setting of asymmetric bilinear groups, under the natural modifications of the required assumption for asymmetric groups.

Oblivious Trapdoor Generation for Quasi-Arguments Similar to the case of no-signaling SSB commitments we define a storngerand easier to work with (in our context) notion that implies the no-signaling property of quasi arguments, *oblivious trapdoor generation.*

We require that there exists an *oblivious* key generation algorithm that takes as input (1) a crs_S that allows extraction for a set S, and (2) the description of a subset $S' \subseteq S$, and generates a $\mathsf{crs}_{S'}$ for some set S' *and* a trapdoor[12] for extracting local witnesses associated to the set S' *obliviously* of $S \setminus S'$. We emphasize that the oblivious trapdoor generation algorithm knows neither the description of S nor any information about the trapdoor associated with it. We require that the new crs is *statistically close* to the crs_S given as input. The fact that this property implies no-signaling commitments is identical to the case of SSB commitments.

Quasi-Arguments of Membership in a Linear Space We define a quasi-argument of knowledge of some vector $[\boldsymbol{x}] \in \mathbb{G}^\ell$ belonging to the image of a matrix $[\mathbf{U}] \in \mathbb{G}^{\ell \times n}$, where $\boldsymbol{x}$ is committed using an SSB commitment. Consider a commitment $[\boldsymbol{c}]$ that is statistically binding on the set S. We show that there exists a local and no-signaling extractor which, given some $S \subseteq [n]$ of size $t \leq K$, extracts $[\boldsymbol{x}_S] \in \mathrm{Im}([\mathbf{U}_S])$, where $\boldsymbol{x}_S \in \mathbb{Z}_p^t$ is the vector whose entries are x_i and

[12] We modify the quasi-argument defintion of [34] to admit a fixed extractor algorithm that takes as input the statement-proof pair of the adversary, and additionally some secret state produced during the crs generation, -the trapdoor- and extracts the local witness.

$\mathbf{U}_S \in \mathbb{Z}_p^{t\times n}$ is the matrix whose rows are the rows of $\mathbf{U}$ indexed by i, where i ranges over S in some fixed order. A local constraint $[\boldsymbol{x}_S]$ associated with the set S can be interpreted as satisfying two properties:

(1) $[\boldsymbol{x}_S]$ is consistent with the commitment $[\boldsymbol{c}]$, namely the (uniqe) S-opening of $[\boldsymbol{c}]$ is x_S, and
(2) $[\boldsymbol{x}_S]$ is in the image of $[\mathbf{U}_S]$.

We use the Kiltz and Wee argument of membership in linear spaces [39] to construct a quasi argument for linear relations. Details follow.

The argument. Our construction is Kiltz and Wee linear membership argument [39] for the matrix $[\mathbf{GU}]$, where $\mathbf{G}$ is an SSB commitment key with locality parameter K. For completeness, we describe the protocol for this specific matrix. We note that we present the scheme with proof size $k+1$ of [39], where k is a parameter of the scheme defined by the underlying assumption, but our construction is also sound for the more efficient instantiation of size k. In any case, we emphasize that the parameter is a small constant ($k = 2$).

Let's recall the construction for the matrix $\mathbf{M} = \mathbf{GU}$. The crs contains $[\mathbf{B}] = [\mathbf{U}^\top\mathbf{G}^\top\mathbf{K}]$ and $[\mathbf{C}] = [\mathbf{KA}]$ for some random hash key $\mathbf{K}$ and $\mathbf{A}$ drawn from some distribution satisfying a kernel assumption. A proof is computed as $[\boldsymbol{\pi}] = \boldsymbol{w}^\top[\mathbf{B}]$, and verification is done by checking if $e([\boldsymbol{\pi}], [\mathbf{A}]) = e([\boldsymbol{c}^\top], [\mathbf{C}])$.

Local and No-Signaling extraction. Our strategy to prove local soundness is to show that, apart from extracting $[\boldsymbol{x}_S]$ from $[\boldsymbol{c}]$, we are also able to produce a verifying proof $[\boldsymbol{\pi}^\dagger]$ that $[\boldsymbol{x}_S] \in \mathbf{Im}(\mathbf{U}_S)$. More concretely, on input a crs $\mathsf{crs}_S = ([\mathbf{A}^\dagger], [\mathbf{B}^\dagger], [\mathbf{C}^\dagger])$ for membership in the linear space of $\mathbf{U}_S$, we can construct another crs that is statistically close to the quasi argument crs for $\mathbf{U}$ and, more importantly, we can extract a local opening $[\boldsymbol{x}_S]$ *and* a proof $[\boldsymbol{\pi}^\dagger]$ satisfying the verification equation for crs_S.

We embed the public parameters $[\mathbf{A}^\dagger], [\mathbf{B}^\dagger], [\mathbf{C}^\dagger]$ of the local linear space argument for $\mathbf{U}_S$ in the quasi argument parameters. Although the secret hash key $\mathbf{K}^\dagger$ of the local linear argument is statistically hidden, we can still pick a random hash key for all the coordinates by picking another secret key and implicitly define the full secret key as some composition of the two keys. Concretely, given the trapdoor $\mathbf{T}_S$ for locally opening SSB commitments we implicitly define $\mathbf{K} = \mathbf{T}_S\mathbf{K}^\dagger + \mathbf{R}$, where $\mathbf{R}$ is the additional key, so that the proofs for $\boldsymbol{c} = \mathbf{GP}\left(\begin{smallmatrix}\boldsymbol{x}_S\\ \boldsymbol{x}_{\overline{S}}\end{smallmatrix}\right) = \mathbf{G}_S\boldsymbol{x}_S + \mathbf{G}_{\overline{S}}\boldsymbol{x}_{\overline{S}}$ are of the form $\boldsymbol{\pi} = \boldsymbol{c}^\top\mathbf{K} = (\mathbf{G}_S\boldsymbol{x}_S + \mathbf{G}_{\overline{S}}\boldsymbol{x}_{\overline{S}})^\top(\mathbf{T}_S\mathbf{K}^\dagger + \mathbf{R}) = \boldsymbol{x}_S^\top\mathbf{K}^\dagger + \boldsymbol{c}^\top\mathbf{R}$. In this way a proof for the local argument can be retrieved as $[\boldsymbol{\pi}^\dagger] = [\boldsymbol{\pi}] - [\boldsymbol{c}^\top]\mathbf{R}$. This equivalent way of sampling $\mathbf{K}$ allows to compute the crs of the larger linear argument using only $[\mathbf{A}^\dagger], [\mathbf{B}^\dagger], [\mathbf{C}^\dagger]$ and $\mathbf{T}_S, \mathbf{R}$. Indeed, we can define $[\mathbf{A}] = [\mathbf{A}^\dagger]$, $[\mathbf{B}] = [\mathbf{B}^\dagger] + [\mathbf{U}^\top\mathbf{G}^\top]\mathbf{R}$ and $[\mathbf{C}] = \mathbf{T}_S[\mathbf{C}^\dagger] + \mathbf{R}[\mathbf{A}^\dagger]$.

