

Pinocchio-Based Adaptive zk-SNARKs and Secure/Correct Adaptive Function Evaluation

Meilof Veeningen, meilof.veeningen@philips.com

Philips Research, Eindhoven, The Netherlands

Abstract. Pinocchio is a practical zk-SNARK that allows a prover to perform cryptographically verifiable computations with verification effort potentially less than performing the computation itself. A recent proposal showed how to make Pinocchio adaptive (or “hash-and-prove”), i.e., to enable proofs with respect to computation-independent commitments. This enables computations to be chosen after the commitments have been produced, and for data to be shared between different computations in a flexible way. Unfortunately, this proposal is not zero-knowledge. In particular, it cannot be combined with Trinocchio, a system in which Pinocchio is outsourced to three workers that do not learn the inputs thanks to multi-party computation (MPC). In this paper, we show how to make Pinocchio adaptive in a zero-knowledge way; apply this to make Trinocchio work on computation-independent commitments; present tooling to easily program flexible verifiable computations (with or without MPC); and use it to build a prototype in a medical research case study.

Full version of [Vee17] including appendices

1 Introduction

Recent advances in SNARKs (Succinct Non-interactive ARguments of Knowledge) are making it more and more feasible to outsource computations to the cloud while obtaining cryptographic guarantees about the correctness of their outputs. In particular, the Pinocchio system [GGPR13,PHGR13] achieved for the first time for a practical computation a verification time of a computation proof that was actually faster than performing the computation itself.

In Pinocchio, proofs are verified with respect to plaintext inputs and outputs of the verifier; but in many cases, it is useful to have computation proofs that also refer to committed data, e.g., provided by a third party. Ideally, such proofs should be *adaptive*, i.e., multiple different computations can be performed on the same commitment, that are chosen after the data has been committed to; and *zero-knowledge*, i.e., the commitments and proofs should reveal no information about the committed data. This latter property allows proofs on sensitive data, and it allows extensions like Trinocchio [SVdV16] that additionally hide this sensitive data from provers by multi-party computation.

Although several approaches are known from the literature, no really satisfactory practical adaptive zk-SNARK exists. The recent “hash first” proposal

[FFG⁺16] shows how to make Pinocchio adaptive at low overhead, but is unfortunately not zero-knowledge. On the other hand, Pinocchio’s successor Gепpetto [CFH⁺15] is zero-knowledge but not adaptive: multiple computations can be performed on the same data but they need to be known before committing. The asymptotically best known SNARKS combining the two properties have $\Theta(n \log n)$ non-cryptographic and $\Theta(n)$ cryptographic work for the prover, a $\Theta(n)$ -sized CRS, and constant-time verification (where n is the size of the computation), but with a large practical overhead: [Lip16] because it relies on the impractical subset-sum language; other constructions (e.g., [CFH⁺15,FFG⁺16]) because they rely on including hash evaluation in the computation¹. Finally, [BBFR15] enables Pinocchio proofs on authenticated data with prover complexity as above, but verification time is linear in the number of committed inputs.

In this work, we give a new Pinocchio-based adaptive zk-SNARK that solves the above problems. We match the best asymptotic performance (i.e., $\Theta(n \log n)$ non-cryptographic work and $\Theta(n)$ cryptographic work for the prover; a $\Theta(n)$ -size CRS and constant-time verification); but obtain the first practical solution by adding only minor overhead to “plain” Pinocchio (instead of relying on expensive approaches such as subset-sum or bootstrapping).

As additional contributions, we apply our zk-SNARK in the Trinocchio setting, and present tooling to easily perform verifiable computations. Trinocchio [SVdV16] achieves privacy-preserving outsourcing to untrusted workers by combining the privacy guarantees of multi-party computation with the correctness guarantees of the Pinocchio zk-SNARK. With our adaptive zk-SNARK, computation can be chosen *after* the inputs were provided and more complex functionalities can be achieved by using the output of one computation as input of another. We also improve the generality of [SVdV16] by proving security for *any* suitable MPC protocol and adaptive zk-SNARK. Our tooling consists of a Python frontend and a C++ backend. The frontend allows easy programming of verifiable computations (with libraries for zero testing, oblivious indexing and fixed-point computations), and execution either directly (for normal outsourcing scenarios) or with MPC (for privacy-preserving outsourcing). The backend provides key generation, proving, and verification functionality for both scenarios.

2 Preliminaries

2.1 Algebraic tools, notation, and complexity assumptions

Our constructions aim to prove correctness of computations over a prime order field $\mathbb{F} = \mathbb{F}_p$. We make use of pairings, i.e., groups $(\mathbb{G}, \mathbb{G}', \mathbb{G}_T)$ of order p and an efficient bilinear map $e : \mathbb{G} \times \mathbb{G}' \rightarrow \mathbb{G}_T$, where for any generators $g_1 \in \mathbb{G}, g_2 \in \mathbb{G}'$, $e(g_1, g_2) \neq 1$ and $e(g_1^a, g_2^b) = e(g_1, g_2)^{ab}$. Throughout, we will make use of polynomial evaluations in a secret point $s \in \mathbb{F}$. For $f \in \mathbb{F}[x]$, write $\langle f \rangle_1$ for $f(s) \cdot g_1$ and $\langle f \rangle_2$ for $f(s) \cdot g_2$.

¹ In practice, computing the hash is complex itself. It can be avoided with bootstrapping [Lip16], giving slightly worse asymptotics and again a large practical overhead.

