Prover-Efficient Commit-and-Prove Zero-Knowledge SNARKs

Helger Lipmaa(✉)

Institute of Computer Science, University of Tartu, Tartu, Estonia
helger.lipmaa@gmail.com

Abstract. Zk-SNARKs (succinct non-interactive zero-knowledge arguments of knowledge) are needed in many applications. Unfortunately, all previous zk-SNARKs for interesting languages are either inefficient for the prover, or are non-adaptive and based on a commitment scheme that depends both on the prover's input and on the language, i.e., they are not commit-and-prove (CaP) SNARKs. We propose a proof-friendly extractable commitment scheme, and use it to construct prover-efficient adaptive CaP succinct zk-SNARKs for different languages, that can all reuse committed data. In new zk-SNARKs, the prover computation is dominated by a linear number of cryptographic operations. We use batch-verification to decrease the verifier's computation; importantly, batch-verification can be used also in QAP-based zk-SNARKs.

Keywords: Batch verification · Commit-and-prove · CRS · NIZK · Numerical NP-complete languages · Range proof · Subset-Sum · zk-SNARK

1 Introduction

Recently, there has been a significant surge of activity in studying succinct non-interactive zero knowledge (NIZK) arguments of knowledge (also known as zk-SNARKs) [3–6,12,13,17,19,23,24,28]. The prover of a zk-SNARK outputs a short (ideally, a small number of group elements) argument π that is used to convince many different verifiers in the truth of the same claim without leaking any side information. The verifiers can verify independently the correctness of π, without communicating with the prover. The argument must be efficiently verifiable. Constructing the argument can be less efficient, since it is only done once. Still, prover-efficiency is important, e.g., in a situation where a single server has to create many arguments to different clients or other servers.

Many known zk-SNARKs are non-adaptive, meaning that the common reference string, CRS, can depend on the concrete instance of the language (e.g., the circuit in the case of Circuit-SAT). In an adaptive zk-SNARK, the CRS is independent on the instance and thus can be reused many times. This distinction is important, since generation and distribution of the CRS must be

D. Pointcheval et al. (Eds.): AFRICACRYPT 2016, LNCS 9646, pp. 185–206, 2016.
DOI: 10.1007/978-3-319-31517-1_10

done securely. The most efficient known *non-adaptive* zk-SNARKs for NP-complete languages from [17] are based on either Quadratic Arithmetic Programs (QAP, for arithmetic Circuit-SAT) or Quadratic Span Programs (QSP, for Boolean Circuit-SAT). There, the prover computation is dominated by $\Theta(n)$ cryptographic operations (see the full version [26] for a clarification on cryptographic/non-cryptographic operations), where n is the number of the gates. QAP, QSP [17,24] and other related approaches like SSP [13] have the same asymptotic complexity.

QSP-based Circuit-SAT SNARK can be made adaptive by using universal circuits [33]. Then, the CRS depends on the construction of universal circuit and not on the concrete input circuit itself. However, since the size of a universal circuit is $\Theta(n \log n)$, the prover computation in resulting adaptive zk-SNARKs is $\Theta(n \log^2 n)$ non-cryptographic operations and $\Theta(n \log n)$ cryptographic operations. (In the case of QAP-based arithmetic Circuit-SAT SNARK, one has to use universal arithmetic circuits [30] that have an even larger size $\Theta(r^4 n)$, where r is the degree of the polynomial computed by the arithmetic circuit. Thus, we will mostly give a comparison to the QSP-based approach.)

Since Valiant's universal circuits incur a large constant $c = 19$ in the $\Theta(\cdot)$ expression, a common approach [21,31] is to use universal circuits with the overhead of $\Theta(\log^2 n)$ but with a smaller constant $c = 1/2$ in $\Theta(\cdot)$. The prover computation in the resulting adaptive zk-SNARKs is $\Theta(n \log^3 n)$ non-cryptographic operations and $\Theta(n \log^2 n)$ cryptographic operations.[1]

Another important drawback of the QSP/QAP-based SNARKs is that they use a circuit-dependent commitment scheme. To use the same input data in multiple sub-SNARKs, one needs to construct a single large circuit that implements all sub-SNARKs, making the SNARK and the resulting *new* commitment scheme more complicated. In particular, these SNARKs are not commit-and-prove (CaP [9,20]) SNARKs. We recall that in CaP SNARKs, a commitment scheme C is fixed first, and the statement consists of commitments of the witness using C; see Sect. 2. Hence, a CaP commitment scheme is *instance-independent*. In addition, one would like the commitment scheme to be *language-independent*, enabling one to first commit to the data and only then to decide in what applications (e.g., verifiable computation of a later fixed function) to use it.

See Table 1 for a brief comparison of the efficiency of proposed adaptive zk-SNARKs for NP-complete languages. Subset-Sum is here brought as an example of a wider family of languages; it can be replaced everywhere say with Partition or Knapsack, see the full version [26]. Here, $N = r_3^{-1}(n) = o(n2^{2\sqrt{2\log_2 n}})$, where $r_3(n)$ is the density of the largest progression-free set in $\{1, \ldots, n\}$. According to the current knowledge, $r_3^{-1}(n)$ is comparable to (or only slightly smaller than) n^2 for $n < 2^{12}$; this makes all known CaP SNARKs [15,19,23] arguably impractical unless n is really small. In all cases, the verifier's computation is dominated by either $\Theta(n)$ cryptographic or $\Theta(n \log n)$

[1] Recently, [12] proposed an independent methodology to improve the prover's computational complexity in QAP-based arguments. However, [12] does not spell out their achieved prover's computational complexity.

Table 1. Prover-efficiency of known *adaptive* zk-SNARKs for NP-complete languages. Here, n is the number of the gates (in the case of Circuit-SAT) and the number of the integers (in the case of Subset-Sum). Green background denotes the best known asymptotic complexity of the *concrete* NP-complete language w.r.t. to the concrete parameter. The solutions marked with * use proof bootstrapping from [12]

Paper	Language	Prover computation		|CRS|
		non-crypt. op.	crypt. op.	
Not CaP-s				
QAP, QSP ([14,19,27])	Circuit-SAT	$\Theta(n \log^2 n)$	$\Theta(n \log n)$	$\Theta(n)$
CaP-s				
Gro10 ([21])	Circuit-SAT	$\Theta(n^2)$	$\Theta(n^2)$	$\Theta(n^2)$
Lip12 ([26])	Circuit-SAT	$\Theta(n^2)$	$\Theta(N)$	$\Theta(N)$
Lip14 + Lip12 ([26,28])*	Circuit-SAT	$\Theta(N \log^2 n)$	$\Theta(N \log n)$	$\Theta(N \log n)$
Lip14 + current paper ([28])*	Circuit-SAT	$\Theta(n \log^2 n)$	$\Theta(n \log n)$	$\Theta(n \log n)$
FLZ13 ([16])	Subset-Sum	$\Theta(N \log n)$	$\Theta(N)$	$\Theta(N)$
Current paper	Subset-Sum	$\Theta(n \log n)$	$\Theta(n)$	$\Theta(n)$

non-cryptographic operations (with the verifier's online computation usually being $\Theta(1)$), and the communication consists of a small constant number of group elements.[2] Given all above, it is natural to ask the following question:

The Main Question of This Paper: *Is it possible to construct* adaptive CaP *zk-SNARKs for NP-complete languages where the prover computation is dominated by a linear number of cryptographic operations?*

We answer the "main question" positively by improving on Groth's modular approach [19]. Using the modular approach allows us to modularize the security analysis, first proving the security of underlying building blocks (the product and the shift SNARKs), and then composing them to construct master SNARKs for even NP-complete languages. The security of master SNARKs follows easily from the security of the basic SNARKs. We also use batch verification to speed up verification of almost all known SNARKs.

All new SNARKs use the same commitment scheme, the interpolating commitment scheme. Hence, one can reuse their input data to construct CaP zk-SNARKs for different unrelated languages, chosen only after the commitment was done. Thus, one can first commit to some data, and only later decide in which application and to what end to use it. Importantly, by using CaP zk-SNARKs, one can guarantee that all such applications use exactly the same data.

[2] We emphasize that Circuit-SAT is *not* our focus; the lines corresponding to Circuit-SAT are provided only for the sake of comparison. One can use proof boot-strapping [12] to decrease the length of the resulting Circuit-SAT argument from $\Theta(\log n)$, as stated in [25], to $\Theta(1)$; we omit further discussion.

The resulting SNARKs are not only commit-and-prove, but also very efficient, and often more efficient than any previously known SNARKs. The new CaP SNARKs have prover-computation dominated by $\Theta(n)$ cryptographic operations, with the constant in $\Theta(\cdot)$ being reasonably small. Importantly, we propose the most efficient known succinct range SNARK. Since the resulting zk-SNARKs are sufficiently different from QAP-based zk-SNARKs, we hope that our methodology by itself is of independent interest. Up to the current paper, Groth's modular approach has resulted in significantly less efficient zk-SNARKs than the QSP/QAP-based approach.

In Sect. 3, we construct a new natural extractable trapdoor commitment scheme (the interpolating commitment scheme). Here, commitment to $\boldsymbol{a} \in \mathbb{Z}_p^n$, where n is a power of 2, is a short garbled and randomized version $g_1^{L_{\boldsymbol{a}}(\chi)}(g_1^{\chi^n-1})^r$ of the Lagrange interpolating polynomial $L_{\boldsymbol{a}}(X)$ of $\boldsymbol{a}$, for a random secret key χ, together with a knowledge component. This commitment scheme is arguably a very natural one, and in particular its design is not influenced by the desire to tailor it to one concrete application. Nevertheless, as we will see, using it improves the efficiency of many constructions while allowing to reuse many existing results.

The new CaP zk-SNARKs are based on the interpolating commitment scheme and two CaP witness-indistinguishable SNARKs: a product SNARK (given commitments to vectors $\boldsymbol{a}$, $\boldsymbol{b}$, $\boldsymbol{c}$, it holds that $c_i = a_i b_i$; see [15,19,23]), and a shift SNARK (given commitments to $\boldsymbol{a}$, $\boldsymbol{b}$, it holds that $\boldsymbol{a}$ is a coordinate-wise shift of $\boldsymbol{b}$; see [15]). One can construct an adaptive Circuit-SAT CaP zk-SNARK from $\Theta(\log n)$ product and shift SNARKs [19,25], or adaptive CaP zk-SNARKs for NP-complete languages like Subset-Sum (and a similar CaP range SNARK) by using a constant number of product and shift SNARKs [15].

In Sect. 4, we propose a CaP product SNARK, that is an argument of knowledge under a computational and a knowledge (needed solely to achieve extractability of the commitment scheme) assumption. Its prover computation is dominated by $\Theta(n \log n)$ non-cryptographic and $\Theta(n)$ cryptographic operations. This can be compared to $r_3^{-1}(n)$ non-cryptographic operations in [15]. The speed-up is mainly due to the use of the interpolating commitment scheme.

In Sect. 5, we propose a variant of the CaP shift SNARK of [15], secure when combined with the interpolating commitment scheme. We prove that this SNARK is an adaptive argument of knowledge under a computational and a knowledge assumption. It only requires the prover to perform $\Theta(n)$ cryptographic and non-cryptographic operations.

Product and shift SNARKs are already very powerful by itself. E.g., a prover can commit to her input vector $\boldsymbol{a}$. Then, after agreeing with the verifier on a concrete application, she can commit to a different yet related input vector (that say consists of certain permuted subset of $\boldsymbol{a}$'s coefficients), and then use the basic SNARKs to prove that this was done correctly. Here, she may use the permutation SNARK [25] that consists of $O(\log n)$ product and shift SNARKs. Finally, she can use another, application-specific, SNARK (e.g., a range SNARK) to prove that the new committed input vector has been correctly formed.