We also show that the crs is indistinguishable for different sets and that there is an oblivious trapdoor generation strategy, and hence we also have a no-signaling extraction strategy. The indistinguishability of the crs follows directly from the indistinguishability of SSB commitment keys; it is enough to note that only the commitment key depends on S and all other values can be

efficiently computed given only the commitment key[13]. For oblivious trapdoor generation, we use the fact that we can sample an identically distributed commitment key along with a trapdoor -this follows by the oblivious key generation of the commitment scheme- and then we argue in the same way as before: given the commitment key we can sample the rest of crs honestly.

Extension to Knowledge Transfer, Bilateral Spaces and Sum Arguments. We also construct variations of the above protocol, specifically a knowledge transfer version based on [24] and two construction suitable for asymmetric bilinear groups.

First we consider the knowledge transfer construction. We first describe the local constraints. Consider two matrices $[\mathbf{M}], [\mathbf{N}]$, and two commitment keys $[\mathbf{G}], [\mathbf{H}]$ statistically binding at S. The statement consists of two commitments $[\boldsymbol{c}], [\boldsymbol{d}]$. For the local extraction guarantee w.r.t. set S we require that, given an accepting proof π *and* an opening $\boldsymbol{w}$, we can extract values $[\boldsymbol{x}_S]$, $[\boldsymbol{y}_S]$ such that

(1) $[\boldsymbol{x}_S]$, $[\boldsymbol{y}_S]$ are the unique S-openings of $[\boldsymbol{c}], [\boldsymbol{d}]$ w.r.t. commitment keys $\mathbf{G}, \mathbf{H}$ respectively, and
(2) if $[\boldsymbol{x}_S] = [\mathbf{M}_S]\boldsymbol{w}$, then $[\boldsymbol{y}_S] = [\mathbf{N}_S]\boldsymbol{w}$.

The construction and the analysis are identical to the previous case. We simply use the [39] construction for the matrix with upper part $\mathbf{GM}$ and lower part $\mathbf{HN}$. The only difference in the analysis is on the local extraction case. We argue that we can extract an accepting proof for a crs for the language of linear knowledge transfer for the matrices $\mathbf{M}_S, \mathbf{N}_S$ and, thus, we also require that the $\mathcal{M}_S^\top$-MDDH assumption holds for every S, where $\mathcal{M}_S$ is the distribution from which we sample $\mathbf{M}_S$.

Finally, we also consider constructions in asymmetric bilinear groups. A variant of the linear subspace QA-NIZK argument given in [22], and extended to knowledge transfer arguments in [24], considers the statement as well as the matrix split between the two groups. We call this argument a linear argument for bilateral spaces. We also consider a particular type of argument for bilateral linear spaces defined in [22] and called "sum in subspace argument". In this case, the statement is $[\boldsymbol{x}]_1, [\boldsymbol{y}]_2$ and soundness captures that $\boldsymbol{x}+\boldsymbol{y} \in \mathrm{Im}(\mathbf{M}+\mathbf{N})$ given $[\mathbf{M}]_1, [\mathbf{N}]_2$ in the two different source groups. We construct quasi arguments for all these variants with knowledge transfer soundness. Luckily, the constructions as well as the security proofs are minor modifications of the original argument.

Quasi-Argument of Hadamard Products The next quasi arguement construction shows that some vector $\boldsymbol{c}$ is the Hadamard product of two vectors $\boldsymbol{a}, \boldsymbol{b}$, namely $\boldsymbol{c} = \boldsymbol{a} \circ \boldsymbol{b}$. We can naturally define the local constraints here as $\boldsymbol{c}_S = \boldsymbol{a}_S \circ \boldsymbol{b}_S$ for every set $S \subseteq [n]$, where n is the dimension of the vectors. As

[13] Here, we assume the distribution $\mathcal{U}$ that outputs the matrix $[\mathbf{U}]$ is witness samplable, meaning that during sampling, we can also sample the discrete logarithms of $[\mathbf{U}]$ which is usually the case. In this work, we only consider such distributions.

in the linear case, we care about committed values, that is, the vectors $\boldsymbol{a}, \boldsymbol{b}, \boldsymbol{c}$ are committed and we claim that the openings satisfy the claimed relation.

Our starting point is the "bit-string" argument of [22]. We observe that it is implicitly a quasi-argument with locality parameter $K = 1$ for the set of equations $b_i(b_i - 1) = 0$ for all $i \in [n]$. Next we describe this construction and after that we show it indeed satisfies the no-signaling local soundness property. It will be convenient to directly work with equations of the form $x_i y_i = z_i$ instead of the bit-string argument equations.

The common reference string in [22] contains what we interpret as three SSB commitment keys $[\mathbf{G}], [\mathbf{H}], [\mathbf{F}]$ with locality parameter $K = 1$. It additionally includes the product $[\mathbf{G} \otimes \mathbf{H}]$. The prover gives three commitments $[\boldsymbol{a}], [\boldsymbol{b}], [\boldsymbol{c}]$ w.r.t. $\mathbf{G}, \mathbf{H}, \mathbf{F}$ and claims that the openings satisfy the Hadamard relation. We first note that it is easy to construct an arguement for a related language. Consider the elements $\mathbf{G} \otimes \mathbf{H}$ as a commitment key. The prover can give a commitment to the Kronecker product $\boldsymbol{z} = \boldsymbol{a} \otimes \boldsymbol{b}$ by computing $[\boldsymbol{t}] = [\mathbf{G} \otimes \mathbf{H}]\boldsymbol{z}$. The verifier can then use the pairing to verify the Kronecker product relation, namely it tests that $e([\boldsymbol{c}], [\boldsymbol{d}]) = e([\boldsymbol{t}], [1])$ where $[\boldsymbol{c}] = [\mathbf{G}]\boldsymbol{a}, [\boldsymbol{d}] = [\mathbf{H}]\boldsymbol{b}$ are commitment to some vectors and are part of the statement. Some simple calculations show that

$$\boldsymbol{c}\boldsymbol{d} = \boldsymbol{c} \otimes \boldsymbol{d} = \mathbf{G}\boldsymbol{a} \otimes \mathbf{H}\boldsymbol{b} = (\mathbf{G} \otimes \mathbf{H})(\boldsymbol{a} \otimes \boldsymbol{b}) = \boldsymbol{t}$$

The Kronecker product commitment $\boldsymbol{t}$ is included as part of the proof. Now, from this simple Kronecker product argument, it is easy to prove the Hadamard product. It is enough to note that the Hadamard product is a linear funciton of the Kronecker product, thus, the prover and verifier can use the protocol for linear relations of the previous section.

Local and No-Signaling Extraction. The crucial observation to prove local extraction is that if $\mathbf{G}, \mathbf{H}$ are extractable in one position, say i, j respectively, then $\mathbf{G} \otimes \mathbf{H}$ is extractable at position $n(i-1) + j$. More concretely, lettting $\mathbf{T_G}$, $\mathbf{T_H}$ be the trapdoors for $\mathbf{G}, \mathbf{H}$ respectively, the trapdoor for the commitment key $\mathbf{G} \otimes \mathbf{H}$ is simply $\mathbf{T_G} \otimes \mathbf{T_H}$. Some straightforward calculations reveal that applying this trapdoor to a commitment with the key $\mathbf{G} \otimes \mathbf{H}$ indeed yields the $n(i-1) + j$-th coordinate of the committed value, which is uniquely defined. In fact, we generalize this for larger locality parameters and we also show that, for some distributions of commitment keys, the no-signaling/oblivious trapdoor generation properties hold if they hold for $\mathbf{G}$, $\mathbf{H}$.