Let $(\mathbb{G}, \mathbb{G}', \mathbb{G}_T, e) \leftarrow \mathcal{G}(1^\kappa)$ denote parameter generation in this setting. We require the following assumptions from [DFGK14] (that generalise those from [PHGR13] to asymmetric pairings):

Definition 1. *The q -power Diffie Hellman (q -PDH) assumption holds for \mathcal{G} if, for any NUPPT adversary \mathcal{A} : $\Pr[(\mathbb{G}, \mathbb{G}', \mathbb{G}_T, e) \leftarrow \mathcal{G}(1^\kappa); g \in_R \mathbb{G}^*; g' \in_R \mathbb{G}'^*; s \in_R \mathbb{Z}_p^*; y \leftarrow \mathcal{A}(\mathbb{G}, \mathbb{G}', \mathbb{G}_T, e, \{g^{s^i}, g'^{s^i}\}_{i=1, \dots, q, q+2, \dots, 2q}) : y = g^{s^{q+1}}] \approx_\kappa 0$.*

Definition 2. *The q -power knowledge of exponent (q -PKE) assumption holds for \mathcal{G} and a class \mathcal{Z} of auxiliary input generators if, for every NUPPT auxiliary input generator $Z \in \mathcal{Z}$ and any NUPPT adversary \mathcal{A} there exists a NUPPT extractor $\mathcal{E}_\mathcal{A}$ such that: $\Pr[\text{crs} := (\mathbb{G}, \mathbb{G}', \mathbb{G}_T, e) \leftarrow \mathcal{G}(1^\kappa); g \in_R \mathbb{G}^*; s \in_R \mathbb{Z}_p^*; z \leftarrow Z(\text{crs}, g, \dots, g^{s^q}); g' \in_R \mathbb{G}'^*; (c, c' || a_0, \dots, a_q) \leftarrow (\mathcal{A} || \mathcal{E}_\mathcal{A})(\text{crs}, \{g^{s^i}, g'^{s^i}\}_{i=0, \dots, q}, z) : e(c, g') = e(g, c') \wedge c \neq \prod_{i=0}^q g_1^{a_i s^i}] \approx_\kappa 0$.*

Here, $(a||b) \leftarrow (\mathcal{A} || \mathcal{E}_\mathcal{A})(c)$ denotes running both algorithms on the same inputs and random tape, and assigning their results to a respectively b . For certain auxiliary input generators, the q -PKE assumption does not hold, so we have to conjecture that our auxiliary input generators are “benign”, cf. [DFGK14].

Definition 3. *The q -target group strong Diffie Hellman (q -SDH) assumption holds for \mathcal{G} if, for any NUPPT adversary \mathcal{A} , $\Pr[\text{crs} := (\mathbb{G}, \mathbb{G}', \mathbb{G}_T, e) \leftarrow \mathcal{G}(1^\kappa); g \in_R \mathbb{G}^*; g' \in_R \mathbb{G}'^*; s \in_R \mathbb{Z}_p(r, Y) \leftarrow \mathcal{A}(\text{crs}, \{g^{s^i}, g'^{s^i}\}_{i=0, \dots, q}) : r \in \mathbb{Z}_p \setminus \{s\} \wedge Y = e(g, g')^{\frac{1}{s-r}}] \approx_\kappa 0$.*

2.2 Adaptive zk-SNARKs in the CRS model

We now define adaptive zk-SNARKs as in [Lip16] with minor modifications. We first define extractable trapdoor commitment families. This is a straightforward generalisation of an extractable trapdoor commitment scheme [Lip16] that explicitly captures multiple commitment keys generated from the same CRS:

Definition 4. *Let $(\text{G0}, \text{Gc}, \text{C})$ be a scheme where $(\text{crs}, \text{td}) \leftarrow \text{G0}(1^\kappa)$ outputs a system-wide CRS and a trapdoor; $(\text{ck}, \text{ctd}) \leftarrow \text{Gc}(\text{crs})$ outputs a commitment key and a trapdoor; and $c \leftarrow \text{C}_{\text{ck}}(m; r)$ outputs a commitment with the given key. Such a scheme is called an extractable trapdoor commitment family if:*

- (Computationally binding) For every NUPPT \mathcal{A} , $\Pr[(\text{crs}, \cdot) \leftarrow \text{G0}(1^\kappa); (\text{ck}, \cdot) \leftarrow \text{Gc}(\text{crs}); (v; r; v'; r') \leftarrow \mathcal{A}(\text{crs}; \text{ck}) : \text{C}_{\text{ck}}(v; r) = \text{C}_{\text{ck}}(v'; r')] \approx 0$.
- (Perfectly hiding) Letting $(\text{crs}, \cdot) \leftarrow \text{G0}(1^\kappa); (\text{ck}, \cdot) \leftarrow \text{Gc}(\text{crs})$, for all $v, v', \text{C}_{\text{ck}}(v; r)$ and $\text{C}_{\text{ck}}(v'; r')$ are identically distributed given random r, r'
- (Trapdoor) There exists a NUPPT algorithm \mathcal{T} such that if $(\text{crs}, \text{td}) \leftarrow \text{G0}(1^\kappa); (\text{ck}, \text{ctd}) \leftarrow \text{Gc}(\text{crs}); (u; t) \leftarrow \mathcal{T}(\text{crs}; \text{td}; \text{ck}; \text{ck}); r \leftarrow \mathcal{T}(t; u; v)$, then u is distributed identically to real commitments and $\text{C}_{\text{ck}}(v; r) = u$.
- (Extractable) For every NUPPT committer \mathcal{A} , there exists a NUPPT extractor $\mathcal{E}_\mathcal{A}$ such that $\Pr[(\text{crs}, \cdot) \leftarrow \text{G0}(1^\kappa); (\text{ck}, \cdot) \leftarrow \text{Gc}(\text{crs}); (u||v; r) \leftarrow (\mathcal{A} || \mathcal{E}_\mathcal{A})(\text{crs}; \text{ck}) : u \in \text{Range}(\text{C}_{\text{ck}}) \wedge u \neq \text{C}_{\text{ck}}(v; r)] \approx 0$.

Given relation \mathcal{R} and commitment keys ck_1, \dots, ck_n from the same commitment family, define: $\mathcal{R}_{ck_1, \dots, ck_n} := \{(\mathbf{u}; \mathbf{v}, \mathbf{r}, \mathbf{w}) : \mathbf{u}_i = C_{ck_i}(\mathbf{v}_i; \mathbf{r}_i) \wedge (\mathbf{v}, \mathbf{w}) \in \mathcal{R}\}$. Intuitively, an adaptive zk-SNARK is a zk-SNARK for relation $\mathcal{R}_{ck_1, \dots, ck_n}$.