In Sect. 6, we describe a modular adaptive CaP zk-SNARK, motivated by [15], for the NP-complete language, Subset-Sum. (Subset-Sum was chosen by us mainly due to the simplicity of the SNARK; the rest of the paper considers more applications.) This SNARK consists of three commitments, one application of the shift SNARK, and three applications of the product SNARK. It is a zk-SNARK given that the commitment scheme, the shift SNARK, and the product SNARK are secure. Its prover computation is strongly dominated by $\Theta(n)$ cryptographic operations, where n is the instance size, the number of integers. More precisely, the prover has to perform only nine ($\approx n$)-wide multi-exponentiations, which makes the SNARK efficient not only asymptotically (to compare, the size of Valiant's arithmetic circuit has constant 19, and this constant has to be multiplied by the overhead of non-adaptive QSP/QAP/SSP-based solutions). Thus, we answer positively to the stated main question of the current paper. Moreover, the prover computation is highly parallelizable, while the *online* verifier computation is dominated by 17 pairings (this number will be decreased later).

In Sect. 7, we propose a new CaP range zk-SNARK that the committed value belongs to a range $[L\,..\,H]$. This SNARK looks very similar to the Subset-Sum SNARK, but with the integer set $\boldsymbol{S}$ of the Subset-Sum language depending solely on the range length. Since here the prover has a committed input, the simulation of the range SNARK is slightly more complicated than of the Subset-Sum SNARK. Its prover-computation is similarly dominated by $\Theta(n)$ cryptographic operations, where this time $n := \lceil \log_2(H - L) \rceil$. Differently from the Subset-Sum SNARK, the verifier computation is dominated only by $\Theta(1)$ cryptographic operations, more precisely, by 19 pairings (also this number will be decreased later). Importantly, this SNARK is computationally more efficient than any of the existing *succinct* range SNARKs either in the standard model (i.e., random oracle-less) or in the random oracle model. E.g., the prover computation in [22] is $\Theta(n^2)$ under the Extended Riemann Hypothesis, and the prover computation in [15] is $\Theta(r^{-3}(n) \log r^{-3}(n))$. It is also significantly simpler than the range SNARKs of [11,15], mostly since we do not have to consider different trade-offs between computation and communication.

In the full version [26], we outline how to use the new basic SNARKs to construct efficient zk-SNARKs for several other NP-complete languages like Boolean and arithmetic Circuit-SAT, Two-Processor Scheduling, Subset-Product, Partition, and Knapsack [16]. Table 1 includes the complexity of Subset-Sum and Circuit-SAT, the complexity of most other SNARKs is similar to that of Subset-Sum zk-SNARK. It is an interesting open problem why some NP-complete languages like Subset-Sum have more efficient zk-SNARKs in the modular approach (equivalently, why their verification can be performed more efficiently in the parallel machine model that consists of Hadamard product and shift) than languages like Circuit-SAT. We note that [14] used recently some of the ideas from the current paper to construct an efficient shuffle argument. However, they did not use product or shift arguments.

In the full version [26], we show that by using batch-verification [2], one can decrease the verifier's computation of all presented SNARKs. In particular, one can

decrease the verifier's computation in the new Range SNARK from 19 pairings to 8 pairings, one 4-way multi-exponentiation in $\mathbb{G}_1$, two 3-way multi-exponentiations in $\mathbb{G}_1$, one 2-way multi-exponentiation in $\mathbb{G}_1$, three exponentiations in $\mathbb{G}_1$, and one 3-way multi-exponentiation in $\mathbb{G}_2$. Since one exponentiation is much cheaper than one pairing [8] and one m-way multi-exponentiation is much cheaper than m exponentiations [29,32], this results in a significant win for the verifier. A similar technique can be used to also speed up other SNARKs; a good example here is the Circuit-SAT argument from [25] that uses $\Theta(\log n)$ product and shift arguments. To compare, in Pinocchio [28] and Geppetto [12], the verifier has to execute 11 pairings; however, batch-verification can also be used to decrease this to 8 pairings and a small number of (multi-)exponentiations.

Finally, all resulting SNARKs work on data that has been committed to by using the interpolating commitment scheme. This means that one can repeatedly reuse committed data to compose different zk-SNARKs (e.g., to show that we know a satisfying input to a circuit, where the first coefficient belongs to a certain range). This is not possible with the known QSP/QAP-based zk-SNARKs where one would have to construct a single circuit of possibly considerable size, say n'. Moreover, in the QSP/QAP-based SNARKs, one has to commit to the vector, the length of which is equal to the total length of the input and witness (e.g., n' is the number of wires in the case of Circuit-SAT). By using a modular solution, one can instead execute several zk-SNARKs with smaller values of the input and witness size; this can make the SNARK more prover-efficient since the number of non-cryptographic operations is superlinear. This emphasizes another benefit of the modular approach: one can choose the value n, the length of the vectors, accordingly to the desired tradeoff, so that larger n results in faster verifier computation, while smaller n results in faster prover computation. We are not aware of such a tradeoff in the case of the QSP/QAP-based approach.

We provide some additional discussion (about the relation between n and then input length, and about possible QSP/QAP-based solutions) in the full version [26]. Due to the lack of space, many proofs and details are only given in the full version [26]. We note that an early version of this paper, [26], was published in May 2014 and thus predates [12]. The published version differs from this early version mainly by exposition, and the use of proof bootstrapping (from [12]) and batching.

2 Preliminaries

By default, all vectors have dimension n. Let $\boldsymbol{a} \circ \boldsymbol{b}$ denote the Hadamard (i.e., element-wise) product of two vectors, with $(\boldsymbol{a} \circ \boldsymbol{b})_i = a_i b_i$. We say that $\boldsymbol{a}$ is a *shift-right-by-z* of $\boldsymbol{b}$, $\boldsymbol{a} = \boldsymbol{b} \gg z$, iff $(a_n, \ldots, a_1) = (0, \ldots, 0, b_n, \ldots, b_{1+z})$. For a tuple of polynomials $\mathcal{F} \subseteq \mathbb{Z}_p[X, Y_1, \ldots, Y_{m-1}]$, define $Y_m\mathcal{F} = (Y_m \cdot f(X, Y_1, \ldots, Y_{m-1}))_{f\in\mathcal{F}} \subseteq \mathbb{Z}_p[X, Y_1, \ldots, Y_m]$. For a tuple of polynomials $\mathcal{F}$ that have the same domain, denote $h^{\mathcal{F}(\boldsymbol{a})} := (h^{f(\boldsymbol{a})})_{f\in\mathcal{F}}$. For a group $\mathbb{G}$, let $\mathbb{G}^*$ be the set of its invertible elements. Since the direct product $\mathbb{G}_1 \times \ldots \times \mathbb{G}_m$ of groups is also a group, we use notation like $(g_1, g_2)^c = (g_1^c, g_2^c) \in \mathbb{G}_1 \times \mathbb{G}_2$ without

prior definition. Let κ be the security parameter. We denote $f(\kappa) \approx_\kappa g(\kappa)$ if $|f(\kappa) - g(\kappa)|$ is negligible in κ.

On input 1^κ, a *bilinear map generator* BP returns $\mathsf{gk} = (p, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, \hat{e})$, where $\mathbb{G}_1, \mathbb{G}_2$ and $\mathbb{G}_T$ are three multiplicative cyclic groups of prime order p (with $\log p = \Omega(\kappa)$), and $\hat{e}$ is an efficient bilinear map $\hat{e}\colon \mathbb{G}_1 \times \mathbb{G}_2 \to \mathbb{G}_T$ that satisfies in particular the following two properties, where g_1 (resp., g_2) is an arbitrary generator of $\mathbb{G}_1$ (resp., $\mathbb{G}_2$): (i) $\hat{e}(g_1, g_2) \neq 1$, and (ii) $\hat{e}(g_1^a, g_2^b) = \hat{e}(g_1, g_2)^{ab}$. Thus, if $\hat{e}(g_1^a, g_2^b) = \hat{e}(g_1^c, g_2^d)$ then $ab \equiv cd \pmod p$. We also give BP another input, n (intuitively, the input length), and allow p to depend on n. We assume that all algorithms that handle group elements verify by default that their inputs belong to corresponding groups and reject if they do not. In the case of many practically relevant pairings, arithmetic in (say) $\mathbb{G}_1$ is considerably cheaper than in $\mathbb{G}_2$; hence, we count separately exponentiations in both groups.

For $\kappa = 128$, the current recommendation is to use an optimal (asymmetric) Ate pairing over Barreto-Naehrig curves [1]. In that case, at security level of $\kappa = 128$, an element of $\mathbb{G}_1/\mathbb{G}_2/\mathbb{G}_T$ can be represented in respectively 256/512/3072 bits. To speed up interpolation, we will additionally need the existence of the n-th, where n is a power of 2, primitive root of unity modulo p (under this condition, one can interpolate in time $\Theta(n \log n)$, otherwise, interpolation takes time $\Theta(n \log^2 n)$). For this, it suffices that $(n+1) \mid (p-1)$ (recall that p is the elliptic curve group order). Fortunately, given κ and a practically relevant value of n, one can easily find a Barreto-Naehrig curve such that $(n+1) \mid (p-1)$ holds; such an observation was made also in [5]. For example, if $\kappa = 128$ and $n = 2^{10}$, one can use Algorithm 1 of [1] to find an elliptic curve group of prime order $N(x_0)$ over a finite field of prime order $P(-x_0)$ for $x_0 = 1753449050$, where $P(x) = 36x^4 + 36x^3 + 24x^2 + 6x + 1$, $T(x) = 6x^2 + 1$, and $N(x) = P(x) + 1 - T(x)$. One can then use the curve $E : y^2 = x^3 + 6$.

In proof bootstrapping [12], one needs an additional elliptic curve group $\tilde{E}$ over a finite field of order $N(x_0)$ (see [12] for additional details). Such elliptic curve group can be found by using the Cocks-Pinch method; note that $\tilde{E}$ has somewhat less efficient arithmetic than E.

The security of the new commitment scheme and of the new SNARKs depends on the following q-type assumptions, variants of which have been used in many previous papers. The assumptions are parameterized but non-interactive in the sense that q is related to the parameters of the language (most generally, to the input length) and not to the number of the adversarial queries. All known (to us) adaptive zk-SNARKs are based on q-type assumptions about BP.

Let $d(n) \in poly(n)$ be a function. Then, BP is

- *$d(n)$-PDL (Power Discrete Logarithm) secure* if for any $n \in poly(\kappa)$ and any non-uniform probabilistic polynomial-time (NUPPT) adversary A, $\Pr[\mathsf{gk} \leftarrow \mathsf{BP}(1^\kappa, n), (g_1, g_2, \chi) \leftarrow_r \mathbb{G}_1^* \times \mathbb{G}_2^* \times \mathbb{Z}_p : \mathsf{A}(\mathsf{gk}; ((g_1, g_2)^{\chi^i})_{i=0}^{d(n)}) = \chi] \approx_\kappa 0$.
- *n-TSDH (Target Strong Diffie-Hellman) secure* if for any $n \in poly(\kappa)$ and any NUPPT adversary A, $\Pr[\mathsf{gk} \leftarrow \mathsf{BP}(1^\kappa, n), (g_1, g_2, \chi) \leftarrow_r \mathbb{G}_1^* \times \mathbb{G}_2^* \times \mathbb{Z}_p : \mathsf{A}(\mathsf{gk}; ((g_1, g_2)^{\chi^i})_{i=0}^{n}) = (r, \hat{e}(g_1, g_2)^{1/(\chi - r)})] \approx_\kappa 0$.