Consider the simple case of $K = 1$ and let all three commitments $\mathbf{G}, \mathbf{H}, \mathbf{F}$ be extractable at the same position i. We show that we can extract local openings $[x_i] = \mathbf{T_G}[\boldsymbol{a}], [y_i] = \mathbf{T_H}[\boldsymbol{b}], [z_i] = \mathbf{T_F}[\boldsymbol{c}]$ as well as $[w_i] = \mathbf{T_{G \otimes H}}[\boldsymbol{t}]$ such that $z_i = x_i y_i$. Assume for the sake of a contradiction that $z_i \neq z_i' = a_i b_i$. Since the columns $\boldsymbol{g}_i, \boldsymbol{h}_i, \boldsymbol{f}_i$ are linearly independent from the other columns in $\mathbf{G}, \mathbf{H}, \mathbf{F}$, respectively, if the commitments $[\boldsymbol{c}], [\boldsymbol{d}], [\boldsymbol{t}]$ satisfies $[\boldsymbol{c}] \otimes [\boldsymbol{d}] = e([\boldsymbol{t}], [1])$, then the unique openings at coordinate i satisfy $z_i = x_i y_i$. Now, if $z_i \neq z_i'$, the linear relation does not hold and we can break the underlying QA-NIZK for membership in linear spaces.

For oblivious trapdoor generation, it is enough to note that if the commitment key satisfies this property, so does the above constructions. Indeed, note that using the commitment key, it is enough to produce a crs for membership in subspace language to create the full crs of the protocol.

Extension to Knowledge Transfer Arguments. We extend the quasi-argument local soundness to offer a "knowledge transfer" guarantee. In this case, we essentially commit to commitments. That is, we use an SSB commitment key to commit to multiple commitments and the local openings are commitments themselves. Namely we extract values $[x_i], [y_i], [z_i]$ which are interpreted as commitments w.r.t. some (not necessarily SSB) commitments keys $\mathbf{U}, \mathbf{V}, \mathbf{W}$. We require that no PPT adversary can produce openings $\boldsymbol{a}, \boldsymbol{b}$ such that $x_i = \mathbf{U}_i\boldsymbol{a}, y_i = \mathbf{V}_i\boldsymbol{b}$ but $z_i \neq \mathbf{W}_i\boldsymbol{a} \circ \boldsymbol{b}$. The constraint language for a set S is parametrized by SSB commitments $\mathbf{G}, \mathbf{H}, \mathbf{F}$ binding at S as well as some matrices $\mathbf{U}, \mathbf{V}, \mathbf{W}$. We require that given an accepting proof π for a statement $[\boldsymbol{c}], [\boldsymbol{d}], [\boldsymbol{f}]$ *and* openings $\boldsymbol{a}, \boldsymbol{b}$, we can extract values $[\boldsymbol{x}_S]$, $[\boldsymbol{y}_S]$, $[\boldsymbol{z}_S]$ such that

(1) $[\boldsymbol{x}_S]$, $[\boldsymbol{y}_S]$, $[\boldsymbol{z}_S]$ are the unique S-openings of $[\boldsymbol{c}], [\boldsymbol{d}]$, $[\boldsymbol{f}]$ w.r.t. commitment keys $\mathbf{G}, \mathbf{H}, \mathbf{F}$ respectively, and
(2) if $[\boldsymbol{x}_S] = [\mathbf{U}_S]\boldsymbol{a}$ and $[\boldsymbol{y}_S] = [\mathbf{V}_S]\boldsymbol{b}$, then $[\boldsymbol{z}_S] = [\mathbf{W}_S]\boldsymbol{a} \circ \boldsymbol{b}$.

One might wonder at this point how we commit to commitments which naturally requires multiplication of group elements which is assumed computationally hard. To achieve that, we simply include in the crs the products $[\mathbf{GU}], [\mathbf{HV}], [\mathbf{FW}]$. Now, we can commit to the n commitments $\mathbf{U}_i\boldsymbol{a}$ as $[\mathbf{GU}]\boldsymbol{a}$ and similarly for the other keys.

The knowledge transfer version is essentially the same as in the previous case. The only difference is that we also need to include some additional elements in the crs to allow to the prover to compute the Kronecker product, namely the values $[\mathbf{Q}] = [(\mathbf{G} \otimes \mathbf{H})(\mathbf{U} \otimes \mathbf{V})]$. As in the previous case, we can then exploit the linear relation between the Hadamard product and the Kronecker product. From a correct commitment $[\mathbf{Q}](\boldsymbol{a} \otimes \boldsymbol{b})$, we can use the linear knowledge transfer to get a commitment to the Hadamard products w.r.t. the third commitment key, namely $[\mathbf{FW}](\boldsymbol{a} \circ \mathbf{b})$. To show this, we first show that the $\mathcal{G} \otimes \mathcal{H}$-MDDH assumption holds if $\mathcal{G}$-MDDH and $\mathcal{H}$-MDDH hold, where $\mathcal{G}, \mathcal{H}$ are the distributions of $\mathbf{G}, \mathbf{H}$ respectively.

We are also able to extend these techniques to work in asymmetric bilinear groups as well. The construction is somewhat technical, but the core idea is to construct SSB commitments suitable for asymmetric groups, where we "split" the commitments between the two groups, and use the bilateral variants of the linear quasi-arguments discussed in the previous sections.

2.3 From our Quasi-Arguments to Delegation.

Using the ideas of [44, 34], we can derive delegation of computation from quasi arguments for languages encoding the computation. The local constraints capture that each step of the computation was done correctly. First, we present the

high level idea for the delegation construction from quasi-arguments. We first show how to delegate low-space TMs/low-width circuits and then we show how to overcome the dependence on space/width.

Delegating bounded space TM/bounded width circuits We first recall the high-level ideas to construct a delegation scheme from quasi arguments of [44, 34] in the simpler case of bounded space computation. Consider some polynomial time sequential computation which on input x outputs y, for example a Turing Machine or an arithmetic circuit. The computation goes through a sequence of states $\mathtt{st}_0, \mathtt{st}_1, \ldots, \mathtt{st}_d$ such that $\mathtt{st}_0$ is consistent with the input, state $\mathtt{st}_d$ contains the output y, and there's a functional relation between states $\mathtt{st}_i, \mathtt{st}_{i+1}$ where $\mathtt{st}_{i+1} = f(\mathtt{st}_i)$ and f is determined by the description of the computation. We first consider the case of bound space computation and discuss later how to remove this constraint. Consider a quasi arguement of locality $K = 2|\mathtt{st}|$ where local constraints require that $\mathtt{st}_i, \mathtt{st}_{i+1}$ are consistent w.r.t. f. The goal is to show that an adversary that makes the quasi-argument verifier accept must (w.o.p) sample x, y such that y is the result of the computation on input x.