Definition 5. *An adaptive zk-SNARK for extractable trapdoor commitment family (G_0, G_c, C) and relation \mathcal{R} is a scheme (G, P, V) where²:*

- $(\text{crsp}; \text{crsv}; \text{tdp}) \leftarrow G(\text{crs}; \{\text{ck}_i\})$, given a CRS and commitment keys, outputs evaluation and verification keys, and a trapdoor;
- $\pi \leftarrow P(\text{crs}; \{\text{ck}_i\}; \text{crsp}; \mathbf{u}; \mathbf{v}; \mathbf{r}; \mathbf{w})$, given a CRS; commitment keys; an evaluation key; commitments; openings; and a witness, outputs a proof;
- $0/1 \leftarrow V(\text{crs}; \{\text{ck}_i\}; \text{crsv}; \mathbf{u}; \pi)$, given a CRS; commitment keys; a verification key; commitments; and a proof, verifies the proof,

satisfying the following properties (let $\text{setup} := (\text{crs}; \text{td}) \leftarrow G_0(1^\kappa); \forall i : (\text{ck}_i; \text{ctd}_i) \leftarrow G_c(\text{crs}); (\text{crsp}; \text{crsv}; \text{tdp}) \leftarrow G(\text{crs}; \{\text{ck}_i\})$):

- Perfect completeness (“proofs normally verify”) $\Pr[\text{setup}; (\mathbf{u}; \mathbf{v}; \mathbf{r}; \mathbf{w}) \leftarrow \mathcal{R}_{\{\text{ck}_i\}} : V(\text{crs}; \{\text{ck}_i\}; \text{crsv}; \mathbf{u}; \pi) = 1] = 1$.
- Argument of knowledge (“the commitment openings and a valid witness can be extracted from an adversary producing a proof”): for every NUPPT \mathcal{A} there exists NUPPT extractor $\mathcal{E}_{\mathcal{A}}$ such that, for every auxiliary information $\text{aux} \in \{0, 1\}^{\text{poly}(\kappa)} : \Pr[\text{setup}; (\mathbf{u}; \pi || \mathbf{v}; \mathbf{r}; \mathbf{w}) \leftarrow (\mathcal{A} || \mathcal{E}_{\mathcal{A}})(\text{crs}; \{\text{ck}_i\}; \text{crsp}; \text{aux} || \dots; \text{td}; \text{ctd}_1; \dots; \text{ctd}_n; \text{tdp}) : (\mathbf{u}; \mathbf{v}; \mathbf{r}; \mathbf{w}) \notin \mathcal{R}_{\{\text{ck}_i\}} \wedge V(\text{crs}; \{\text{ck}_i\}; \text{crsv}; \mathbf{u}; \pi) = 1] \approx_\kappa 0$. Here, $(\mathcal{A} || \mathcal{E}_{\mathcal{A}})(\cdot || \dots; \cdot')$ is parallel execution with extra input \cdot' for $\mathcal{E}_{\mathcal{A}}$.
- Perfectly composable zero knowledge (“proofs can be simulated using the commitments and trapdoor”): there exists a PPT simulator \mathcal{S} such that, for all stateful NUPPT adversaries \mathcal{A} , $\Pr[\text{setup}; (\mathbf{u}; \mathbf{v}; \mathbf{r}; \mathbf{w}) \leftarrow \mathcal{A}(\text{crs}, \{\text{ck}_i\}, \text{crsp}); \pi \leftarrow P(\text{crs}, \{\text{ck}_i\}, \text{crsp}; \mathbf{u}; \mathbf{v}; \mathbf{r}; \mathbf{w}) : (\mathbf{u}, \mathbf{v}, \mathbf{r}, \mathbf{w}) \in \mathcal{R}_{\{\text{ck}_i\}} \wedge \mathcal{A}(\pi) = 1] = \Pr[\text{setup}; (\mathbf{u}; \mathbf{v}; \mathbf{r}; \mathbf{w}) \leftarrow \mathcal{A}(\text{crs}, \{\text{ck}_i\}, \text{crsp}); \pi \leftarrow \mathcal{S}(\text{crs}, \{\text{ck}_i\}, \text{crsp}; \mathbf{u}; \text{td}, \{\text{ctd}_i\}, \text{tdp}) : (\mathbf{u}, \mathbf{v}, \mathbf{r}, \mathbf{w}) \in \mathcal{R}_{\{\text{ck}_i\}} \wedge \mathcal{A}(\pi) = 1]$.

We base our definitions on [Lip16] because it is closest to what we want to achieve. Unlike in [CFH⁺15], we do not guarantee security when relation \mathcal{R} is chosen adaptively based on the commitment keys; this is left as future work.

2.3 The Pinocchio zk-SNARK construction from [PHGR13]

QAPs Pinocchio models computations as *quadratic arithmetic programs* (QAPs) [GGPR13]. A QAP over a field \mathbb{F} is a triple $(\mathbf{V}, \mathbf{W}, \mathbf{Y}) \in (\mathbb{F}^{d \times k})^3$, where d is called the *degree* of the QAP and k is called the *size*. A vector $\mathbf{x} \in \mathbb{F}^k$ is said to be a *solution* to the QAP if $(\mathbf{V} \cdot \mathbf{x}) \times (\mathbf{W} \cdot \mathbf{x}) = \mathbf{Y} \cdot \mathbf{x}$, where \times denotes the pairwise product and \cdot denotes normal matrix-vector multiplication. A QAP Q

² We differ from [Lip16] in three minor ways: 1) we generalise from commitment schemes to families because we need this in Adaptive Trinocchio; 2) we allow witnesses that are not committed to separately, giving a slight efficiency improvement; 3) the extractor has access to the trapdoor, as needed when using Pinocchio [GGPR13].