For algorithms A and X_{A}, we write $(y; y') \leftarrow (\mathsf{A}||X_{\mathsf{A}})(\chi)$ if A on input χ outputs y, and X_{A} on the same input (including the random tape of A) outputs y'. We will need knowledge assumptions w.r.t. several knowledge secrets γ_i. Let m be the number of different knowledge secrets in any concrete SNARK. Let $\mathcal{F} = (P_i)_{i=0}^n$ be a tuple of univariate polynomials, and $\mathcal{G}_1$ (resp., $\mathcal{G}_2$) be a tuple of univariate (resp., m-variate) polynomials. Let $i \in [1\,..\,m]$. Then, BP is $(\mathcal{F}, \mathcal{G}_1, \mathcal{G}_2, i)$-*PKE (Power Knowledge of Exponent) secure* if for any NUPPT adversary A there exists an NUPPT extractor X_{A}, such that

$$\Pr\left[\begin{array}{l} \mathsf{gk} \leftarrow \mathsf{BP}(1^\kappa, n), (g_1, g_2, \chi, \boldsymbol{\gamma}) \leftarrow_r \mathbb{G}_1^* \times \mathbb{G}_2^* \times \mathbb{Z}_p \times \mathbb{Z}_p^m, \\ \boldsymbol{\gamma}_{-i} \leftarrow (\gamma_1, \ldots, \gamma_{i-1}, \gamma_{i+1}, \ldots, \gamma_m), \mathsf{aux} \leftarrow (g_1^{\mathcal{G}_1(\chi)}, g_2^{\mathcal{G}_2(\chi, \boldsymbol{\gamma}_{-i})}), \\ (h_1, h_2; (a_i)_{i=0}^n) \leftarrow (\mathsf{A}||X_{\mathsf{A}})(\mathsf{gk}; (g_1, g_2^{\gamma_i})^{\mathcal{F}(\chi)}, \mathsf{aux}) : \\ \hat{e}(h_1, g_2^{\gamma_i}) = \hat{e}(g_1, h_2) \wedge h_1 \neq g_1^{\sum_{i=0}^n a_i P_i(\chi)} \end{array}\right] \approx_\kappa 0.$$

Here, aux can be seen as the common auxiliary input to A and X_{A} that is generated by using benign auxiliary input generation. If $\mathcal{F} = (X^i)_{i=0}^d$ for some $d = d(n)$, then we replace the first argument in $(\mathcal{F}, \ldots)$-PKE with d. If $m = 1$, then we omit the last argument i in $(\mathcal{F}, \ldots, i)$-PKE. While knowledge assumptions are non-falsifiable, we recall that non-falsifiable assumptions are needed to design succincts SNARKs for interesting languages [18].

By generalizing [7,19,23], one can show that the TSDH, PDL, and PKE assumptions hold in the generic bilinear group model.

Within this paper, $m \leq 2$, and hence we denote γ_1 just by γ, and γ_2 by δ.

An extractable trapdoor commitment scheme in the CRS model consists of two efficient algorithms $\mathsf{G}_{\mathsf{com}}$ (that outputs a CRS ck and a trapdoor) and, (that, given ck, a message m and a randomizer r, outputs a commitment $\mathsf{C}_{\mathsf{ck}}(m; r)$), and must satisfy the following security properties.

Computational Binding: without access to the trapdoor, it is intractable to open a commitment to two different messages.

Trapdoor: given access to the original message, the randomizer and the trapdoor, one can open the commitment to any other message.

Perfect Hiding: commitments of any two messages have the same distribution.

Extractability: given access to the CRS, the commitment, and the random coins of the committer, one can open the commitment to the committed message.

See, e.g., [19] for formal definitions. In the context of the current paper, the message is a vector from $\mathbb{Z}_p^n$. We denote the randomizer space by $\mathfrak{R}$.

Let $\mathcal{R} = \{(u, w)\}$ be an efficiently verifiable relation with $|w| = \mathrm{poly}(|u|)$. Here, u is a statement, and w is a witness. Let $\mathcal{L} = \{u : \exists w, (u, w) \in \mathcal{R}\}$ be an **NP**-language. Let $n = |u|$ be the input length. For fixed n, we have a relation $\mathcal{R}_n$ and a language $\mathcal{L}_n$.

Following [9,20], we will define commit-and-prove (CaP) argument systems. Intuitively, a CaP non-interactive zero knowledge argument system for $\mathcal{R}$ allows to create a common reference string (CRS) crs, commit to some values w_i

(say, $u_i = \mathsf{C}_{\mathsf{ck}}(w_i; r_i)$, where ck is a part of crs), and then prove that a subset $u := (u_{i_j}, w_{i_j}, r_{i_j})_{j=1}^{\ell_m(n)}$ (for publicly known indices i_j) satisfies that u_{i_j} is a commitment of w_{i_j} with randomizer r_{i_j}, and that $(w_{i_j}) \in \mathcal{R}$.

Differently from most of the previous work (but see also [12]), our CaP argument systems will use computationally binding trapdoor commitment schemes. This means that without their openings, commitments $u_i = \mathsf{C}_{\mathsf{ck}}(a_i; r_i)$ themselves do not define a valid relation, since u_i can be a commitment to any a'_i, given a suitable r'_i. Rather, we define a new relation $\mathcal{R}_{\mathsf{ck}} := \{(\boldsymbol{u}, \boldsymbol{w}, \boldsymbol{r}) : (\forall i, u_i = \mathsf{C}_{\mathsf{ck}}(w_i; r_i)) \wedge \boldsymbol{w} \in \mathcal{R}\}$, and construct argument systems for $\mathcal{R}_{\mathsf{ck}}$.

Within this subsubsection, we let vectors $\boldsymbol{u}$, $\boldsymbol{w}$, and $\boldsymbol{r}$ be of dimension $\ell_m(n)$ for some polynomial $\ell_m(n)$. However, we allow committed messages w_i themselves to be vectors of dimension n. Thus, $\ell_m(n)$ is usually very small. In some argument systems (like the Subset-Sum SNARK in Sect. 6), also the argument will include some commitments. In such cases, technically speaking, $\boldsymbol{w}$ and $\boldsymbol{r}$ are of higher dimension than $\boldsymbol{u}$. To simplify notation, we will ignore this issue.

A *commit-and-prove non-interactive zero-knowledge argument system* [9, 20] Π for $\mathcal{R}$ consists of an ($\mathcal{R}$-independent) trapdoor commitment scheme $\Gamma = (\mathsf{G}_{\mathsf{com}}, \mathsf{C})$ and of a non-interactive zero-knowledge argument system $(\mathsf{G}, \mathsf{P}, \mathsf{V})$, that are combined as follows: 1. the CRS generator G (that, in particular, invokes $(\mathsf{ck}, \mathsf{td}_{\mathsf{C}}) \leftarrow \mathsf{G}_{\mathsf{com}}(1^\kappa, n)$) outputs $(\mathsf{crs} = (\mathsf{crs}_p, \mathsf{crs}_v), \mathsf{td}) \leftarrow \mathsf{G}(1^\kappa, n)$, where both crs_p and crs_v include ck, and td includes td_{C}. 2. the prover P produces an argument π, $\pi \leftarrow \mathsf{P}(\mathsf{crs}_p; \boldsymbol{u}; \boldsymbol{w}, \boldsymbol{r})$, where presumably $u_i = \mathsf{C}_{\mathsf{ck}}(w_i; r_i)$. 3. the verifier V, $\mathsf{V}(\mathsf{crs}_v; \boldsymbol{u}, \pi)$, outputs either 1 (accept) or 0 (reject). [(i)] Now, Π is *perfectly complete*, if for all $n = \mathrm{poly}(\kappa)$, $\Pr[(\mathsf{crs}, \mathsf{td}) \leftarrow \mathsf{G}(1^\kappa, n), (\boldsymbol{u}, \boldsymbol{w}, \boldsymbol{r}) \leftarrow \mathcal{R}_{\mathsf{ck},n} : \mathsf{V}(\mathsf{crs}_v; \boldsymbol{u}, \mathsf{P}(\mathsf{crs}_p; \boldsymbol{u}, \boldsymbol{w}, \boldsymbol{r})) = 1] = 1$.

Since Γ is computationally binding and trapdoor (and hence u_i can be commitments to *any* messages), soundness of the CaP argument systems only makes sense together with the argument of knowledge property.

Let $b(X)$ be a non-negative polynomial. Π is *a (b-bounded-auxiliary-input) argument of knowledge* for $\mathcal{R}$, if for all $n = \mathrm{poly}(\kappa)$ and every NUPPT A, there exists an NUPPT extractor X_{A}, such that for every auxiliary input $\mathsf{aux} \in \{0,1\}^{b(\kappa)}, \Pr[(\mathsf{crs}, \mathsf{td}) \leftarrow \mathsf{G}(1^\kappa, n), ((\boldsymbol{u}, \pi); \boldsymbol{w}, \boldsymbol{r}) \leftarrow (\mathsf{A}||X_{\mathsf{A}})(\mathsf{crs}; \mathsf{aux}) : (\boldsymbol{u}, \boldsymbol{w}, \boldsymbol{r}) \notin \mathcal{R}_{\mathsf{ck},n} \wedge \mathsf{V}(\mathsf{crs}_v; \boldsymbol{u}, \pi) = 1] \approx_\kappa 0$. As in the definition of PKE, we can restrict the definition of an argument of knowledge to benign auxiliary information generators, where aux is known to come from; we omit further discussion.

Π is *perfectly witness-indistinguishable*, if for all $n = \mathrm{poly}(\kappa)$, it holds that if $(\mathsf{crs}, \mathsf{td}) \in \mathsf{G}(1^\kappa, n)$ and $((\boldsymbol{u}; \boldsymbol{w}, \boldsymbol{r}), (\boldsymbol{u}; \boldsymbol{w}', \boldsymbol{r}')) \in \mathcal{R}^2_{\mathsf{ck},n}$ with $r_i, r'_i \leftarrow_r \mathfrak{R}$, then the distributions $\mathsf{P}(\mathsf{crs}_p; \boldsymbol{u}; \boldsymbol{w}, \boldsymbol{r})$ and $\mathsf{P}(\mathsf{crs}_p; \boldsymbol{u}; \boldsymbol{w}', \boldsymbol{r}')$ are equal. Note that a witness-indistinguishable argument system does not have to have a trapdoor.

Π is *perfectly composable zero-knowledge*, if there exists a probabilistic poly-time simulator S, s.t. for all stateful NUPPT adversaries A and $n = \mathrm{poly}(\kappa), \Pr(\mathsf{crs}, \mathsf{td}) \leftarrow \mathsf{G}(1^\kappa, n), (\boldsymbol{u}, \boldsymbol{w}, \boldsymbol{r}) \leftarrow \mathsf{A}(\mathsf{crs}), \pi \leftarrow \mathsf{P}(\mathsf{crs}_p; \boldsymbol{u}; \boldsymbol{w}, \boldsymbol{r}) : (\boldsymbol{u}, \boldsymbol{w}, \boldsymbol{r}) \in \mathcal{R}_{\mathsf{ck},n} \wedge \mathsf{A}(\pi) = 1] = \Pr[(\mathsf{crs}, \mathsf{td}) \leftarrow \mathsf{G}(1^\kappa, n), (\boldsymbol{u}, \boldsymbol{w}, \boldsymbol{r}) \leftarrow \mathsf{A}(\mathsf{crs}), \pi \leftarrow \mathsf{S}(\mathsf{crs}; \boldsymbol{u}, \mathsf{td}) : (\boldsymbol{u}, \boldsymbol{w}, \boldsymbol{r}) \in \mathcal{R}_{\mathsf{ck},n} \wedge \mathsf{A}(\pi) = 1]$. Here, the prover and the simulator use the same CRS, and thus we have *same-string zero knowledge*. Same-string

statistical zero knowledge allows to use the same CRS an unbounded number of times.

An argument system that satisfies above requirements is known as *adaptive*. An argument system where the CRS depends on the statement is often called *non-adaptive*. It is not surprising that non-adaptive SNARKs can be much more efficient than adaptive SNARKs.

A non-interactive argument system is *succinct* if the output length of P and the running time of V are polylogarithmic in the P's input length (and polynomial in the security parameter). A succinct non-interactive argument of knowledge is usually called *SNARK*. A zero-knowledge SNARK is abbreviated to *zk-SNARK*.