We can first "program" the local extractor extractor to extract $\mathtt{st}_0, \mathtt{st}_1$, i.e. use locality parameter $K = 2|\mathtt{st}|$, where $|\mathtt{st}|$ is a bound on the size of the states (i.e. space of the TM or width of the circuit). Local soundness asserts that state $\mathtt{st}_0$ is consistent with x. Local soundness also implies that $\mathtt{st}_1$ is consistent with $\mathtt{st}_0$ and hence with x (note that the statement $\mathtt{st}_1 = f(\mathtt{st}_0)$ depends only on local variables). Now, to show that $\mathtt{st}_2$ is also consistent, we jump to another game where first the extractor computes only $\mathtt{st}_1$, and in the next game the extractor computes $\mathtt{st}_1, \mathtt{st}_2$. The crucial observation is that $\mathtt{st}_1$ should be still consistent with x in both games. Otherwise, we can distinguish between the common output of extractors for $\mathtt{st}_0, \mathtt{st}_1$ and $\mathtt{st}_1$ or between $\mathtt{st}_1$ and $\mathtt{st}_1, \mathtt{st}_2$, which contradicts the no-signaling property. Importantly, we can efficiently compute the "correct" state $\mathtt{st}_1$ since the computation is deterministic, and thus the no-signaling distinguisher discribed is indeed efficient. Similarly, consistency of $\mathtt{st}_1$ and local soundness imply that $\mathtt{st}_2$ is also consistent. Now, we can inductively continue until we reach the last state, $\mathtt{st}_d$, which corresponds to the output of the computatiaon.

Small width circuit delegation from DLIN. Let C be an arithmetic circuit with width w and depth d. We consider the input to correspond to level 0. Without loss of generality, assume that the circuit has w input and w output wires. In this section we consider the width w to be small, or alternatively, efficiency will depend on w.

We follow the circuit arithmetization of [24]. The multiplication gates are partitioned in d levels. Each level groups the gates at the same distance from the inputs, without counting linear gates. In this way, the inputs of level $i + 1$ are linear combinations of outputs of the i previous levels. We can then express this as constraints describing the computation as

$$\boldsymbol{a}_i \circ \boldsymbol{b}_i = \boldsymbol{c}_i \qquad \text{for } i = 1 \text{ to } d, \tag{1}$$

$$\begin{pmatrix} \boldsymbol{a}_{i+1} \\ \boldsymbol{b}_{i+1} \end{pmatrix} = \sum_{0 \le j \le i} \begin{pmatrix} \mathbf{D}_{i,j} \\ \mathbf{E}_{i,j} \end{pmatrix} \boldsymbol{c}_j = \begin{pmatrix} \mathbf{D}_i \ \mathbf{0} \\ \mathbf{E}_i \ \mathbf{0} \end{pmatrix} \boldsymbol{c} \qquad \text{for } i = 0 \text{ to } d-1, \tag{2}$$

$$\boldsymbol{c}_0 = \boldsymbol{x} \in \mathbb{Z}_p^w \text{ and } \boldsymbol{c}_d = \boldsymbol{y} \in \mathbb{Z}_p^w. \tag{3}$$

Vectors $\boldsymbol{a}_i, \boldsymbol{b}_i, \boldsymbol{c}_i$ denote respectively the left, right and output wires of multiplication gates in level i. Matrices $\mathbf{D}_{i,j}, \mathbf{E}_{i,j}$ can be naturally derived from the circuit's linear gates. Equation (1) states the relation between output wires and the input wires of a level of multiplication gates.

Now consider a symmetric bilinear group described by gk and consider three SSB commitments $\mathbf{G}, \mathbf{H}, \mathbf{F}$ with locality $K = |w|$ for committing to wd-dimensional vectors. We publish in the crs the commitment keys and we we also compute two quasi argument crs:

(1) for membership in linear space crs for the matrix $[\mathbf{M}_1] = \begin{bmatrix} \mathbf{F} \\ \mathbf{GD} \\ \mathbf{HE} \end{bmatrix}$. Here, $\mathbf{D}$, $\mathbf{E}$ are the matrices for the linear relations as a whole (note per level). That is, for left and output wires it should hold $\boldsymbol{a} = \mathbf{D}\boldsymbol{c}$, and similarly for right wires.
(2) for hadamard relation for $\mathbf{G}, \mathbf{H}, \mathbf{F}$. Note that, essentially, this corresponds to yet another quasi argument for membership in linear spaces for $[\mathbf{M}_2] = \begin{bmatrix} (\mathbf{G} \otimes \mathbf{H}) \\ \mathbf{F\Delta} \end{bmatrix}$ where $\boldsymbol{\Delta}$ captures the linear relation between the Kronecker and Hadamard product, that is $(\boldsymbol{a} \circ \boldsymbol{b}) = \boldsymbol{\Delta}(\boldsymbol{a} \otimes \boldsymbol{b})$.

The prover gives the commitments to the left, right, output wires, namely $[\boldsymbol{L}] = [\mathbf{G}]\boldsymbol{a}, [\boldsymbol{R}] = [\mathbf{H}]\boldsymbol{b}, [\boldsymbol{O}] = [\mathbf{F}]\boldsymbol{c}$. Note that these commitments are of size $\mathcal{O}(\mathsf{poly}(\kappa)w)$ but independent of d. Next, it proves that $[\mathbf{O}], [\mathbf{L}], [\mathbf{R}]$

- lie in the image of $[\mathbf{M}_1]$ using the witness $\boldsymbol{c}$.
- satisfy the Hadamard relations. To do so, it computes a commitment $[\boldsymbol{Z}] = [(\mathbf{G}\otimes\mathbf{H})](\boldsymbol{a}\otimes\boldsymbol{b})$ and shows using the linear argument that the vector $\left[\begin{pmatrix} \boldsymbol{Z} \\ \boldsymbol{O} \end{pmatrix}\right]$ lies in the image of $\mathbf{M}_2$ using the witness $\boldsymbol{a} \otimes \boldsymbol{b}$.

The verifier checks that (1) the linear proofs verify and (2) that $e([\boldsymbol{L}], [\boldsymbol{R}]) = e([\boldsymbol{Z}], [1])$. It also does some additional input/output consistency check which we omit for now and describe next.

Now, let's see the core of the extraction argument. The inductive claim goes as follows: If we set $[\mathbf{F}]$ extractable for the i-th level, namely we the set $S_i = \{iw + 1, \ldots, (i+1)w\}$, then -conditioned on an accepting proof- extracting the level i-th level wires corresponds to the correct values $[\boldsymbol{c}_i]$ w.r.t. the input $\boldsymbol{c}_0$. We will handle the base case later when we discuss input/output consistency. For the inductive step, assume the statement is true for i. We show that it is true for $i + 1$. We proceed as follows:

(1) We first set $\mathbf{G}, \mathbf{H}$ extractable at set S_{i+1} corresponding to the $i+1$-th level in addition to the $\mathbf{F}$ extractable at S_i. By the no-signaling guarantees the value $[\boldsymbol{c}_i]$ extracted by $[\boldsymbol{O}]$ is still correct.
(2) By the local soundness of the linear quasi argument, the extracted values $[\boldsymbol{c}_i], [\boldsymbol{a}_{i+1}], [\boldsymbol{b}_{i+1}]$ must lie in the image of the submatrix of $\mathbf{M}_1$ corresponding to these values. This matrix contains the blocks $\mathbf{I}, \mathbf{D}_{i+1}, \mathbf{E}_{i+1}$. Hence the values extracted correspond to the correct values $[\boldsymbol{a}_{i+1}], [\boldsymbol{b}_{i+1}]$ w.r.t the input $\boldsymbol{c}_0$.
(3) We only set $\mathbf{G}, \mathbf{H}$ extractable at set S_{i+1} and leave $\mathbf{F}$ extractable at the empty set. By the no-signaling guarantees the extracted wires for left and right values $[\boldsymbol{a}_{i+1}], [\mathbf{b}_{i+1}]$ are still correct.
(4) In addition to $\mathbf{G}, \mathbf{H}$ extractable at set S_{i+1}, we set $\mathbf{F}$ extractable at S_{i+1}. Now we argue about local constraint of the Hadamard product. We proceed in two steps:
   - By the pairing test $e([\boldsymbol{L}], [\boldsymbol{R}]) = e([\boldsymbol{Z}], [1])$ and the assumption that $[\boldsymbol{a}_{i+1}]$, $[\boldsymbol{b}_{i+1}]$ are correct we get that

$$\mathbf{T}_{\mathbf{G}}\boldsymbol{L} \otimes \mathbf{T}_{\mathbf{H}}\boldsymbol{R} = (\mathbf{T}_{\mathbf{G}} \otimes \mathbf{T}_{\mathbf{H}})(\boldsymbol{L} \otimes \boldsymbol{R}) = (\mathbf{T}_{\mathbf{G}} \otimes \mathbf{T}_{\mathbf{H}})\boldsymbol{Z} = \mathbf{T}_{\mathbf{G}\otimes\mathbf{H}}\boldsymbol{Z}$$

     which implies that $\boldsymbol{z}_{i+1} = \boldsymbol{a}_{i+1} \otimes \boldsymbol{b}_{i+1}$. This means that the extracted value of the Kronecker commitment corresponds to the Kronecker product $\boldsymbol{a}_{i+1} \otimes \boldsymbol{b}_{i+1}$ of left and right wires in level $i+1$.
   - Working similarly to the step (2), we get that the extracted values $\boldsymbol{Z}_{i+1}, \boldsymbol{O}_{i+1}$ live in the image of $\mathbf{M}_2$. It should then be the case that we extract $[\boldsymbol{c}_{i+1}]$ which is the Hadamard product $\boldsymbol{a}_{i+1} \circ \boldsymbol{b}_{i+1}$. This correspond to the correct assignment of output wires in level $i+1$.
(5) Finally, we only set $\mathbf{F}$ extractable at set S_{i+1} and leave $\mathbf{G}, \mathbf{H}$ extractable at the empty set. By the no-signaling guarantees the extracted value $[\boldsymbol{c}_{i+1}]$ is still correct.

We note that proving this is technically more involved. We need to show that the quasi arguments can be composed well, and they still satisfy the no-signaling properties despite the fact that they share commitment keys. Equivalently one could define and analyze a unified quasi argument to directly work with the circuit "transition funciton". In any case, we omit these details from these technical overview.

Input/Output Consistency. We modify the commitment $\mathbf{F}$ by making it trivially extractable at the input/output levels $0, d$ always, regardless of the extraction set. That is, we "use" the identity matrix $\mathbf{I}_w$ for committing to the output wires at the first and last level. This corresponds to augmenting $\mathbf{F}$ with some identity rows. Thus, the verifier can always trivially check the consistency with input/output. Note that the final commitment size grows by $2|w|$, the size of input and output, but these values are part of the statement and don't need to be included in the proof. We stress out the "trivial" identity commitment satisfies the properties needed to be used in our quasi-arguments.

Assumptions. We next discuss the assumptions we use. For the specific matrices used in the reduction, one can prove soundness of the QA-NIZK argument under falsifiable assumptions since the S-submatrices $\mathbf{M}_1, \mathbf{M}_2$ produce a non-trivial subspace. This means that we rely on the kernel assumption we use for instantiating the QA-NIZK. Noting that MDDH assumptions implies the corresponding kernel assumptions, we can instantiate the quasi argument using the DLIN assumption. Furthermore, the no-signaling property of the commitment keys (the only computational property we use) reduces to an MDDH which we chose on instantiation. Noting that DDH does not hold in symmetric groups we resort to the DLIN assumption which makes the commitments larger by 1 group element. Thus, soundness of the above delegation scheme reduces to the DLIN assumption.

Overcoming the dependence on space/width. The issue with the above construction is that setting $K = O(|\mathtt{st}|)$ yields a proof whose size is linear in the space of the computation. To achieve succinctness in the general case, we need to also perform some "compressing" of the state/width. Kalai et. al. overcome this by considering delegation of RAM computation [33] using collision-resistant hash function to compress the width. They use a notion similar to the knowledge transfer notion, namely that no PPT adversary can produce digests $\mathtt{h}, \mathtt{h}'$ and state $\mathtt{st}$ such that $\mathtt{h} = \mathsf{Hash}(\mathtt{st})$ but $\mathtt{h}' \neq \mathsf{Hash}(f(\mathtt{st}))$. Now, a quasi argument for the local constraints $\mathtt{h}_i = \mathsf{Hash}(f(\mathtt{st}_i))$ and $\mathtt{h}_{i+1} = \mathsf{Hash}(f(\mathtt{st}_i))$ is enough for delegation in the general case.

While previous works achieve this by essentially encoding the computation of generic hash functions in the computation, we use hash functions that are based on Pedersen commitments and have nice algebraic structure and properties. This allows to avoid the concrete cost of encoding arbitrary hash functions in the arithmetic circuit. To this end, we use techniques from [24] to derive a structure preserving construction. We present next the basic ideas of their (low depth) delegation construction.

Structure Preserving Delegation for Bounded-Depth Circuits. González and Ràfols [24] constructed a delegation scheme with proof-size $O(d\kappa)$ and verification requiring n plus $O(d)$ cryptographic operations, where n is the size of the input, d the depth of the circuit and κ a security parameter. Interestingly, the verification procedure of [24] can be described completely as a set of pairing product equations. As shown by Abe et al.[1], cryptographic primitives whose correctness can be stated as equations over bilinear groups are more suited for practically efficient arguments without resorting to generic reductions to a circuit or a 3CNF formula.

In the heart of the delegation scheme of [24] lie the two knowledge transfer arguments for linear and quadratic relations described before. To delegate the computation of an arithmetic circuit, the multiplication gates are partitioned in d levels. Each level groups the gates at the same distance from the inputs, without counting linear gates. In this way, the inputs of level $i + 1$ are linear

combinations of outputs of the i previous levels. A prover commits to the left, right, and output wires of each level as L_i, R_i, O_i. In the first d arguments f is a linear function and the argument handles the linear relations between the input wires (the openings of L_i, R_i) of level i and the output wires of all previous levels (the openings of $O_1, \ldots, O_{i-1}$). In the next d arguments f is the hadamard product so that the opening of O_i is the the hadamard product of the openings of L_i and R_i. The fact that the verifier can check the commitment to the first level using the public input and a simple inductive argument over the levels shows that the output must be correct.