is said to *compute* function $f : \mathbb{F}^i \rightarrow \mathbb{F}^j$ if $\mathbf{b} = f(\mathbf{a})$ if and only if there exists a *witness* \mathbf{w} such that $(\mathbf{a}; \mathbf{b}; \mathbf{w})$ is a solution to Q . For example, consider the QAP

$$\mathbf{V} = \begin{pmatrix} 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 \end{pmatrix}, \quad \mathbf{W} = \begin{pmatrix} 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \quad \mathbf{Y} = \begin{pmatrix} 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix}.$$

Intuitively, the first row of this QAP represents equation $(x_1 + x_2) \cdot (x_1 + x_2) = x_4$ in variables (x_1, x_2, x_3, x_4) whereas the second row represents equation $(x_1 + x_2) \cdot x_4 = x_3$. Note that $x_3 = (x_1 + x_2)^3$ if and only if there exists x_4 satisfying the two equations, so this QAP computes function $f : (x_1, x_2) \mapsto x_3$.³

Fixing d distinct, public $\omega_1, \dots, \omega_d \in \mathbb{F}$, then a QAP can equivalently be described by a collection of interpolating polynomials in these points. Namely, let $\{v_i(x)\}$ be the unique polynomials of degree $< d$ such that $v_i(\omega_j) = \mathbf{V}_{j,i}$, and similarly for $\{w_i(x)\}$, $\{y_i(x)\}$. Then $\{v_i(x), w_i(x), y_i(x)\}$ is an equivalent description of the QAP. Defining $t(x) = (x - \omega_1) \cdot \dots \cdot (x - \omega_d) \in \mathbb{F}[x]$, note that $\mathbf{x}_1, \dots, \mathbf{x}_n$ is a solution to Q if and only if, for all j , $(\sum_i \mathbf{x}_i \cdot v_i(\omega_j)) \cdot (\sum_i \mathbf{x}_i \cdot w_i(\omega_j)) = (\sum_i \mathbf{x}_i \cdot y_i(\omega_j))$, or equivalently, if $t(x)$ divides $p(x) := (\sum_i \mathbf{x}_i \cdot v_i(x)) \cdot (\sum_i \mathbf{x}_i \cdot w_i(x)) - (\sum_i \mathbf{x}_i \cdot y_i(x)) \in \mathbb{F}[x]$.

Security Guarantees Pinocchio is a zk-SNARK, which is essentially the same as an adaptive zk-SNARK except proving and verifying are with respect to plaintext values instead of commitments. In Pinocchio, relation \mathcal{R} is that, for given \mathbf{v} , there exists witness \mathbf{w} such that $(\mathbf{v}; \mathbf{w})$ is a solution to a given QAP Q . Pinocchio replies on a pairing secure under the $(4d + 4)$ -PDH, d -PKE and $(8d + 8)$ -SDH assumptions discussed above, where d is the degree of the QAP.

Construction Fix random, secret $s, \alpha_v, \alpha_w, \alpha_y, \beta, r_v, r_w, r_y(x) := r_v r_w$. The central idea of Pinocchio is to prove satisfaction of all QAP equations using evaluations of the interpolating polynomials in a secret point. Namely, the prover computes quotient polynomial $h = p/t$ and basically provides evaluations “in the exponent” of h , $\sum_i \mathbf{x}_i \cdot v_i$, $\sum_i \mathbf{x}_i \cdot w_i$, $\sum_i \mathbf{x}_i \cdot y_i$ in the point s that is unknown to him, that can then be verified using the pairing. Precisely, the prover algorithm, given solution $\mathbf{x} = (\mathbf{v}; \mathbf{w})$ to the QAP, generates random $\delta_v, \delta_w, \delta_y$; computes coefficients \mathbf{h} of the polynomial $(\sum_i \mathbf{x}_i \cdot v_i(x) + \delta_v t(x)) \cdot (\sum_i \mathbf{x}_i \cdot w_i(x) + \delta_w t(x)) - (\sum_i \mathbf{x}_i \cdot y_i(x) + \delta_y t(x)) / t(x)$ (with δ . terms added to make the proof zero-knowledge), and outputs (all \sum_i over witness indices $|\mathbf{v}| + 1, \dots, |\mathbf{x}|$; recall

³ In Pinocchio, the linear terms corresponding to \mathbf{V} , \mathbf{W} , \mathbf{Y} can also contain constant values. This is achieved by assigning special meaning to a “constant” wire with value 1. We do not formalise this separately, instead leaving it up to the user to include a special variable and an equation $x_i \cdot x_i = x_i$ that forces this variable to be one.

that for polynomial f , $\langle f \rangle_1 := f(s) \cdot g_1$ and $\langle f \rangle_2 := f(s) \cdot g_2$:

$$\begin{aligned}\langle V \rangle_1 &= \sum_i \mathbf{x}_i \langle r_v v_i \rangle_1 + \delta_v \langle r_v t \rangle_1, \langle \alpha_v V \rangle_2 = \sum_i \mathbf{x}_i \langle \alpha_v r_v v_i \rangle_2 + \delta_v \langle \alpha_v r_v t \rangle_2, \\ \langle W \rangle_2 &= \sum_i \mathbf{x}_i \langle r_w w_i \rangle_2 + \delta_w \langle r_w t \rangle_2, \langle \alpha_w W \rangle_1 = \sum_i \mathbf{x}_i \langle \alpha_w r_w w_i \rangle_1 + \delta_w \langle \alpha_w r_w t \rangle_1, \\ \langle Y \rangle_1 &= \sum_i \mathbf{x}_i \langle r_y y_i \rangle_1 + \delta_y \langle r_y t \rangle_1, \langle \alpha_y Y \rangle_2 = \sum_i \mathbf{x}_i \langle \alpha_y r_y y_i \rangle_2 + \delta_y \langle \alpha_y r_y t \rangle_2. \\ \langle Z \rangle_1 &= \sum_i \mathbf{x}_i \langle r_v \beta v_i + r_w \beta w_i + r_y \beta y_i \rangle_1 + \delta_v \langle r_v \beta t \rangle_1 + \delta_w \langle r_w \beta t \rangle_1 + \delta_y \langle r_y \beta t \rangle_1, \\ \langle H \rangle_1 &= \sum_{j=0}^d \mathbf{h}_j \langle x^j \rangle_1.\end{aligned}$$