3 New Extractable Trapdoor Commitment Scheme

We now define a new extractable trapdoor commitment scheme. It uses the following polynomials. Assume n is a power of two, and let ω be the n-th primitive root of unity modulo p. Then,

- $Z(X) := \prod_{i=1}^{n}(X-\omega^{i-1}) = X^n - 1$ is the unique degree n monic polynomial, such that $Z(\omega^{i-1}) = 0$ for all $i \in [1\,..\,n]$.
- $\ell_i(X) := \prod_{j\neq i}((X-\omega^{j-1})/(\omega^{i-1}-\omega^{j-1}))$, the *ith Lagrange basis polynomial*, is the unique degree $n-1$ polynomial, such that $\ell_i(\omega^{i-1}) = 1$ and $\ell_i(\omega^{j-1}) = 0$ for $j \neq i$.

Clearly, $L_{\boldsymbol{a}}(X) = \sum_{i=1}^{n} a_i\ell_i(X)$ is the interpolating polynomial of $\boldsymbol{a}$ at points ω^{i-1}, with $L_{\boldsymbol{a}}(\omega^{i-1}) = a_i$, and can thus be computed by executing an inverse Fast Fourier Transform. Moreover, $(\ell_i(\omega^{j-1}))_{j=1}^{n} = \boldsymbol{e}_i$ (the ith unit vector) and $(Z(\omega^{j-1}))_{j=1}^{n} = \boldsymbol{0}_n$. Thus, $Z(X)$ and $(\ell_i(X))_{i=1}^{n}$ are $n+1$ linearly independent degree $\leq n$ polynomials, and hence $\mathcal{F}_{\mathsf{C}} := (Z(X), (\ell_i(X))_{i=1}^{n})$ is a basis of such polynomials. Clearly, $Z^{-1}(0) = \{j : Z(j) = 0\} = \{\omega^{i-1}\}_{i=1}^{n}$.

Definition 1 (Interpolating Commitment Scheme). *Let $n = \mathrm{poly}(\kappa)$, $n > 0$, be a power of two. First, $\mathsf{G}_{\mathsf{com}}(1^\kappa, n)$ sets $\mathsf{gk} \leftarrow \mathsf{BP}(1^\kappa, n)$, picks $g_1 \leftarrow_r \mathbb{G}_1^*$, $g_2 \leftarrow_r \mathbb{G}_2^*$, and then outputs the CRS $\mathsf{ck} \leftarrow (\mathsf{gk}; (g_1^{f(\chi)}, g_2^{\gamma f(\chi)})_{f\in\mathcal{F}_{\mathsf{C}}})$ for $\chi \leftarrow_r \mathbb{Z}_p \backslash Z^{-1}(0)$ and $\gamma \leftarrow_r \mathbb{Z}_p^*$. The trapdoor is equal to χ.*

The commitment of $\boldsymbol{a} \in \mathbb{Z}_p^n$, given a randomizer $r \leftarrow_r \mathbb{Z}_p$, is $\mathsf{C}_{\mathsf{ck}}(\boldsymbol{a}; r) := (g_1^{Z(\chi)}, g_2^{\gamma Z(\chi)})^r \cdot \prod_{i=1}^{n}(g_1^{\ell_i(\chi)}, g_2^{\gamma\ell_i(\chi)})^{a_i} \in \mathbb{G}_1 \times \mathbb{G}_2$, i.e., $\mathsf{C}_{\mathsf{ck}}(\boldsymbol{a}; r) := (g_1, g_2^{\gamma})^{r(\chi^n-1)+L_a(\chi)}$. The validity of a commitment (A_1, A_2^{γ}) is checked by verifying that $\hat{e}(A_1, g_2^{\gamma Z(\chi)}) = \hat{e}(g_1^{Z(\chi)}, A_2^{\gamma})$. To open a commitment, the committer sends $(\boldsymbol{a}, r)$ to the verifier.

The condition $Z(\chi) \neq 0$ is needed in Theorem 1 to get perfect hiding and the trapdoor property. The condition $\gamma \neq 0$ is only needed in Theorem 5 to get perfect zero knowledge. Also, (a function of) γ is a part of the trapdoor in the range SNARK of Sect. 7.

Clearly, $\log_{g_1} A_1 = \log_{g_2^{\gamma}} A_2^{\gamma} = rZ(\chi) + \sum_{i=1}^{n} a_i\ell_i(\chi)$. The second element, A_2^{γ}, of the commitment is known as the knowledge component.

Theorem 1. *The interpolating commitment scheme is perfectly hiding and trapdoor. If* BP *is* n*-PDL secure, then it is computationally binding. If* BP *is* $(n, \emptyset, \emptyset)$*-PKE secure, then it is extractable.*

Proof. PERFECT HIDING: since $Z(\chi) \neq 0$, then $rZ(\chi)$ (and thus also $\log_{g_1} A_1$) is uniformly random in $\mathbb{Z}_p$. Hence, (A_1, A_2^γ) is a uniformly random element of the multiplicative subgroup $\langle (g_1, g_2^\gamma) \rangle \subset \mathbb{G}_1^* \times \mathbb{G}_2^*$ generated by (g_1, g_2^γ), independently of the committed value. TRAPDOOR: given χ, $\boldsymbol{a}$, r, $\boldsymbol{a}^*$, and $c = \mathsf{C}_{\mathsf{ck}}(\boldsymbol{a}; r)$, we compute r^* s.t. $(r^* - r)Z(\chi) + \sum_{i=1}^n (a_i^* - a_i)\ell_i(\chi) = 0$. This is possible since $Z(\chi) \neq 0$. Clearly, $c = \mathsf{C}_{\mathsf{ck}}(\boldsymbol{a}^*; r^*)$. EXTRACTABILITY: clear from the statement.

COMPUTATIONAL BINDING: assume that there exists an adversary A_{C} that outputs $(\boldsymbol{a}, r_a)$ and $(\boldsymbol{b}, r_b)$ with $(\boldsymbol{a}, r_a) \neq (\boldsymbol{b}, r_b)$, s.t. the polynomial $d(X) := (r_a Z(X) + \sum_{i=1}^n a_i \ell_i(X)) - (r_b Z(X) + \sum_{i=1}^n b_i \ell_i(X))$ has a root at χ.

Construct now the following adversary A_{pdl} that breaks the PDL assumption. Given an n-PDL challenge, since $\mathcal{F}_{\mathsf{C}}$ consists of degree $\leq n$ polynomials, A_{pdl} can compute a valid ck from (a distribution that is statistically close to) the correct distribution. He sends ck to A_{C}. If A_{C} is successful, then $d(X) \in \mathbb{Z}_p[X]$ is a non-trivial degree-$\leq n$ polynomial. Since the coefficients of d are known, A_{pdl} can use an efficient polynomial factorization algorithm to compute all roots r_i of $d(X)$. One of these roots has to be equal to χ. A_{pdl} can establish which one by comparing each (say) $g_1^{\ell_1(r_i)}$ to the element $g_1^{\ell_1(\chi)}$ given in the CRS. Clearly, $g_1^{\ell_1(r_i)}$ is computed from g_1 (which can be computed, given the CRS, since $1 \in \mathrm{span}(\mathcal{F}_{\mathsf{C}})$), the coefficients of $\ell_1(X)$, and r_i. A_{pdl} has the same success probability as A_{C}, while her running time is dominated by that of A_{C} plus the time to factor a degree-$\leq n$ polynomial. □

Theorem 1 also holds when instead of $Z(X)$ and $\ell_i(X)$ one uses any $n+1$ linearly independent low-degree polynomials (say) $P_0(X)$ and $P_i(X)$. Given the statement of Theorem 1, this choice of the concrete polynomials is very natural: $\ell_i(X)$ interpolate linearly independent vectors (and thus are linearly independent; in fact, they constitute a basis), and the choice to interpolate unit vectors is the conceptually clearest way of choosing $P_i(X)$. Another natural choice of independent polynomials is to set $P_i(X) = X^i$ as in [19], but that choice has resulted in much less efficient (CaP) SNARKs.

In the full version [26] we show how to use batch-verification techniques to speed up simultaneous validity verification of many commitments.

4 New Product SNARK

Assume the use of the interpolating commitment scheme. In a *CaP product SNARK* [19], the prover aims to convince the verifier that she knows how to open three commitments (A, A^γ), (B, B^γ), and (C, C^γ) to vectors $\boldsymbol{a}$, $\boldsymbol{b}$ and $\boldsymbol{c}$ (together with the used randomizers), such that $\boldsymbol{a} \circ \boldsymbol{b} = \boldsymbol{c}$. Thus,

$$\mathcal{R}_{\mathsf{ck},n}^{\times} := \left\{ \begin{array}{l} (u_\times, w_\times, r_\times) : u_\times = ((A_1, A_2^\gamma), (B_1, B_2^\gamma), (C_1, C_2^\gamma)) \wedge \\ w_\times = (\boldsymbol{a}, \boldsymbol{b}, \boldsymbol{c}) \wedge r_\times = (r_a, r_b, r_c) \wedge (A_1, A_2^\gamma) = \mathsf{C}_{\mathsf{ck}}(\boldsymbol{a}; r_a) \wedge \\ (B_1, B_2^\gamma) = \mathsf{C}_{\mathsf{ck}}(\boldsymbol{b}; r_b) \wedge (C_1, C_2^\gamma) = \mathsf{C}_{\mathsf{ck}}(\boldsymbol{c}; r_c) \wedge \boldsymbol{a} \circ \boldsymbol{b} = \boldsymbol{c} \end{array} \right\}.$$

Next, we propose an efficient CaP product SNARK. For this, we need Lemma 1.

Lemma 1. *Let $A(X)$, $B(X)$ and $C(X)$ be polynomials with $A(\omega^{i-1}) = a_i$, $B(\omega^{i-1}) = b_i$ and $C(\omega^{i-1}) = c_i$, $\forall i \in [1\,..\,n]$. Let $Q(X) = A(X)B(X) - C(X)$. Assume that (i) $A(X), B(X), C(X) \in \mathrm{span}\{\ell_i(X)\}_{i=1}^n$, and (ii) there exists a degree $n-2$ polynomial $\pi(X)$, s.t. $\pi(X) = Q(X)/Z(X)$. Then $\boldsymbol{a} \circ \boldsymbol{b} = \boldsymbol{c}$.*

Proof. From (i) it follows that $A(X) = L_{\boldsymbol{a}}(X)$, $B(X) = L_{\boldsymbol{b}}(X)$, and $C(X) = L_{\boldsymbol{c}}(X)$, and thus $Q(\omega^{i-1}) = a_i b_i - c_i$ for all $i \in [1\,..\,n]$. But (ii) iff $Z(X) \mid Q(X)$, which holds iff $Q(X)$ evaluates to 0 at all n values ω^{i-1}. Thus, $\boldsymbol{a} \circ \boldsymbol{b} = \boldsymbol{c}$. Finally, if (i) holds then $\deg Q(X) = 2n-2$ and thus $\deg \pi(X) = n-2$. □

If privacy and succinctness are not needed, one can think of the product argument being equal to $\pi(X)$. We achieve privacy by picking $r_a, r_b, r_c \leftarrow_r \mathbb{Z}_p$, and defining $Q_{wi}(X) := (L_{\boldsymbol{a}}(X) + r_a Z(X))\,(L_{\boldsymbol{b}}(X) + r_b Z(X)) - (L_{\boldsymbol{c}}(X) + r_c Z(X))$. Here, the new addends of type $r_a Z(X)$ guarantee hiding. On the other hand, $Q_{wi}(X)$ remains divisible by $Z(X)$ iff $\boldsymbol{c} = \boldsymbol{a} \circ \boldsymbol{b}$. Thus, $\boldsymbol{a} \circ \boldsymbol{b} = \boldsymbol{c}$ iff

(i') $Q_{wi}(X)$ can be expressed as $Q_{wi}(X) = A(X)B(X) - C(X)$ for some polynomials $A(X)$, $B(X)$ and $C(X)$ that belong to the span of $\mathcal{F}_{\mathsf{C}}$, and
(ii') there exists a polynomial $\pi_{wi}(X)$, such that

$$\pi_{wi}(X) = Q_{wi}(X)/Z(X). \tag{1}$$

The degree of $Q_{wi}(X)$ is $2n$, thus, if $\pi_{wi}(X)$ exists, then it has degree n.