More concretely, starting from a correct commitment O_0 (directly checked for consistency with input x from the verifier) we conclude that L_1, R_1 by the knowledge transfer guarantee of the linear argument. Since L_1, R_1 are correct w.r.t. x, O_1 is also correct w.r.t. x by the knowledge transfer guarantee of the quadratic arguement. We continue this way and we conclude that O_d is a correct commitment to the output of the computation. Now, we simply need to check that the claimed output y is a correct opening for that latter commitment.

As for soundness, the quadratic knowledge transfer arguement requires a specific (not uniform) distribution for the commitment keys where each row of the matrix of the commitment key is the result of evaluating Lagrange polynomials at a different random point. Thus, soundness relies on a width-size assumption, namely "$\mathcal{R}$-Rational Strong Diffie Hellman" assumption [24] which is proven secure in the Generic Group Model. We stress out that we modify the construction of [24] to overcome the need for a q-size assumption and rely only on a constnt-size one, albeit at the cost of having a quadratic crs and prover computation.

Succinct Publicly Verifiable Delegation for polynomial size circuits. We use the technique of [24] to overcome the width dependency in the above construction. The problem with this construction is that we need to rely on simple soundness of the underlying Kiltz and Wee QA-NIZK. However if we try to "shrink" the per-level information to eliminate the width dependence, the subspaces used become trivial and knowledge soundness seems to be needed.

We overcome this by relying on the knowledge transfer analysis of Kiltz and Wee used in [24]. To exploit this to construct delegation, we proceed as follows: we keep the same skeleton of the small-width circuit protocol, but instead of directly committing to the left, right and output wires, we commit to commitments of them. That is, for each level we compute three shrinking commitments -with size independent of the width- corresponding to left, right and output wires for that level, and we commit to these commitments (by including appropriate group elements in the crs). Furthermore, we use the knowledge transfer variants of the quasi arguments.

Now, our no-signaling extractor works as in the small-width case, but instead of the wires for some level, it outputs the commitments for the wires in this level. By the knowledge transfer guarantees, we establish that the extracted values for each level satisfy:

(1) *if* O_i is a commitment to $\boldsymbol{c}_i$ *then* L_{i+1} and R_{i+1} are commitments to $\boldsymbol{a}_{i+1}, \boldsymbol{b}_{i+1}$,

(2) *if* L_{i+1} and R_{i+1} are commitments to $\boldsymbol{a}_{i+1}$ and $\boldsymbol{b}_{i+1}$ respectively, *then* O_{i+1} is a commitment to $\boldsymbol{c}_{i+1}$

Extracting these values in a no-signaling way, as in the bounded space case, yields soundness for the delegation scheme. The analysis is almost the same and the only difference is that the knowledge transfer guarantee implies some hardness assumption (MDDH) on the distribution of matrices used as parameters, in this case, the width commitment keys. To satisfy this using constant size assumptions, we use a simple variation of Pedersen commitments where the commitment keys satisfy the DLIN assumption.

Remark 2 (Uniform vs Non-Uniform Computation). Our construction can be used for any non-uniform computation, namely polynomial size arithmetic circuits, while previous works such as [44, 34] focus on delegating uniform computations: Turing or RAM machines. While this is a stronger result, we achieve it using a long (quadratic in the size/time of computation) crs while the work of [34] achieves a short (i.e. sublinear) crs. One motivation for working directly with poly-size circuits is for practical efficiency: we utilize the rich SNARK toolbox without the need to encode expensive cryptographic operations as arithmetic circuits, namely, we focus on structure preserving constructions. While we have an inefficient (quadratic) prover, in all other aspects we achieve optimal efficiency comparable with SNARGs from non-falsifiable assumptions. We believe that this is a promising direction and an interesting open problem is to improve the prover to quasi-linear using these techniques. This would yield a delegation scheme for poly-size circuits that directly competes with the aforementioned non-falsifiable based constructions in all aspects, effectively making the use of non-falsifiable assumptions unjustifiable in the context of deterministic computation. We also leave as future work exploring to what extend our techniques can be applied for delegating uniform computations and if this would give some improvement over existing constructions.

Remark 3 (On bootstrapping and proof composition). To improve efficiency (crs size), [34] use the bootstrapping technique which involves proof composition. Our techniques seem to be incompatible with the bootstrapping technique. This is because the crs of our construction depends on the circuit and we cannot directly reuse a crs for different computations. We leave as future work to examine if we can modify our techniques to be able to apply the bootstrapping technique. We also stress out that this might prove to be an interesting direction for improvements in practical efficiency as well due to some recent results in proof-composition techniques [10, 9].

2.4 NIZK, SNARKs and Compact NIZK

We can use standard techniques to turn our delegation scheme into a NIZK argument. Essentially, the prover needs to prove knowledge of (additional) secret input wires w and proof that $C(x, w) = y$ for some secret input w. Given the "structure preserving" properties of our delegation scheme, we can directly

apply the Groth Sahai proof system [27][14] on the set of verification equations. In general, all we need to achieve knowledge soundness is an extractable (and hiding) commitment for extracting the witness w. Depending on the properties of the extractable commitment scheme we get different NIZK flavors.

If the commitments to the inputs are succinct, the construction yields a SNARK for NP. Such commitments are widely employed in SNARKs, but their security relies on non-standard assumptions: either knowledge type assumptions such as q-Knowledge of Exponents assumption [19] or the generic group model [25]. If we take for example the zk-SNARK from [15], the size of q is the number of field elements extracted from a valid proof. Indeed, the proof of soundness requires the extraction of all the circuit wires, which are later used to break some falsifiable q-assumption. Consequently, the knowledge assumption is of size $q = O(|C|)$. By reducing the number of extracted values from $O(|C|)$ to $|w|$, we reduce the size of the underlying knowledge assumption to $q = |w| < |C|$.

If we use the "bit-string" argument of [22] to show knowledge of $\boldsymbol{b} \in \{0,1\}^n$, we get extractable commitments of size $n + O(1)$ group elements based on a constant-size falsifiable assumption. Combining this extractable commitment with our delegation scheme yields a NIZK argument for circuit satisfiability with proof size $n + O(1)$ groups elements, or equivalently of size $O(n\kappa)$.

Finally, we can then use the techniques of Katsumata et al. [38, 37] to construct a compact NIZK. The construction of Katsumata et al. is based on a non-compact NIZK argument for NC^1 plus a symmetric key encryption scheme $(\mathsf{K}, \mathsf{E}, \mathsf{D})$ where the size of $\mathsf{E}(K, m)$ is $|m| + \mathsf{poly}(\kappa)$. Instead of committing to the input $\boldsymbol{x}$ of a circuit C, we need to compute $K \leftarrow \mathsf{K}(1^\kappa)$ to obtain $ct \leftarrow \mathsf{E}(K, \boldsymbol{x})$ and give a NIZK argument of knowledge of some $K \in \{0,1\}^{\mathsf{poly}(\kappa)}$ such that $C(\mathsf{D}(K, ct)) = 1$. We note that we can straightforward use this idea to construct compact NIZK for any circuit by simply plugging our NIZK argument based on the commitments of [22]. The final proof is of size $|ct| + |K|\mathsf{poly}(\kappa) + |\pi| = n + \mathsf{poly}(\kappa)$ and is sound for any polynomial size circuit.