The evaluation key consists of all $\langle \cdot \rangle_1$, $\langle \cdot \rangle_2$ items used in the formulas above.⁴

The verification algorithm, given statement \mathbf{v} , extends $\langle V \rangle_1$, $\langle W \rangle_1$, $\langle Y \rangle_1$ to include also the input/output wires (\sum_i over I/O wire indices $1, \dots, |\mathbf{v}|$): $\langle V^+ \rangle_1 = \langle V \rangle_1 + \sum_i \mathbf{x}_i \langle r_v v_i \rangle_1$, $\langle W^+ \rangle_2 = \langle W \rangle_2 + \sum_i \mathbf{x}_i \langle r_w w_i \rangle_2$, $\langle Y^+ \rangle_1 = \langle Y \rangle_1 + \sum_i \mathbf{x}_i \langle r_y y_i \rangle_1$. Then, it checks (the verification key are the needed $\langle \cdot \rangle_1$, $\langle \cdot \rangle_2$ items):

$$\begin{aligned}e(\langle V \rangle_1, \langle \alpha_v \rangle_2) &= e(\langle 1 \rangle_1, \langle \alpha_v V \rangle_2); & \mathbf{(V)} \\ e(\langle \alpha_w \rangle_1, \langle W \rangle_2) &= e(\langle \alpha_w W \rangle_1, \langle 1 \rangle_2); & \mathbf{(W)} \\ e(\langle Y \rangle_1, \langle \alpha_y \rangle_2) &= e(\langle 1 \rangle_1, \langle \alpha_y Y \rangle_2); & \mathbf{(Y)} \\ e(\langle V \rangle_1 + \langle Y \rangle_1, \langle \beta \rangle_2) \cdot e(\langle \beta \rangle_1, \langle W \rangle_2) &= e(\langle Z \rangle_1, \langle 1 \rangle_2); & \mathbf{(Z)} \\ e(\langle V^+ \rangle_1, \langle W^+ \rangle_2) \cdot e(\langle Y^+ \rangle_1, \langle 1 \rangle_2)^{-1} &= e(\langle H \rangle_1, \langle r_y t \rangle_2). & \mathbf{(H)}\end{aligned}$$

At a high level, checks **(V)**, **(W)**, **(Y)** guarantee that the proof is a proof of knowledge of the witness \mathbf{w} ; check **(Z)** guarantees that the same witness \mathbf{w} was used for $\langle V \rangle_1, \langle W \rangle_2, \langle Y \rangle_1$; and check **(H)** guarantees that indeed, $p(x) = h(x) \cdot t(x)$ holds, which implies a solution to the QAP.

3 Adaptive zk-SNARKs based on Pinocchio

This section presents the central contribution of this paper: an adaptive zk-SNARK based on Pinocchio. We obtain our Pinocchio-based adaptive zk-SNARK by generalising the role of the $\langle Z \rangle_1$ element of the Pinocchio proof. Recall that in Pinocchio, proof elements $\langle V \rangle_1$, $\langle W \rangle_1$, and $\langle Y \rangle_1$ are essentially weighted sums $\sum_j \mathbf{x}_j \langle v_j \rangle_1$, $\sum_j \mathbf{x}_j \langle w_j \rangle_2$, $\sum_j \mathbf{x}_j \langle y_j \rangle_1$ over elements $\langle v_j \rangle_1$, $\langle w_j \rangle_2$, $\langle y_j \rangle_1$ from the CRS, with the weights given by the witness part of the QAP's solution vector \mathbf{x} . The $\langle Z \rangle_1$ element ensures that these weighted sums consistently use the same witness. This is done by forcing the prover to come up essentially with $\beta \cdot (\langle V \rangle_1 + \langle W \rangle_2 + \langle Y \rangle_1)$ given only elements $\langle \beta \cdot (v_j + w_j + y_j) \rangle_1$ in which v_j , w_j , and y_j occur together. The essential idea of our construction is to use the $\langle Z \rangle_1$ element also to ensure consistency to external commitments.

In more detail, in earlier works [CFH⁺15, SVdV16], it was noted that the Pinocchio $\langle V \rangle_1, \langle W \rangle_2, \langle Y \rangle_1$ elements can be divided into multiple ‘‘blocks’’ ($\langle V_i \rangle_1, \langle W_i \rangle_2, \langle Y_i \rangle_1, \langle Z_i \rangle_1$). Each block contains the values of a number of variables of

⁴ We use $\langle \alpha_v V \rangle_2$ etc. instead of $\langle \alpha_v V \rangle_1$ from [SVdV16], so that we can rely on the asymmetric q -PKE assumption from [DFGK14] (which [SVdV16] did not spell out).

Extractable Trapdoor Commitment Scheme Family (G^1, Gc^1, C^1):

- G^1 : Fix $\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_3$ and random s . Return $\text{crs} = (\{\langle x^i \rangle_1, \langle x^i \rangle_2\}_{i=0, \dots, d})$, $\text{td} = s$.
- Gc^1 : Pick random α . Return $\text{ck} = (\langle 1 \rangle_1, \langle \alpha \rangle_2, \langle x \rangle_1, \langle \alpha x \rangle_2, \dots, \langle x^d \rangle_1, \langle \alpha x^d \rangle_2)$
- C^1 : Return $(r \langle 1 \rangle_1 + \mathbf{v}_1 \langle x \rangle_1 + \mathbf{v}_2 \langle x^2 \rangle_1 + \dots, r \langle \alpha \rangle_2 + \mathbf{v}_1 \langle \alpha x \rangle_2 + \mathbf{v}_2 \langle \alpha x^2 \rangle_2 + \dots)$