However, $|\pi_{wi}(X)|$ is not sublinear in n. To minimize communication, we let the prover transfer a "garbled" evaluation of $\pi_{wi}(X)$ at a random secret point χ. More precisely, the prover computes $\pi_\times := g_1^{\pi_{wi}(\chi)}$, using the values $g_1^{\chi^i}$ (given in the CRS) and the coefficients π_i of $\pi_{wi}(X) = \sum_{i=0}^n \pi_i X^i$, as follows:

$$\pi_\times := g_1^{\pi_{wi}(\chi)} \leftarrow \prod_{i=0}^{n} (g_1^{\chi^i})^{\pi_i}. \tag{2}$$

Similarly, instead of (say) $L_{\boldsymbol{a}}(X) + r_a Z(X)$, the verifier has the succinct interpolating commitment $\mathsf{C}_{\mathsf{ck}}(\boldsymbol{a}; r_a) = (g_1, g_2^\gamma)^{L_{\boldsymbol{a}}(\chi) + r_a Z(\chi)}$ of $\boldsymbol{a}$.

We now give a full description of the new product SNARK $\Pi_\times$, given the interpolating commitment scheme $(\mathsf{G}_{\mathsf{com}}, \mathsf{C})$ and the following tuple of algorithms, $(\mathsf{G}_\times, \mathsf{P}_\times, \mathsf{V}_\times)$. Note that $\mathsf{C}_{\mathsf{ck}}(\boldsymbol{1}_{\boldsymbol{n}}; 0) = (g_1, g_2^\gamma)$.

CRS Generation: $\mathsf{G}_\times(1^\kappa, n)$**:** Let $\mathsf{gk} \leftarrow \mathsf{BP}(1^\kappa)$, $(g_1, g_2, \chi, \gamma) \leftarrow_r \mathbb{G}_1^* \times \mathbb{G}_2^* \times \mathbb{Z}_p^2$ with $Z(\chi) \neq 0$ and $\gamma \neq 0$. Let $\mathsf{crs}_p = \mathsf{ck} \leftarrow (\mathsf{gk}; (g_1, g_2^\gamma)^{\mathcal{F}_{\mathsf{C}}(\chi)})$ and $\mathsf{crs}_v \leftarrow (\mathsf{gk}; g_2^{\gamma Z(\chi)})$. Output $\mathsf{crs}_\times = (\mathsf{crs}_p, \mathsf{crs}_v)$.
Common Input: $u_\times = ((A_1, A_2^\gamma), (B_1, B_2^\gamma), (C_1, C_2^\gamma))$.
Proving: $\mathsf{P}_\times(\mathsf{crs}_p; u_\times; w_\times = (\boldsymbol{a}, \boldsymbol{b}, \boldsymbol{c}), r_\times = (r_a, r_b, r_c))$**:** Compute $\pi_{wi}(X) = \sum_{i=0}^n \pi_i X^i$ as in Eq. (1) and $\pi_\times$ as in Eq. (2). Output $\pi_\times$.
Verification: $\mathsf{V}_\times(\mathsf{crs}_v; u_\times; \pi_\times)$**:** accept if $\hat{e}(A_1, B_2^\gamma) = \hat{e}(g_1, C_2^\gamma) \cdot \hat{e}(\pi_\times, g_2^{\gamma Z(\chi)})$.

Since one can recompute it from ck, inclusion of $g_2^{\gamma Z(\chi)}$ in the CRS is only needed to speed up the verification. Here as in the shift SNARK of Sect. 5, validity of the commitments will be verified in the master SNARK. This is since the master SNARKs use some of the commitments in several sub-SNARKs, while it suffices to verify the validity of every commitment only once.

To obtain an argument of knowledge, we use knowledge assumptions in all following proofs. This SNARK is not zero-knowledge since the possible simulator gets three commitments as inputs but not their openings; to create an accepting argument the simulator must at least know how to open the commitment $(A_1B_1/C_1, A_2^{\gamma}B_2^{\gamma}/C_2^{\gamma})$ to $\boldsymbol{a} \circ \boldsymbol{b} - \boldsymbol{c}$. It is witness-indistinguishable, and this suffices for the Subset-Sum and other master SNARKs to be zero-knowledge.

Theorem 2. *$\Pi_{\times}$ is perfectly complete and witness-indistinguishable. If the input consists of valid commitments, and BP is n-TSDH and $(n, \emptyset, \emptyset)$-PKE secure, then $\Pi_{\times}$ is an ($\Theta(n)$-bounded-auxiliary-input) adaptive argument of knowledge.*

Proof. Perfect completeness: follows from the discussion in the beginning of this section. Perfect witness-indistinguishability: since the argument $\pi_{\times}$ that satisfies the verification equations is unique, all witnesses result in the same argument, and thus this argument is witness-indistinguishable.

Argument of knowledge: Assume that A_{aok} is an adversary that, given $\mathsf{crs}_{\times}$, returns $(u_{\times}, \pi)$ such that $\mathsf{V}_{\times}(\mathsf{crs}_v; u_{\times}, \pi) = 1$. Assume that the PKE assumption holds, and let X_{A} be the extractor that returns openings of the commitments in $u_{\times}$, i.e., $(\boldsymbol{a}, r_a)$, $(\boldsymbol{b}, r_b)$, and $(\boldsymbol{c}, r_c)$. We now claim that X_{A} is also the extractor needed to achieve the argument of knowledge property.

Assume that this is not the case. We construct an adversary A_{tsdh} against n-TSDH. Given an n-TSDH challenge $ch = (\mathsf{gk}, ((g_1, g_2)^{\chi^i})_{i=0}^n)$, A_{tsdh} first generates $\gamma \leftarrow_r \mathbb{Z}_p^*$, and then computes (this is possible since $\mathcal{F}_{\mathsf{C}}$ consists of degree $\leq n$ polynomials) and sends $\mathsf{crs}_{\times}$ to A_{aok}. Assume $(\mathsf{A}_{aok}||X_{\mathsf{A}})(\mathsf{crs}_{\times})$ returns $((u_{\times} = ((A_1, A_2^{\gamma}), (B_1, B_2^{\gamma}), (C_1, C_2^{\gamma})), \pi), (w_{\times} = (\boldsymbol{a}, \boldsymbol{b}, \boldsymbol{c}), r_{\times} = (r_a, r_b, r_C)))$, s.t. $u_i = \mathsf{C}_{\mathsf{ck}}(w_i; r_i)$ but $(u_{\times}, w_{\times}, r_{\times}) \notin \mathcal{R}^{\times}_{\mathsf{ck},n}$. Since the openings are correct, $\boldsymbol{a} \circ \boldsymbol{b} \neq \boldsymbol{c}$ but π is accepting. According to Lemma 1, thus $Z(X) \nmid Q_{wi}(X)$.

Since $Z(X) \nmid Q_{wi}(X)$, then for some $i \in [1\,..\,n]$, $(X - \omega^{i-1}) \nmid Q_{wi}(X)$. Write $Q_{wi}(X) = q(X)(X - \omega^{i-1}) + r$ for $r \in \mathbb{Z}_p^*$. Clearly, $\deg q(X) \leq 2n - 1$. Moreover, we write $q(X) = q_1(X)Z(X) + q_2(X)$ with $\deg q_i(X) \leq n - 1$. Since the verification succeeds, $\hat{e}(g_1, g_2^{\gamma})^{Q_{wi}(\chi)} = \hat{e}(\pi_{\times}, g_2^{\gamma Z(\chi)})$, or $\hat{e}(g_1, g_2^{\gamma})^{q(\chi)(\chi - \omega^{i-1}) + r} = \hat{e}(\pi_{\times}, g_2^{\gamma Z(\chi)})$, or $\hat{e}(g_1, g_2^{\gamma})^{q(\chi) + r/(\chi - \omega^{i-1})} = \hat{e}(\pi_{\times}, g_2^{\gamma Z(\chi)/(\chi - \omega^{i-1})})$, or $\hat{e}(g_1, g_2^{\gamma})^{1/(\chi - \omega^{i-1})} = (\hat{e}(\pi_{\times}, g_2^{\gamma Z(\chi)/(\chi - \omega^{i-1})}) / \hat{e}(g_1^{q(\chi)}, g_2^{\gamma}))^{r^{-1}}$.

Now, $\hat{e}(g_1^{q(\chi)}, g_2^{\gamma}) = \hat{e}(g_1^{q_1(\chi)}, g_2^{\gamma Z(\chi)}) \hat{e}(g_1^{q_2(\chi)}, g_2^{\gamma})$, and thus it can be efficiently computed from $((g_1^{\chi^i})_{i=0}^{n-1}, g_2^{\gamma}, g_2^{\gamma Z(\chi)}) \subset \mathsf{crs}$. Moreover, $Z(X)/(X - \omega^{i-1}) = \ell_i(X) \cdot \prod_{j \neq i}(\omega^{i-1} - \omega^{j-1})$, and thus $g_2^{\gamma Z(\chi)/(\chi - \omega^{i-1})}$ can be computed from $g_2^{\gamma \ell_i(\chi)}$ by using generic group operations. Hence, $\hat{e}(g_1, g_2^{\gamma})^{1/(\chi - \omega^{i-1})}$ can be

computed from $((g_1^{\chi^i})_{i=0}^{n-1}, g_2^{\gamma}, g_2^{\gamma Z(\chi)}, (g_2^{\gamma \ell_i(\chi)})_{i=1}^{n})$ (that can be computed from ch), by using generic group operations. Thus, the adversary has computed $(r = \omega^{i-1}, \hat{e}(g_1, g_2^{\gamma})^{1/(\chi - r)})$, for $r \neq \chi$. Since A_{tsdh} knows $\gamma \neq 0$, he can finally compute $(r, \hat{e}(g_1, g_2)^{1/(\chi - r)})$, and thus break the n-TSDH assumption.

Hence, the argument of knowledge property follows. □

We remark that the product SNARK (but not the shift SNARK of Sect. 5) can be seen as a QAP-based SNARK [17], namely for the relation $\boldsymbol{a} \circ \boldsymbol{b} - \boldsymbol{c}$. (Constructing a QAP-based shift SNARK is possible, but results in using different polynomials and thus in a different commitment scheme.)

The prover computation is dominated by the following: (i) one $(n+1)$-wide multi-exponentiation in $\mathbb{G}_1$. By using the Pippenger's multi-exponentiation algorithm for *large* n this means approximately $n+1$ bilinear-group multiplications, see [29]. For small values of n, one can use the algorithm by Straus [32]; then one has to execute $\Theta(n/\log n)$ bilinear-group exponentiations. (ii) three polynomial interpolations, one polynomial multiplication, and one polynomial division to compute the coefficients of the polynomial $\pi_{wi}(X)$. Since polynomial division can be implemented as 2 polynomial multiplications (by using pre-computation and storing some extra information in the CRS, [24]), this part is dominated by two inverse FFT-s and three polynomial multiplications.

The verifier computation is dominated by 3 pairings. (We will count the cost of validity verifications separately in the master SNARKs.) In the special case $C_1 = A_1$ (e.g., in the *Boolean SNARK*, where we need to prove that $\boldsymbol{a} \circ \boldsymbol{a} = \boldsymbol{a}$, or in the *restriction SNARK* [19], where we need to prove that $\boldsymbol{a} \circ \boldsymbol{b} = \boldsymbol{a}$ for a *public* Boolean vector $\boldsymbol{b}$), the verification equation can be simplified to $\hat{e}(A_1, B_2^{\gamma}/g_2^{\gamma}) = \hat{e}(\pi_{\times}, g_2^{\gamma Z(\chi)})$, which saves one more pairing. In the full version [26], we will describe a batch-verification technique that allows to speed up simultaneous verification of several product SNARKs.

Excluding gk, the prover CRS together with ck consists of $2(n+1)$ group elements, while the verifier CRS consists of 1 group element. The CRS can be computed in time $\Theta(n)$, by using an algorithm from [3].