References

1. Abe, M., Fuchsbauer, G., Groth, J., Haralambiev, K., Ohkubo, M.: Structure-preserving signatures and commitments to group elements **29**(2), 363–421. https://doi.org/10.1007/s00145-014-9196-7
2. Abe, M., Fuchsbauer, G., Groth, J., Haralambiev, K., Ohkubo, M.: Structure-preserving signatures and commitments to group elements. In: Rabin, T. (ed.) CRYPTO 2010. LNCS, vol. 6223, pp. 209–236. Springer, Heidelberg. https://doi.org/10.1007/978-3-642-14623-7_12
3. Ananth, P., Chen, Y.C., Chung, K.M., Lin, H., Lin, W.K.: Delegating RAM computations with adaptive soundness and privacy. In: Hirt, M., Smith, A.D. (eds.) TCC 2016-B, Part II. LNCS, vol. 9986, pp. 3–30. Springer, Heidelberg. https://doi.org/10.1007/978-3-662-53644-5_1

[14] This can be also achieved in a more efficient way (concretely) by directly using hiding commitments for the delegation scheme.

4. Backes, M., Döttling, N., Hanzlik, L., Kluczniak, K., Schneider, J.: Ring signatures: Logarithmic-size, no setup - from standard assumptions. In: Ishai, Y., Rijmen, V. (eds.) EUROCRYPT 2019, Part III. LNCS, vol. 11478, pp. 281–311. Springer, Heidelberg. https://doi.org/10.1007/978-3-030-17659-4_10
5. Badrinarayanan, S., Kalai, Y.T., Khurana, D., Sahai, A., Wichs, D.: Succinct delegation for low-space non-deterministic computation. In: Diakonikolas, I., Kempe, D., Henzinger, M. (eds.) 50th ACM STOC. pp. 709–721. ACM Press. https://doi.org/10.1145/3188745.3188924
6. Bitansky, N., Garg, S., Lin, H., Pass, R., Telang, S.: Succinct randomized encodings and their applications. Cryptology ePrint Archive, Report 2015/356, https://eprint.iacr.org/2015/356
7. Boneh, D., Boyen, X., Shacham, H.: Short group signatures. In: Franklin, M. (ed.) CRYPTO 2004. LNCS, vol. 3152, pp. 41–55. Springer, Heidelberg. https://doi.org/10.1007/978-3-540-28628-8_3
8. Brakerski, Z., Holmgren, J., Kalai, Y.T.: Non-interactive delegation and batch NP verification from standard computational assumptions. In: Hatami, H., McKenzie, P., King, V. (eds.) 49th ACM STOC. pp. 474–482. ACM Press. https://doi.org/10.1145/3055399.3055497
9. Bünz, B., Chiesa, A., Lin, W., Mishra, P., Spooner, N.: Proof-carrying data without succinct arguments. In: Malkin, T., Peikert, C. (eds.) CRYPTO 2021, Part I. LNCS, vol. 12825, pp. 681–710. Springer, Heidelberg. https://doi.org/10.1007/978-3-030-84242-0_24
10. Bünz, B., Chiesa, A., Mishra, P., Spooner, N.: Recursive proof composition from accumulation schemes. In: Pass, R., Pietrzak, K. (eds.) TCC 2020, Part II. LNCS, vol. 12551, pp. 1–18. Springer, Heidelberg. https://doi.org/10.1007/978-3-030-64378-2_1
11. Canetti, R., Chen, Y., Holmgren, J., Lombardi, A., Rothblum, G.N., Rothblum, R.D., Wichs, D.: Fiat-Shamir: from practice to theory. In: Charikar, M., Cohen, E. (eds.) 51st ACM STOC. pp. 1082–1090. ACM Press. https://doi.org/10.1145/3313276.3316380
12. Canetti, R., Holmgren, J., Jain, A., Vaikuntanathan, V.: Succinct garbling and indistinguishability obfuscation for RAM programs. In: Servedio, R.A., Rubinfeld, R. (eds.) 47th ACM STOC. pp. 429–437. ACM Press. https://doi.org/10.1145/2746539.2746621
13. Chen, Y.C., Chow, S.S.M., Chung, K.M., Lai, R.W.F., Lin, W.K., Zhou, H.S.: Cryptography for parallel RAM from indistinguishability obfuscation. In: Sudan, M. (ed.) ITCS 2016. pp. 179–190. ACM. https://doi.org/10.1145/2840728.2840769
14. Cramer, R., Shoup, V.: Universal hash proofs and a paradigm for adaptive chosen ciphertext secure public-key encryption. In: Knudsen, L.R. (ed.) EUROCRYPT 2002. LNCS, vol. 2332, pp. 45–64. Springer, Heidelberg. https://doi.org/10.1007/3-540-46035-7_4
15. Danezis, G., Fournet, C., Groth, J., Kohlweiss, M.: Square span programs with applications to succinct NIZK arguments. In: Sarkar, P., Iwata, T. (eds.) ASIACRYPT 2014, Part I. LNCS, vol. 8873, pp. 532–550. Springer, Heidelberg. https://doi.org/10.1007/978-3-662-45611-8_28
16. Daza, V., González, A., Pindado, Z., Ràfols, C., Silva, J.: Shorter quadratic QA-NIZK proofs. In: Lin, D., Sako, K. (eds.) PKC 2019, Part I. LNCS, vol. 11442, pp. 314–343. Springer, Heidelberg. https://doi.org/10.1007/978-3-030-17253-4_11
17. Escala, A., Herold, G., Kiltz, E., Ràfols, C., Villar, J.: An algebraic framework for Diffie-Hellman assumptions. In: Canetti, R., Garay, J.A. (eds.)