Key generation G^1 : Fix a QAP of degree at most d , and let $v_j(x), w_j(x), y_j(x)$ be as in Pinocchio. Fix random, secret $\alpha_v, \alpha_w, \alpha_y, \beta, r_v, r_w$. Let $r_y = r_v r_w$. Let $z_j(x) = x^j + r_v v_j + r_w w_j + r_y y_j$ if $j \leq W$ and $z_j(x) = r_v v_j + r_w w_j + r_y y_j$ otherwise. Evaluation key ($i = 1, \dots, n, j = 1, \dots, d$):

$$\langle x^j \rangle_1, \langle r_v v_j \rangle_1, \langle r_w t \rangle_1, \langle \alpha_v r_v v_j \rangle_2, \langle \alpha_v r_v t \rangle_2, \langle r_w w_j \rangle_1, \langle r_w t \rangle_1, \langle \alpha_w r_w w_j \rangle_1, \langle \alpha_w r_w t \rangle_1, \langle r_y y_j \rangle_1, \langle r_y t \rangle_1, \langle \alpha_y r_y y_j \rangle_2, \langle \alpha_y r_y t \rangle_2, \langle \beta_i z_{(i-1)d+j} \rangle_1, \langle \beta_i z_{nd+j} \rangle_1, \langle \beta_i \rangle_1, \langle \beta_i r_v t \rangle_1, \langle \beta_i r_w t \rangle_1, \langle \beta_i r_y t \rangle_1$$

Verification key ($i = 1, \dots, n$): $(\langle \alpha_v \rangle_2, \langle \alpha_w \rangle_1, \langle \alpha_y \rangle_2, \langle \beta_i \rangle_2, \langle \beta_i \rangle_1, \langle r_y t \rangle_2)$.

Proof generation P^1 : Let $\mathbf{u}_i = C_{\text{ck}_i}^1(\mathbf{v}_i; r_i)$, and let \mathbf{w} be the witness such that $(\mathbf{v}_1, \dots, \mathbf{v}_n; \mathbf{w})$ is a solution to the QAP. Generate random $\delta_{v,i}, \delta_{w,i}, \delta_{y,i}$. Compute \mathbf{h} as the coefficients of polynomial $((\sum_j \mathbf{x}_j \cdot v_j(x) + \delta_v \cdot t(x)) \cdot (\sum_j \mathbf{x}_j \cdot w_j(x) + \delta_w \cdot t(x)) - (\sum_j \mathbf{x}_j \cdot y_j(x) + \delta_y \cdot t(x)))/t(x)$. Return ($i = 1, \dots, n; [\cdot]$ means only if $i = 1$):

$$\begin{aligned} \langle V_i \rangle_1 &= \sum_{j=1}^d \mathbf{v}_{i,j} \langle r_v v_{(i-1)d+j} \rangle_1 \left[+ \sum_{j=1}^N \mathbf{w}_j \langle r_v v_{nd+j} \rangle_1 \right] + \delta_{v,i} \langle r_v t \rangle_1, \langle \alpha_v V_i \rangle_2 = \dots \\ \langle W_i \rangle_1 &= \sum_{j=1}^d \mathbf{v}_{i,j} \langle r_w w_{(i-1)d+j} \rangle_1 \left[+ \sum_{j=1}^N \mathbf{w}_j \langle r_w w_{nd+j} \rangle_1 \right] + \delta_{w,i} \langle r_w t \rangle_1, \langle \alpha_w W_i \rangle_1 = \dots \\ \langle Y_i \rangle_1 &= \sum_{j=1}^d \mathbf{v}_{i,j} \langle r_y y_{(i-1)d+j} \rangle_1 \left[+ \sum_{j=1}^N \mathbf{w}_j \langle r_y y_{nd+j} \rangle_1 \right] + \delta_{y,i} \langle r_y t \rangle_1, \langle \alpha_y Y_i \rangle_2 = \dots \\ \langle Z_i \rangle_1 &= \sum_{j=1}^d \mathbf{v}_{i,j} \langle \beta_i z_{(i-1)d+j} \rangle_1 \left[+ \sum_{j=1}^N \mathbf{w}_j \langle \beta_i z_{nd+j} \rangle_1 \right] + r_i \langle \beta_i \rangle_1 + \delta_{v,i} \langle \beta_i r_v t \rangle_1 \\ \langle H \rangle_1 &= \sum_j \mathbf{h}_j \langle x^j \rangle_1. \qquad \qquad \qquad + \delta_{w,i} \langle \beta_i r_w t \rangle_1 + \delta_{y,i} \langle \beta_i r_y t \rangle_1 \end{aligned}$$

Proof verification V^1 : Letting $\text{ck}_i = (\dots, \langle \alpha_i \rangle_2)$, $\mathbf{u}_i = (\langle C_i \rangle_1, \langle \alpha_i C_i \rangle_2)$, check that:

$$\begin{aligned} e(\langle C_i \rangle_1, \langle \alpha_i \rangle_2) &= e(\langle 1 \rangle_1, \langle \alpha_i C_i \rangle_2); & e(\langle V_i \rangle_1, \langle \alpha_v \rangle_2) &= e(\langle \alpha_v V_i \rangle_1, \langle 1 \rangle_2); & (\mathbf{C}, \mathbf{V}) \\ e(\langle \alpha_w \rangle_1, \langle W_i \rangle_2) &= e(\langle 1 \rangle_1, \langle \alpha_w W_i \rangle_2); & e(\langle Y_i \rangle_1, \langle \alpha_y \rangle_2) &= e(\langle 1 \rangle_1, \langle \alpha_y Y_i \rangle_2); & (\mathbf{W}, \mathbf{Y}) \\ e(\langle V_i \rangle_1 + \langle Y_i \rangle_1 + \langle C_i \rangle_1, \langle \beta_i \rangle_2) \cdot e(\langle \beta_i \rangle_1, \langle W_i \rangle_2) &= e(\langle Z_i \rangle_1, \langle 1 \rangle_2); & (\mathbf{Z}) \\ e(\langle V \rangle_1, \langle W \rangle_2) \cdot e(\langle Y \rangle_1, \langle 1 \rangle_2)^{-1} &= e(\langle H \rangle_1, \langle r_y t \rangle_2). & (\mathbf{H}) \end{aligned}$$