5 New Shift SNARK

In a *shift-right-by-z* SNARK [15] (shift SNARK, for short), the prover aims to convince the verifier that for 2 commitments (A, A^{γ}) and (B, B^{γ}), he knows how to open them as $(A, A^{\gamma}) = \mathsf{C}_{\mathsf{ck}}(\boldsymbol{a}; r_a)$ and $(B, B^{\gamma}) = \mathsf{C}_{\mathsf{ck}}(\boldsymbol{b}; r_b)$, s.t. $\boldsymbol{a} = \boldsymbol{b} \gg z$. I.e., $a_i = b_{i+z}$ for $i \in [1 \,..\, n-z]$ and $a_i = 0$ for $i \in [n-z+1 \,..\, n]$. Thus,

$$\mathcal{R}_{\mathsf{ck},n}^{\mathsf{rsft}} := \left\{ \begin{array}{l} (u_{\times}, w_{\times}, r_{\times}) : u_{\times} = ((A_1, A_2^{\gamma}), (B_1, B_2^{\gamma})) \wedge w_{\times} = (\boldsymbol{a}, \boldsymbol{b}) \wedge \\ r_{\times} = (r_a, r_b) \wedge (A_1, A_2^{\gamma}) = \mathsf{C}_{\mathsf{ck}}(\boldsymbol{a}; r_a) \wedge \\ (B_1, B_2^{\gamma}) = \mathsf{C}_{\mathsf{ck}}(\boldsymbol{b}; r_b) \wedge (\boldsymbol{a} = \boldsymbol{b} \gg z) \end{array} \right\}.$$

An efficient shift SNARK was described in [15]. We now reconstruct this SNARK so that it can be used together with the interpolating commitment

scheme. We can do it since the shift SNARK of [15] is *almost* independent of the commitment scheme. We also slightly optimize the resulting SNARK; in particular, the verifier has to execute one less pairing compared to [15].

Our strategy of constructing a shift SNARK follows the strategy of [19,23]. We start with a concrete verification equation that also contains the argument, that we denote by π_1. We write the discrete logarithm of π_1 (that follows from this equation) as $F_\pi(\chi) + F_{con}(\chi)$, where χ is a secret key, and $F_\pi(X)$ and $F_{con}(X)$ are two polynomials. The first polynomial, $F_\pi(X)$, is identically zero iff the prover is honest. Since the spans of certain two polynomial sets do not intersect, this results in an efficient adaptive shift SNARK that is an argument of knowledge under (two) PKE assumptions.

Now, for a non-zero polynomial $Z^*(X)$ to be defined later, consider the verification equation $\hat{e}(A_1, g_2^{\gamma Z^*(\chi)})/\hat{e}(B_1\pi_1, g_2^{\gamma}) = 1$ (due to the properties of pairing, this is equivalent to verifying that $\pi_1 = A_1^{Z^*(\chi)}/B_1$), with (A_1, A_2^{γ}) and (B_1, B_2^{γ}) being interpolating commitments to $\boldsymbol{a}$ and $\boldsymbol{b}$, and $\pi_1 = g_1^{\pi(\chi)}$ for some polynomial $\pi(X)$. Denote $r(X) := (r_a Z^*(X) - r_b)Z(X)$. Taking a discrete logarithm of the verification equation, we get that $\pi(X) = (r_a Z(X) + \sum_{i=1}^{n} a_i \ell_i(X))\, Z^*(X) - (r_b Z(X) + \sum_{i=1}^{n} b_i \ell_i(X)) = Z^*(X) \sum_{i=1}^{n} a_i \ell_i(X) - \sum_{i=1}^{n} b_i \ell_i(X) + r(X) = \left(\sum_{i=1}^{n-z} a_i \ell_i(X) + \sum_{i=n-z+1}^{n} a_i \ell_i(X)\right) Z^*(X) + r(X) - \sum_{i=1}^{n-z} b_{i+z} \ell_{i+z}(X) - \sum_{i=1}^{z} b_i \ell_i(X)$. Hence, $\pi(X) = F_\pi(X) + F_{con}(X)$, where

$$F_\pi(X) = \left(\sum_{i=1}^{n-z} (a_i - b_{i+z})\ell_i(X) + \sum_{i=n-z+1}^{n} a_i \ell_i(X)\right) \cdot Z^*(X),$$
$$F_{con}(X) = \left(\sum_{i=z+1}^{n} b_i (\ell_{i-z}(X) Z^*(X) - \ell_i(X)) - \sum_{i=1}^{z} b_i \ell_i(X)\right) + r(X).$$

Clearly, the prover is honest iff $F_\pi(X) = 0$, which holds iff $\pi(X) = F_{con}(X)$, i.e., $\pi(X)$ belongs to the span of $\mathcal{F}_{z-\mathsf{rsft}} := (\ell_{i-z}(X)Z^*(X) - \ell_i(X))_{i=z+1}^{n}, (\ell_i(X))_{i=1}^{z}, Z(X)Z^*(X), Z(X))$. For the shift SNARK to be an argument of knowledge, we need that

(i) $(\ell_i(X)Z^*(X))_{i=1}^{n}$ is linearly independent, and
(ii) $F_\pi(X) \cap \mathrm{span}(\mathcal{F}_{z-\mathsf{rsft}}) = \emptyset$.

Together, (i) and (ii) guarantee that from $\pi(X) \in \mathrm{span}(\mathcal{F}_{z-\mathsf{rsft}})$ it follows that $\boldsymbol{a}$ is a shift of $\boldsymbol{b}$.

We guarantee that $\pi(X) \in \mathrm{span}(\mathcal{F}_{z-\mathsf{rsft}})$ by a knowledge assumption (w.r.t. another knowledge secret δ); for this we will also show that $\mathcal{F}_{z-\mathsf{rsft}}$ is linearly independent. As in the case of the product SNARK, we also need that (A_1, A_2^{γ}) and (B_1, B_2^{γ}) are actually commitments of n-dimensional vectors (w.r.t. γ), i.e., we rely on two PKE assumptions.

Denote $\mathcal{F}_\pi := \{\ell_i(X)Z^*(X)\}_{i=1}^{n}$. For a certain choice of $Z^*(X)$, both (i) and (ii) follow from the next lemma.

Lemma 2. *Let $Z^*(X) = Z(X)^2$. Then $\mathcal{F}_\pi \cup \mathcal{F}_{z-\mathsf{rsft}}$ is linearly independent.*

Proof. Assume that there exist $\boldsymbol{a} \in \mathbb{Z}_p^n$, $\boldsymbol{b} \in \mathbb{Z}_p^n$, $c \in \mathbb{Z}_p$, and $d \in \mathbb{Z}_p$, s.t. $f(X) := \sum_{i=1}^{n} a_i \ell_i(X) Z^*(X) + \sum_{i=z+1}^{n} b_i \left(\ell_{i-z}(X) Z^*(X) - \ell_i(X)\right) -$

$\sum_{i=1}^{z} b_i \ell_i(X) + cZ(X)Z^*(X) + dZ(X) = 0$. But then also $f(\omega^{j-1}) = 0$, for $j \in [1\,..\,n]$. Thus, due to the definition of $\ell_i(X)$ and $Z(X)$, $\sum_{i=1}^{n} b_i \boldsymbol{e}_i = \boldsymbol{0}_n$ which is only possible if $b_i = 0$ for all $i \in [1\,..\,n]$. Thus also $f'(X) := f(X)/Z(X) = \sum_{i=1}^{n} a_i \ell_i(X) Z^*(X)/Z(X) + cZ^*(X) + d = 0$. But then also $f'(\omega^{j-1}) = 0$ for $j \in [1\,..\,n]$. Hence, $cZ^*(\omega^{j-1}) + d = d = 0$. Finally, $f''(X) := f(X)/Z^*(X) = \sum_{i=1}^{n} a_i \ell_i(X) + cZ(X) = 0$, and from $f''(\omega^{j-1}) = 0$ for $j \in [1\,..\,n]$, we get $\boldsymbol{a} = \boldsymbol{0}_n$. Thus also $c = 0$. This finishes the proof. □

Since the argument of knowledge property of the new shift SNARK relies on $\pi(X)$ belonging to a certain span, similarly to [15], we will use an additional knowledge assumption. That is, it is necessary that there exists an extractor that outputs a witness that $\pi(X) = F_{con}(X)$ belongs to the span of $\mathcal{F}_{z-\mathsf{rsft}}$.

Similarly to the product SNARK, the shift SNARK does not contain $\pi(X) = F_{con}(X)$, but the value $\pi_{\mathsf{rsft}} = (g_1, g_2^{\delta})^{\pi(\chi)}$ for random χ and δ (necessary due to the use of the second PKE assumption), computed as

$$\begin{aligned}\pi_{\mathsf{rsft}} \leftarrow & (\pi_1, \pi_2^{\delta}) = (g_1, g_2^{\delta})^{\pi(\chi)} \\ = & \textstyle\prod_{i=z+1}^{n} ((g_1, g_2^{\delta})^{\ell_{i-z}(\chi)Z^*(\chi) - \ell_i(\chi)})^{b_i} \cdot \prod_{i=1}^{z} ((g_1, g_2^{\delta})^{\ell_i(\chi)})^{-b_i} \cdot \\ & ((g_1, g_2^{\delta})^{Z(\chi)Z^*(\chi)})^{r_a} \cdot ((g_1, g_2^{\delta})^{Z(\chi)})^{-r_b}.\end{aligned} \tag{3}$$

We are now ready to state the new shift-right-by-z SNARK Π_{rsft}. It consists of the interpolating commitment scheme and of the following three algorithms:

CRS Generation: $\mathsf{G}_{\mathsf{rsft}}(1^{\kappa}, n)$**:** Let $Z^*(X) = Z(X)^2$. Let $\mathsf{gk} \leftarrow \mathsf{BP}(1^{\kappa})$, $(g_1, g_2, \chi, \gamma, \delta) \leftarrow \mathbb{G}_1^* \times \mathbb{G}_2^* \times \mathbb{Z}_p^3$, s.t. $Z(\chi) \neq 0$, $\gamma \neq 0$. Set $\mathsf{ck} \leftarrow (\mathsf{gk}; (g_1, g_2^{\gamma})^{\mathcal{F}_{\mathsf{C}}(\chi)})$, $\mathsf{crs}_p \leftarrow (\mathsf{gk}; (g_1, g_2^{\delta})^{\mathcal{F}_{z-\mathsf{rsft}}(\chi)})$, $\mathsf{crs}_v \leftarrow (\mathsf{gk}; (g_1, g_2^{\delta})^{Z(\chi)}, g_2^{\delta Z(\chi) Z^*(\chi)})$. Return $\mathsf{crs}_{\mathsf{rsft}} = (\mathsf{ck}, \mathsf{crs}_p, \mathsf{crs}_v)$.
Common Input: $u_{\mathsf{rsft}} = ((A_1, A_2^{\gamma}), (B_1, B_2^{\gamma}))$.
Proving: $\mathsf{P}_{\mathsf{rsft}}(\mathsf{crs}_p; u_{\mathsf{rsft}}; w_{\mathsf{rsft}} = (\boldsymbol{a}, \boldsymbol{b}), r_{\mathsf{rsft}} = (r_a, r_b))$**:** return $\pi_{\mathsf{rsft}} \leftarrow (\pi_1, \pi_2^{\delta})$ from Eq. (3).
Verification: $\mathsf{V}_{\mathsf{rsft}}(\mathsf{crs}_v; u_{\mathsf{rsft}}; \pi_{\mathsf{rsft}} = (\pi_1, \pi_2^{\delta}))$**:** accept if $\hat{e}(\pi_1, g_2^{\delta Z(\chi)}) = \hat{e}(g_1^{Z(\chi)}, \pi_2^{\delta})$ and $\hat{e}(B_1 \pi_1, g_2^{\delta Z(\chi)}) = \hat{e}(A_1, g_2^{\delta Z(\chi) Z^*(\chi)})$.