CRYPTO 2013, Part II. LNCS, vol. 8043, pp. 129–147. Springer, Heidelberg. https://doi.org/10.1007/978-3-642-40084-1_8
18. Fauzi, P., Lipmaa, H., Pindado, Z., Siim, J.: Somewhere statistically binding commitment schemes with applications. Cryptology ePrint Archive, Report 2020/652, https://eprint.iacr.org/2020/652
19. Gennaro, R., Gentry, C., Parno, B., Raykova, M.: Quadratic span programs and succinct NIZKs without PCPs. In: Johansson, T., Nguyen, P.Q. (eds.) EUROCRYPT 2013. LNCS, vol. 7881, pp. 626–645. Springer, Heidelberg. https://doi.org/10.1007/978-3-642-38348-9_37
20. Gentry, C., Wichs, D.: Separating succinct non-interactive arguments from all falsifiable assumptions. In: Fortnow, L., Vadhan, S.P. (eds.) 43rd ACM STOC. pp. 99–108. ACM Press. https://doi.org/10.1145/1993636.1993651
21. Goldwasser, S., Kalai, Y.T., Rothblum, G.N.: Delegating computation: interactive proofs for muggles. In: Ladner, R.E., Dwork, C. (eds.) 40th ACM STOC. pp. 113–122. ACM Press. https://doi.org/10.1145/1374376.1374396
22. González, A., Hevia, A., Ràfols, C.: QA-NIZK arguments in asymmetric groups: New tools and new constructions. In: Iwata, T., Cheon, J.H. (eds.) ASIACRYPT 2015, Part I. LNCS, vol. 9452, pp. 605–629. Springer, Heidelberg. https://doi.org/10.1007/978-3-662-48797-6_25
23. González, A., Ràfols, C.: New techniques for non-interactive shuffle and range arguments. In: Manulis, M., Sadeghi, A.R., Schneider, S. (eds.) ACNS 16. LNCS, vol. 9696, pp. 427–444. Springer, Heidelberg. https://doi.org/10.1007/978-3-319-39555-5_23
24. González, A., Ràfols, C.: Shorter pairing-based arguments under standard assumptions. In: Galbraith, S.D., Moriai, S. (eds.) ASIACRYPT 2019, Part III. LNCS, vol. 11923, pp. 728–757. Springer, Heidelberg. https://doi.org/10.1007/978-3-030-34618-8_25
25. Groth, J.: On the size of pairing-based non-interactive arguments. In: Fischlin, M., Coron, J.S. (eds.) EUROCRYPT 2016, Part II. LNCS, vol. 9666, pp. 305–326. Springer, Heidelberg. https://doi.org/10.1007/978-3-662-49896-5_11
26. Groth, J., Ostrovsky, R., Sahai, A.: Perfect non-interactive zero knowledge for NP. In: Vaudenay, S. (ed.) EUROCRYPT 2006. LNCS, vol. 4004, pp. 339–358. Springer, Heidelberg. https://doi.org/10.1007/11761679_21
27. Groth, J., Sahai, A.: Efficient non-interactive proof systems for bilinear groups. In: Smart, N.P. (ed.) EUROCRYPT 2008. LNCS, vol. 4965, pp. 415–432. Springer, Heidelberg. https://doi.org/10.1007/978-3-540-78967-3_24
28. Hubacek, P., Wichs, D.: On the communication complexity of secure function evaluation with long output. In: Roughgarden, T. (ed.) ITCS 2015. pp. 163–172. ACM. https://doi.org/10.1145/2688073.2688105
29. Jawale, R., Kalai, Y.T., Khurana, D., Zhang, R.: SNARGs for bounded depth computations and PPAD hardness from sub-exponential LWE. Cryptology ePrint Archive, Report 2020/980, https://eprint.iacr.org/2020/980
30. Jutla, C.S., Roy, A.: Shorter quasi-adaptive NIZK proofs for linear subspaces. In: Sako, K., Sarkar, P. (eds.) ASIACRYPT 2013, Part I. LNCS, vol. 8269, pp. 1–20. Springer, Heidelberg. https://doi.org/10.1007/978-3-642-42033-7_1
31. Jutla, C.S., Roy, A.: Switching lemma for bilinear tests and constant-size NIZK proofs for linear subspaces. In: Garay, J.A., Gennaro, R. (eds.) CRYPTO 2014, Part II. LNCS, vol. 8617, pp. 295–312. Springer, Heidelberg. https://doi.org/10.1007/978-3-662-44381-1_17

32. Kalai, Y., Paneth, O., Yang, L.: On publicly verifiable delegation from standard assumptions. Cryptology ePrint Archive, Report 2018/776, https://eprint.iacr.org/2018/776
33. Kalai, Y.T., Paneth, O.: Delegating RAM computations. In: Hirt, M., Smith, A.D. (eds.) TCC 2016-B, Part II. LNCS, vol. 9986, pp. 91–118. Springer, Heidelberg. https://doi.org/10.1007/978-3-662-53644-5_4
34. Kalai, Y.T., Paneth, O., Yang, L.: How to delegate computations publicly. In: Charikar, M., Cohen, E. (eds.) 51st ACM STOC. pp. 1115–1124. ACM Press. https://doi.org/10.1145/3313276.3316411
35. Kalai, Y.T., Raz, R., Rothblum, R.D.: Delegation for bounded space. In: Boneh, D., Roughgarden, T., Feigenbaum, J. (eds.) 45th ACM STOC. pp. 565–574. ACM Press. https://doi.org/10.1145/2488608.2488679
36. Kalai, Y.T., Raz, R., Rothblum, R.D.: How to delegate computations: the power of no-signaling proofs. In: Shmoys, D.B. (ed.) 46th ACM STOC. pp. 485–494. ACM Press. https://doi.org/10.1145/2591796.2591809
37. Katsumata, S., Nishimaki, R., Yamada, S., Yamakawa, T.: Compact NIZKs from standard assumptions on bilinear maps. In: Canteaut, A., Ishai, Y. (eds.) EUROCRYPT 2020, Part III. LNCS, vol. 12107, pp. 379–409. Springer, Heidelberg. https://doi.org/10.1007/978-3-030-45727-3_13
38. Katsumata, S., Nishimaki, R., Yamada, S., Yamakawa, T.: Exploring constructions of compact NIZKs from various assumptions. In: Boldyreva, A., Micciancio, D. (eds.) CRYPTO 2019, Part III. LNCS, vol. 11694, pp. 639–669. Springer, Heidelberg. https://doi.org/10.1007/978-3-030-26954-8_21
39. Kiltz, E., Wee, H.: Quasi-adaptive NIZK for linear subspaces revisited. In: Oswald, E., Fischlin, M. (eds.) EUROCRYPT 2015, Part II. LNCS, vol. 9057, pp. 101–128. Springer, Heidelberg. https://doi.org/10.1007/978-3-662-46803-6_4
40. Koppula, V., Lewko, A.B., Waters, B.: Indistinguishability obfuscation for turing machines with unbounded memory. In: Servedio, R.A., Rubinfeld, R. (eds.) 47th ACM STOC. pp. 419–428. ACM Press. https://doi.org/10.1145/2746539.2746614
41. Libert, B., Peters, T., Joye, M., Yung, M.: Linearly homomorphic structure-preserving signatures and their applications. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013, Part II. LNCS, vol. 8043, pp. 289–307. Springer, Heidelberg. https://doi.org/10.1007/978-3-642-40084-1_17
42. Morillo, P., Ràfols, C., Villar, J.L.: The kernel matrix Diffie-Hellman assumption. In: Cheon, J.H., Takagi, T. (eds.) ASIACRYPT 2016, Part I. LNCS, vol. 10031, pp. 729–758. Springer, Heidelberg. https://doi.org/10.1007/978-3-662-53887-6_27
43. Okamoto, T., Pietrzak, K., Waters, B., Wichs, D.: New realizations of somewhere statistically binding hashing and positional accumulators. In: Iwata, T., Cheon, J.H. (eds.) ASIACRYPT 2015, Part I. LNCS, vol. 9452, pp. 121–145. Springer, Heidelberg. https://doi.org/10.1007/978-3-662-48797-6_6
44. Paneth, O., Rothblum, G.N.: On zero-testable homomorphic encryption and publicly verifiable non-interactive arguments. In: Kalai, Y., Reyzin, L. (eds.) TCC 2017, Part II. LNCS, vol. 10678, pp. 283–315. Springer, Heidelberg. https://doi.org/10.1007/978-3-319-70503-3_9
45. Reingold, O., Rothblum, G.N., Rothblum, R.D.: Constant-round interactive proofs for delegating computation. In: Wichs, D., Mansour, Y. (eds.) 48th ACM STOC. pp. 49–62. ACM Press. https://doi.org/10.1145/2897518.2897652