(where $\langle V \rangle_1 = \langle V_1 \rangle_1 + \dots + \langle V_n \rangle_1$, $\langle W \rangle_2 = \langle W_1 \rangle_2 + \dots + \langle W_n \rangle_2$, $\langle Y \rangle_1 = \langle Y_1 \rangle_1 + \dots$)

Fig. 1. Pinocchio-Based Adaptive zk-SNARK (G^1, P^1, V^1)

the QAP solution, which is enforced by providing $\langle z_j \rangle_1 = \langle \beta_i \cdot (v_j + w_j + y_j) \rangle_1$ elements only for the indices j of those variables. Our core idea is use external commitments of the form $\sum_k \mathbf{v}_k \cdot \langle x^k \rangle_1$ (that can be re-used across Pinocchio computations) and link the k th component of this commitment to the j th variable of the block using a modified $\langle z_j \rangle_1 = \langle \beta_i \cdot (x^k + v_j + w_j + y_j) \rangle_1$. We use one block per external commitment that the proof refers to. The witness (which is not committed to externally) is included in the first block, with the normal

Pinocchio element $\langle z_j \rangle_1 = \langle \beta_1 \cdot (v_j + w_j + y_j) \rangle_1$ just checking internal consistency as usual. The verification procedure changes slightly: $\langle V \rangle_1$ is no longer extended to $\langle V^+ \rangle_1$ to include public I/O (which we do not have); instead, the **(Z)** check ensures consistency with the corresponding commitment, for which there is a new correctness check **(C)**.

The precise construction is shown in Figure 1. This construction contains details on how to add randomness to make the proof zero-knowledge; and it shows how additional $\langle \alpha \cdot \rangle_1$ elements are added to obtain an extractable trapdoor commitment family $(\mathbf{G0}^1, \mathbf{Gc}^1, \mathbf{C}^1)$. In Appendix A, we show that:

Theorem 1. *Under the $(4d + 3)$ -PDH, d -PKE, and $(8d + 6)$ -SDH assumptions (with d the maximal QAP degree), $(\mathbf{G0}^1, \mathbf{Gc}^1, \mathbf{C}^1)$ is an extractable trapdoor commitment scheme family, and $(\mathbf{G}^1, \mathbf{P}^1, \mathbf{V}^1)$ is an adaptive zk-SNARK.*

4 Smaller Proofs and Comparison to Literature

We now present two optimizations that decrease the size of the above zk-SNARK, and compare the concrete efficiency of our three proposals to two related proposals from the literature. Note that, in the above construction, seven Pinocchio proof elements $\langle V \rangle_1, \langle \alpha_v V \rangle_2, \langle W \rangle_2, \langle \alpha_w W \rangle_1, \langle Y \rangle_1, \langle \alpha_y Y \rangle_2, \langle Z \rangle_1$ are repeated for each input commitment. We present two different (but, unfortunately, mutually incompatible) ways in which this can be avoided.

In our first optimization, inspired by a similar proposal to reduce verification work in Pinocchio ([CFH⁺15], later corrected by [Par15]), we decrease proof size and verification time at the expense of needing a larger-degree QAP. Namely, suppose that all variables in a given commitment occur only in the right-hand side of QAP equations. In this case, $v_j(x) = w_j(x) = 0$ for all j , so proof elements $\langle V_i \rangle_1, \langle \alpha_v V_i \rangle_2, \langle W_i \rangle_2, \langle \alpha_w W_i \rangle_1, \langle Y_i \rangle_1, \langle \alpha_y Y_i \rangle_2$ contain only randomness and, setting $\delta_{v,j} = \delta_{w,j} = 0$, can be omitted. As a consequence, the marginal costs per commitment used decrease from 7 to 3; the **(V)** and **(W)** verification steps can be skipped and the **(Z)** step simplified. To guarantee that a committed variable a only occurs in the right-hand of equations, we can introduce a witness b and equation $0 \cdot 0 = a - b$, slightly increasing the overall QAP size and degree. (This cannot be done for the first commitment since $\langle V_1 \rangle_1, \dots$ also contain the witness, which occur in the left-hand side of equations as well.)

Our second proposal is a modified zk-SNARK that also reduces the marginal cost per commitment from 7 to 3, but gives more efficient verification when using many commitments. The core idea is to first concatenate all commitments $\mathbf{u}_1, \dots, \mathbf{u}_n$ into one “intermediate commitment \mathbf{u}' ”, and then use our original zk-SNARK with respect to \mathbf{u}' . More precisely, we build intermediate commitment \mathbf{u}'_1 with the first ℓ_1 values of \mathbf{u}_1 ; \mathbf{u}'_2 with ℓ_1 zeros followed by the first ℓ_2 values of \mathbf{u}_2 ; etcetera. Then, $\mathbf{u}' = \sum_i \mathbf{u}'_i$ is a commitment to the first ℓ_1, \dots, ℓ_n values of the respective commitments $\mathbf{u}_1, \dots, \mathbf{u}_n$. To avoid ambiguity between normal and intermediate commitments, to normal commitments we add a random factor r_c , i.e. $(r \langle r_c \rangle_1 + \sum_i \mathbf{v}_i \langle r_c x^i \rangle_1, r \langle \alpha r_c \rangle_2 + \dots)$ and intermediate commitments are

Construction	Comm. Proof size		Prover computation		Verif comp
	Size		non-crypt. op.	crypt. op.	
Geppetto	3 gr. el.	8 gr. el.	$\Theta(D \log D)$	$\Theta(D)$	$4n + 12$ pair.
Hash First+Pinocchio	2 gr. el.	$9n+1$ gr. el.	$\Theta(d \log d)$	$\Theta(d)$	$13n + 3$ pair.
Hash First+Pinocchio*	2 gr. el.	$5n+5$ gr. el.	$\Theta(d' \log d')$	$\Theta(d')$	$8n + 8$ pair.
Our zk-SNARK I	2 gr. el.	$7n+1$ gr. el.	$\Theta(d \log d)$	$\Theta(d)$	$11n + 3$ pair.
Our zk-SNARK I*	2 gr. el.	$3n+5$ gr. el.	$\Theta(d' \log d')$	$\Theta(d')$	$7n + 7$ pair.
Our zk-SNARK II	2 gr. el.	$3n+8$ gr. el.	$\Theta(d \log d)$	$\Theta(d)$	$6n + 12$ pair.