Since crs_v can be recomputed from $\mathsf{ck} \cup \mathsf{crs}_p$, then clearly it suffices to take CRS to be $\mathsf{crs}_{\mathsf{rsft}} = (\mathsf{gk}; g_1^{\mathcal{F}_{\mathsf{C}}(\chi) \cup \mathcal{F}_{z-\mathsf{rsft}}(\chi)}, g_2^{\gamma \mathcal{F}_{\mathsf{C}}(\chi) \cup \delta \mathcal{F}_{z-\mathsf{rsft}}(\chi)})$.

Theorem 3. *Let $Z^*(X) = Z(X)^2$, $y = \deg(Z(X)Z^*(X)) = 3n$. Π_{rsft} is perfectly complete and witness-indistinguishable. If the input consists of valid commitments, and BP is y-PDL, $(n, \mathcal{F}_{z-\mathsf{rsft}}, Y_2 \mathcal{F}_{z-\mathsf{rsft}}, 1)$-PKE, and $(\mathcal{F}_{z-\mathsf{rsft}}, \mathcal{F}_{\mathsf{C}}, Y_1 \mathcal{F}_{\mathsf{C}}, 2)$-PKE secure, then Π_{rsft} is an $(\Theta(n)$-bounded-auxiliary-input) adaptive argument of knowledge.*

The prover computation is dominated by two $(n+2)$-wide multi-exponentiations (one in $\mathbb{G}_1$ and one in $\mathbb{G}_2$); there is no need for polynomial interpolation, multiplication or division. The communication is 2 group elements. The verifier computation is dominated by 4 pairings. In the full version [26], we describe a

batch-verification technique that allows to speed up simultaneous verification of several shift SNARKs. Apart from gk, the prover CRS and ck together contain $4n+6$ group elements, and the verifier CRS contains 3 group elements.

A shift-left-by-z (necessary in [25] to construct a permutation SNARK) SNARK can be constructed similarly. A rotation-left/right-by-z SNARK (one committed vector is a *rotation* of another committed vector) requires only small modifications, see [15].

6 New Subset-Sum SNARK

For fixed n and $p = n^{\omega(1)}$, the NP-complete language SUBSET-SUM over $\mathbb{Z}_p$ is defined as the language $\mathcal{L}_n^{\text{Subset-Sum}}$ of tuples $(\boldsymbol{S} = (S_1, \ldots, S_n), s)$, with $S_i, s \in \mathbb{Z}_p$, such that there exists a vector $\boldsymbol{b} \in \{0,1\}^n$ with $\sum_{i=1}^n S_i b_i = s$ in $\mathbb{Z}_p$. SUBSET-SUM can be solved in pseudo-polynomial time $O(pn)$ by using dynamic programming. In the current paper, since $n = \kappa^{o(1)}$ and $p = 2^{O(\kappa)}$, pn is not polynomial in the input size $n \log_2 p$.

In a SUBSET-SUM SNARK, the prover aims to convince the verifier that he knows how to open commitment (B_1, B_2^γ) to a vector $\boldsymbol{b} \in \{0,1\}^n$, such that $\sum_{i=1}^n S_i b_i = s$. We show that by using the new product and shift SNARKs, one can design a prover-efficient adaptive SUBSET-SUM zk-SNARK Π_{ssum}. We emphasize that SUBSET-SUM is just one of the languages for which we can construct an efficient zk-SNARK; Sect. 7 and the full version [26] have more examples.

First, we use the interpolating commitment scheme. The CRS generation $\mathsf{G}_{\mathsf{ssum}}$ invokes CRS generations of the commitment scheme, the product SNARK and the shift SNARK, sharing the same gk, g_1, g_2, γ, and trapdoor $\mathsf{td} = \chi$ between the different invocations. (Since here the argument must be zero knowledge, it needs a trapdoor.) Thus, $\mathsf{crs}_{\mathsf{ssum}} = \mathsf{crs}_{\mathsf{rsft}}$ for $z = 1$.

Let $\boldsymbol{e}_i$ be the ith unit vector. The prover's actions are depicted by Fig. 1 (a precise explanation of this SNARK will be given in the completeness proof in Theorem 4). This SNARK, even without taking into account the differences in the product and shift SNARKs, is both simpler and moth efficient than the

Let $\boldsymbol{b} \in \{0,1\}^n$ be such that $\sum_{i=1}^n S_i b_i = s$.
Let (B_1, B_2^γ) be a commitment to $\boldsymbol{b}$.
Construct a product argument π_1 to show that $\boldsymbol{b} \circ \boldsymbol{b} = \boldsymbol{b}$.
Let (C_1, C_2^γ) be a commitment to $\boldsymbol{c} \leftarrow \boldsymbol{S} \circ \boldsymbol{b}$.
Construct a product argument π_2 to show that $\boldsymbol{c} = \boldsymbol{S} \circ \boldsymbol{b}$.
Let (D_1, D_2^γ) be a commitment to $\boldsymbol{d}$, where $d_i = \sum_{j \geq i} c_j$.
Construct a shift-right-by-1 argument $(\pi_{31}, \pi_{32}^\delta)$ to show that $\boldsymbol{d} = (\boldsymbol{d} - \boldsymbol{c}) \gg 1$.
Construct a product argument π_4 to show that $\boldsymbol{e}_1 \circ (\boldsymbol{d} - s\boldsymbol{e}_1) = \boldsymbol{0}_n$.
Output $\pi_{\mathsf{ssum}} = (B_1, B_2^\gamma, C_1, C_2^\gamma, D_1, D_2^\gamma, \pi_1, \pi_2, \pi_{31}, \pi_{32}^\delta, \pi_4)$.

Fig. 1. The new SUBSET-SUM SNARK Π_{ssum} (prover's operations)

Subset-Sum SNARK presented in [15] where one needed an additional step of proving that $\boldsymbol{b} \neq \boldsymbol{0}_n$.

We remark that the vector $\boldsymbol{d}$, with $d_i = \sum_{j \geq i} c_j$, is called either a *vector scan*, an *all-prefix-sums*, or a *prefix-sum* of $\boldsymbol{c}$, and $(\pi_{31}, \pi_{32}^{\delta})$ can be thought of as a *scan SNARK* [15] that $\boldsymbol{d}$ is a correct scan of $\boldsymbol{c}$.

After receiving π_{ssum}, the verifier computes $S_1' \leftarrow \prod_i (g_1^{\ell_i(\chi)})^{S_i}$ as the first half of a commitment to $\boldsymbol{S}$, and then performs the following verifications: (i) Three commitment validations: $\hat{e}(B_1, g_2^{\gamma}) = \hat{e}(g_1, B_2^{\gamma})$, $\hat{e}(C_1, g_2^{\gamma}) = \hat{e}(g_1, C_2^{\gamma})$, $\hat{e}(D_1, g_2^{\gamma}) = \hat{e}(g_1, D_2^{\gamma})$. (ii) Three product argument verifications: $\hat{e}(B_1/g_1, B_2^{\gamma}) = \hat{e}(\pi_1, g_2^{\gamma Z(\chi)})$, $\hat{e}(S_1', B_2^{\gamma}) = \hat{e}(g_1, C_2^{\gamma}) \cdot \hat{e}(\pi_2, g_2^{\gamma Z(\chi)})$, $\hat{e}(g_1^{\ell_1(\chi)}, D_2^{\gamma}/(g_2^{\gamma \ell_1(\chi)})^s) = \hat{e}(\pi_4, g_2^{\gamma Z(\chi)})$. (iii) One shift argument verification, consisting of two equality tests: $\hat{e}(\pi_{31}, g_2^{\delta Z(\chi)}) = \hat{e}(g_1^{Z(\chi)}, \pi_{32}^{\delta})$, $\hat{e}(D_1/C_1\pi_{31}, g_2^{\delta Z(\chi)}) = \hat{e}(D_1, g_2^{\delta Z(\chi) Z^*(\chi)})$.

Theorem 4. *Π_{ssum} is perfectly complete and perfectly composable zero-knowledge. It is an $(\Theta(n)$-bounded-auxiliary-input) adaptive argument of knowledge if* BP *satisfies n-TSDH and the same assumptions as in Theorem 3 (for $z = 1$).*

The prover computation is dominated by three commitments and the application of 3 product SNARKs and 1 shift SNARK, i.e., by $\Theta(n \log n)$ non-cryptographic operations and $\Theta(n)$ cryptographic operations. The latter is dominated by nine ($\approx n$)-wide multi-exponentiations (2 in commitments to $\boldsymbol{c}$ and $\boldsymbol{d}$ and in the shift argument, and 1 in each product argument), 7 in $\mathbb{G}_1$ and 4 in $\mathbb{G}_2$. The argument size is constant (11 group elements), and the verifier computation is dominated by *offline* computation of two $(n+1)$-wide multi-exponentiations (needed to once commit to $\boldsymbol{S}$) and *online* computation of 17 pairings (3 pairings to verify π_2, 2 pairings to verify each of the other product arguments, 4 pairings to verify the shift argument, and 6 pairings to verify the validity of 3 commitments). In the full version [26], we will describe a batch-verification technique that allows to speed up on-line part of the verification of the Subset-Sum SNARK.

As always, multi-exponentiation can be sped up by using algorithms from [29, 32]; it can also be highly parallelized, potentially resulting in very fast parallel implementations of the zk-SNARK.

7 New Range SNARK

In a *range SNARK*, given public range $[L \,..\, H]$, the prover aims to convince the verifier that he knows how to open commitment (A_1, A_2^{γ}) to a value $a \in [L \,..\, H]$. That is, that the common input (A_1, A_2^{γ}) is a commitment to vector $\boldsymbol{a}$ with $a_1 = a$ and $a_i = 0$ for $i > 1$.

We first remark that instead of the range $[L \,..\, H]$, one can consider the range $[0 \,..\, H - L]$, and then use the homomorphic properties of the commitment scheme to add L to the committed value. Hence, we will just assume that the range is equal to $[0 \,..\, H]$ for some $H \geq 1$. Moreover, the efficiency of the following SNARK depends on the range length.

The new range SNARK Π_{rng} is very similar to Π_{ssum}, except that one has to additionally commit to a value $a \in [0\,..\,H]$, use a specific sparse $\boldsymbol{S}$ with $S_i = \lfloor (H + 2^{i-1})/2^i \rfloor$ [10,27], and prove that $a = \sum_{i=1}^{n} S_i b_i$ for the committed a. Since $\boldsymbol{S} = (S_i)_{i=1}^{n}$ does not depend on the instance (i.e., on a), the verifier computation is $\Theta(1)$. On the other hand, since the commitment to a is given as an input to the prover (and not created by prover as part of the argument), Π_{rng} has a more complex simulation strategy, with one more element in the trapdoor.

Let $n = \lfloor \log_2 H \rfloor + 1$. Define $S_i = \lfloor (H + 2^{i-1})/2^i \rfloor$ for $i \in [1\,..\,n]$ and $\boldsymbol{S} = (S_i)$. We again use the interpolating commitment scheme. To prove that $a \in [0\,..\,H]$, we do the following.

The CRS generation $\mathsf{G}_{\mathsf{rng}}$ invokes the CRS generations of the commitment scheme, the product SNARK and the shift SNARK, sharing the same gk and trapdoor $\mathsf{td} = (\chi, \delta/\gamma)$ between the different invocations. In this case, the trapdoor has to include δ/γ (which is well defined, since $\gamma \neq 0$) since the simulator does not know how to open (A_1, A_2^{γ}); see the proof of Theorem 5 for more details. We note that the trapdoor only has to contain δ/γ, and not γ and δ separately. The CRS also contains the first half of a commitment $S_1' \leftarrow \prod (g_1^{\ell_i(\chi)})^{S_i}$ to $\boldsymbol{S}$, needed for a later efficient verification of the argument π_2. Clearly, the CRS can be computed efficiently from $\mathsf{crs}_{\mathsf{rsft}}$ (for $z = 1$).

1 Let $a = \sum_{i=1}^{n} S_i b_i$ for $b_i \in \{0, 1\}$.
Let (B_1, B_2^{γ}) be a commitment to $\boldsymbol{b}$.
Construct a product argument π_1 to show that $\boldsymbol{b} = \boldsymbol{b} \circ \boldsymbol{b}$.
Let (C_1, C_2^{γ}) be a commitment to $\boldsymbol{c} \leftarrow \boldsymbol{S} \circ \boldsymbol{b}$.
Construct a product argument π_2 to show that $\boldsymbol{c} = \boldsymbol{S} \circ \boldsymbol{b}$.
Let (D_1, D_2^{γ}) be a commitment to $\boldsymbol{d}$, where $d_i = \sum_{j \geq i} c_i$.
Construct a shift argument $(\pi_{31}, \pi_{32}^{\delta})$ to show that $\boldsymbol{d} = (\boldsymbol{d} - \boldsymbol{c}) \gg 1$.
2 Construct a product argument π_4 to show that $\boldsymbol{e}_1 \circ (\boldsymbol{d} - \boldsymbol{a}) = \boldsymbol{0}_n$.
Output $\pi_{\mathsf{rng}} = (B_1, B_2^{\gamma}, C_1, C_2^{\gamma}, D_1, D_2^{\gamma}, \pi_1, \pi_2, \pi_{31}, \pi_{32}^{\delta}, \pi_4)$.

Fig. 2. The new range argument Π_{rng}

The prover's actions on input (A_1, A_2^{γ}) are depicted by Fig. 2 (further explanations are given in the concise completeness proof in Theorem 5). The only differences, compared to the prover computation of Π_{ssum}, are the computation of b_i on step 1, and of π_4 on step 2. After receiving π_{rng}, the verifier performs the following checks: (i) Four commitment validations: $\hat{e}(A_1, g_2^{\gamma}) = \hat{e}(g_1, A_2^{\gamma})$, $\hat{e}(B_1, g_2^{\gamma}) = \hat{e}(g_1, B_2^{\gamma})$, $\hat{e}(C_1, g_2^{\gamma}) = \hat{e}(g_1, C_2^{\gamma})$, $\hat{e}(D_1, g_2^{\gamma}) = \hat{e}(g_1, D_2^{\gamma})$. (ii) Three product argument verifications: $\hat{e}(B_1/g_1, B_2^{\gamma}) = \hat{e}(\pi_1, g_2^{\gamma Z(\chi)})$, $\hat{e}(S_1', B_2^{\gamma}) = \hat{e}(g_1, C_2^{\gamma}) \cdot \hat{e}(\pi_2, g_2^{\gamma Z(\chi)})$, $\hat{e}(g_1^{\ell_1(\chi)}, D_2^{\gamma}/A_2^{\gamma}) = \hat{e}(\pi_4, g_2^{\gamma Z(\chi)})$. (iii) One shift argument verification, consisting of two equality tests: $\hat{e}(\pi_{31}, g_2^{\delta Z(\chi)}) = \hat{e}(g_1^{Z(\chi)}, \pi_{32}^{\delta})$, $\hat{e}(D_1/C_1\pi_{31}, g_2^{\delta Z(\chi)}) = \hat{e}(D_1, g_2^{\delta Z(\chi) Z^*(\chi)})$.

Theorem 5. *Π_{rng} is perfectly complete and composable zero-knowledge. If* BP *satisfies n-TSDH and the assumptions of Theorem 3, then Π_{rng} is an adaptive ($\Theta(n)$-bounded-auxiliary-input) argument of knowledge.*

The prover computation is dominated by three commitments and the application of three product arguments and one shift argument, that is, by $\Theta(n \log n)$ non-cryptographic operations and $\Theta(n)$ cryptographic operations. The latter is dominated by nine ($\approx n$)-wide multi-exponentiations (2 in commitments to $\boldsymbol{c}$ and $\boldsymbol{d}$ and in the shift argument, and 1 in each product argument), seven in $\mathbb{G}_1$ and four in $\mathbb{G}_2$. The argument size is constant (11 group elements), and the verifier computation is dominated by 19 pairings (3 pairings to verify π_2, 2 pairings to verify each of the other product arguments, 4 pairings to verify the shift argument, and 8 pairings to verify the validity of 4 commitments). In this case, since the verifier does not have to commit to $\boldsymbol{S}$, the verifier computation is dominated by $\Theta(1)$ cryptographic operations.

The new range SNARK is significantly more computation-efficient for the prover than the previous range SNARKs [11,15] that have prover computation $\Theta(r_3^{-1}(n) \log n)$. Π_{rng} has better communication (11 versus 31 group elements in [15]), and verification complexity (19 versus 65 pairings in [15]). Moreover, Π_{rng} is also simpler: since the prover computation is quasi-linear, we do not have to consider various trade-offs (though they are still available) between computation and communication as in [11,15]. In the full version [26], we will use batch verification to further speed up the verification of the Range SNARK.

Acknowledgments. We would like to thank Diego Aranha, Paulo Barreto, Markulf Kohlweiss, and Prastudy Fauzi for useful comments. This work was supported by the European Union's Horizon 2020 research and innovation programme under grant agreement No. 653497 (project PANORAMIX), and the Estonian Research Council.

References

1. Barreto, P.S.L.M., Naehrig, M.: Pairing-friendly elliptic curves of prime order. In: Preneel, B., Tavares, S. (eds.) SAC 2005. LNCS, vol. 3897, pp. 319–331. Springer, Heidelberg (2006)
2. Bellare, M., Garay, J.A., Rabin, T.: Batch verification with applications to cryptography and checking. In: Lucchesi, C.L., Moura, A.V. (eds.) LATIN 1998. LNCS, vol. 1380, pp. 170–191. Springer, Heidelberg (1998)
3. Ben-Sasson, E., Chiesa, A., Genkin, D., Tromer, E., Virza, M.: SNARKs for C: verifying program executions succinctly and in zero knowledge. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013, Part II. LNCS, vol. 8043, pp. 90–108. Springer, Heidelberg (2013)
4. Ben-Sasson, E., Chiesa, A., Tromer, E., Virza, M.: Scalable zero knowledge via cycles of elliptic curves. In: Garay, J.A., Gennaro, R. (eds.) CRYPTO 2014, Part II. LNCS, vol. 8617, pp. 276–294. Springer, Heidelberg (2014)
5. Ben-Sasson, E., Chiesa, A., Tromer, E., Virza, M.: Succinct non-interactive zero knowledge for a von Neumann architecture. In: USENIX, pp. 781–796 (2014)

6. Bitansky, N., Chiesa, A., Ishai, Y., Ostrovsky, R., Paneth, O.: Succinct non-interactive arguments via linear interactive proofs. In: Sahai, A. (ed.) TCC 2013. LNCS, vol. 7785, pp. 315–333. Springer, Heidelberg (2013)
7. Boneh, D., Boyen, X.: Short signatures without random oracles and the SDH assumption in bilinear groups. J. Cryptol. **21**(2), 149–177 (2008)
8. Bos, J.W., Costello, C., Naehrig, M.: Exponentiating in pairing groups. In: Lange, T., Lauter, K., Lisoněk, P. (eds.) SAC 2013. LNCS, vol. 8282, pp. 438–455. Springer, Heidelberg (2014)
9. Canetti, R., Lindell, Y., Ostrovsky, R., Sahai, A.: Universally composable two-party and multi-party secure computation. In: STOC, pp. 494–503 (2002)
10. Chaabouni, R., Lipmaa, H., Shelat, A.: Additive combinatorics and discrete logarithm based range protocols. In: Steinfeld, R., Hawkes, P. (eds.) ACISP 2010. LNCS, vol. 6168, pp. 336–351. Springer, Heidelberg (2010)
11. Chaabouni, R., Lipmaa, H., Zhang, B.: A non-interactive range proof with constant communication. In: Keromytis, A.D. (ed.) FC 2012. LNCS, vol. 7397, pp. 179–199. Springer, Heidelberg (2012)
12. Costello, C., Fournet, C., Howell, J., Kohlweiss, M., Kreuter, B., Naehrig, M., Parno, B., Zahur, S.: Geppetto: versatile verifiable computation. In: IEEE SP, pp. 253–270 (2015)
13. 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. LNCS, vol. 8873, pp. 532–550. Springer, Heidelberg (2014)
14. Fauzi, P., Lipmaa, H.: Efficient culpably sound NIZK shuffle argument without random oracles. CT-RSA 2016. LNCS, vol. 9610. Springer, switzerland (2016)
15. Fauzi, P., Lipmaa, H., Zhang, B.: Efficient modular NIZK arguments from shift and product. In: Abdalla, M., Nita-Rotaru, C., Dahab, R. (eds.) CANS 2013. LNCS, vol. 8257, pp. 92–121. Springer, Heidelberg (2013)
16. Garey, M.R., Johnson, D.S.: Computers and Intractability: A Guide to the Theory of NP-Completeness. Series of Books in the Mathematical Sciences. W.H. Freeman, New York (1979)
17. 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 (2013)
18. Gentry, C., Wichs, D.: Separating succinct non-interactive arguments from all falsifiable assumptions. In: STOC, pp. 99–108 (2011)
19. Groth, J.: Short pairing-based non-interactive zero-knowledge arguments. In: Abe, M. (ed.) ASIACRYPT 2010. LNCS, vol. 6477, pp. 321–340. Springer, Heidelberg (2010)
20. Kilian, J.: Uses of randomness in algorithms and protocols. Ph.D. thesis, Massachusetts Institute of Technology, USA (1989)
21. Kolesnikov, V., Schneider, T.: A practical universal circuit construction and secure evaluation of private functions. In: Tsudik, G. (ed.) FC 2008. LNCS, vol. 5143, pp. 83–97. Springer, Heidelberg (2008)
22. Lipmaa, H.: On diophantine complexity and statistical zero-knowledge arguments. In: Laih, C.-S. (ed.) ASIACRYPT 2003. LNCS, vol. 2894, pp. 398–415. Springer, Heidelberg (2003)
23. Lipmaa, H.: Progression-free sets and sublinear pairing-based non-interactive zero-knowledge arguments. In: Cramer, R. (ed.) TCC 2012. LNCS, vol. 7194, pp. 169–189. Springer, Heidelberg (2012)

24. Lipmaa, H.: Succinct non-interactive zero knowledge arguments from span programs and linear error-correcting codes. In: Sako, K., Sarkar, P. (eds.) ASIACRYPT 2013, Part I. LNCS, vol. 8269, pp. 41–60. Springer, Heidelberg (2013)
25. Lipmaa, H.: Efficient NIZK arguments via parallel verification of benes networks. In: Abdalla, M., De Prisco, R. (eds.) SCN 2014. LNCS, vol. 8642, pp. 416–434. Springer, Heidelberg (2014)
26. Lipmaa, H.: Prover-efficient commit-and-prove zero-knowledge SNARKs. TR 2014/396, IACR (2014). http://eprint.iacr.org/2014/396
27. Lipmaa, H., Asokan, N., Niemi, V.: Secure vickrey auctions without threshold trust. FC 2002. LNCS, vol. 2357, pp. 87–101. Springer, Heidelberg (2002)
28. Parno, B., Gentry, C., Howell, J., Raykova, M.: Pinocchio: nearly practical verifiable computation. In: IEEE SP, pp. 238–252 (2013)
29. Pippenger, N.: On the evaluation of powers and monomials. SIAM J. Comput. **9**(2), 230–250 (1980)
30. Raz, R.: Elusive functions and lower bounds for arithmetic circuits. Theor. Comput. **6**(1), 135–177 (2010)
31. Sadeghi, A.-R., Schneider, T.: Generalized universal circuits for secure evaluation of private functions with application to data classification. In: Lee, P.J., Cheon, J.H. (eds.) ICISC 2008. LNCS, vol. 5461, pp. 336–353. Springer, Heidelberg (2009)
32. Straus, E.G.: Addition chains of vectors. Amer. Math. Mon. **70**, 806–808 (1964)
33. Valiant, L.G.: Universal circuits (Preliminary report). In: STOC, pp. 196–203 (1976)