Table 1. Comparison between Pinocchio-based SNARKs (n : number of commitments; d is QAP degree; $d' \leq d$ is QAP degree with optimization; $D \geq d$ is fixed QAP degree)

as above⁵. Proving correspondence between normal and intermediate commitments is done similarly to the (\mathbf{Z}) check above: we generate random β_i and give $\langle \beta'_i \cdot (r_c x^j + x^{\ell_1 + \dots + \ell_{i-1} + j})_1 \rangle$ to the prover, who needs to produce proof element $\langle Z'_i \rangle_1$ such that $\langle Z'_i \rangle_1 = \beta'_i \cdot (\langle C_i \rangle_1 + \langle C'_i \rangle_1)$, which he can only do if $\langle C'_i \rangle_1$ is formed correctly. Details and the security proof appear in Appendix B.

In Table 1, we provide a detailed comparison of our zk-SNARKs with two similar constructions: the Geppetto protocol due to [CFH⁺15] (which is also zero-knowledge but not adaptive); and the “hash first” approach applied to Pinocchio [FFG⁺16] (which is adaptive but not zero-knowledge). Geppetto is Protocol 2 from [CFH⁺15]. We assume QAP witnesses of $\mathcal{O}(d)$. In Geppetto, a fixed set of QAPs of degree d_i are combined into one large “MultiQAP” of degree D slightly larger than $\max d_i$. As a consequence, if both small and large computations need to be applied on the same data, then the small computations take over the much worse performance of the large computations. For Hash First+Pinocchio, we took the extractable scheme $\text{XP}_\mathcal{E}$ since the Geppetto and our construction are extractable as well. To make it work on multiple commitments (which is described for neither Hash First nor Pinocchio), we assume natural generalisations of Hash First and of Pinocchio along the lines of [CFH⁺15,SVdV16]. Our first optimization can be applied to this construction; we mark the result with a star and write $d' \geq d$ for the increased degree due to the use of this optimization. Finally, we show our zk-SNARK without and with the first optimization; and our second zk-SNARK construction (to which the optimization does not apply).

In conclusion, Geppetto is the most efficient construction, but apart from not being adaptive, it also requires all computations to be fixed and of the same size, making it inefficient for small computations when they are combined with large ones. Our construction outperforms Hash First+Pinocchio, essentially adding zero knowledge for free; which variant is best depends on n and $d' - d$. Note that Hash First allows using the same commitment in different zk-SNARK schemes; our scheme only allows this for zk-SNARKs based on the kind of polynomial commitments used in Pinocchio.

⁵ Hence this construction can only handle inputs of combined size at most d .

5 Secure/Correct Adaptive Function Evaluation

In this section, we sketch how our zk-SNARK can be used to perform “adaptive function evaluation”: privacy-preserving verifiable computation on committed data. We consider a setting in which multiple mutually distrusting *data owners* want to allow privacy-preserving outsourced computations on their joint data. A *client* asks a computation to be performed on this data by a set of *workers*. The input data is sensitive, so the workers should not learn what data they are computing on (assuming up to a maximum number of workers are passively corrupted). On the other hand, the client wants to be guaranteed the computation result is correct, for instance, with respect to a commitment to the data published by the data owner (making no assumption on which data owners and/or workers are actively corrupted). The difference in assumptions for the privacy and correctness guarantees is motivated by settings where data owners together choose the computation infrastructure (so they feel active corruption is unlikely) but need to convince an external client (e.g. a medical reviewer) of correctness. We work in the CRS model, where a trusted party (who is otherwise not involved in the system) performs one-time key generation.

In Appendix C, we provide a precise security model that captures the above security guarantees by ideal functionalities. We define two ideal functionalities. The first guarantees privacy *and* correctness (*secure adaptive function evaluation*), and is realised by our construction if at most a threshold of workers are passively corrupted (but all other parties can be actively corrupted). The second guarantees only correctness (*correct adaptive function evaluation*), and is realised by our construction regardless of corruptions.

5.1 Our Construction

We now present our general construction based on multi-party computation and *any* adaptive zk-SNARK (as we will see later, our adaptive zk-SNARK gives a particularly efficient instantiation). At a high level, to achieve secure adaptive function evaluation, the workers compute the function using multi-party computation (MPC), guaranteeing privacy and correctness under certain conditions. However, when these conditions are not met, we still want to achieve correct adaptive function evaluation, i.e., we still want to ensure a correct computation result. To achieve this, the workers also produce, using MPC, a zk-SNARK proof of correctness of the result.

We require a MPC protocol in the outsourcing setting, i.e., with separate inputters (in our case, the data owners and the client), recipients (the client) and workers. The protocol needs to be reactive, so that the data owners can provide their input before knowing the function to be computed⁶; and secure even if any number of data owners the client are actively corrupted. Security of the

⁶ Using non-reactive MPC requires is also possible, but then steps 3 and 4 of the protocol need to be swapped. As a consequence, data owners can abort based on the client’s choice of function, leading to a weaker form of correct function evaluation.

Protocol Adaptive Trinocchio

(Data provider has $\mathbf{a}_i \in \mathbb{F}^d$; client has $\mathbf{a}_c \in \mathbb{F}^{d'}$, function $f : (\mathbb{F}^d)^n \times \mathbb{F}^{d'} \rightarrow \mathbb{F}^{d-d'}$.)

1. The trusted party generates a system-wide CRS crs of the trapdoor commitment family, and commitment keys $\text{ck}_1, \dots, \text{ck}_n, \text{ck}_c$ for the data owner and client. This material is distributed to all parties.
2. Each data owner computes commitment $c_i = \text{C}_{\text{ck}_i}(\mathbf{a}_i, r_i)$ to its input $\mathbf{a}_i \in \mathbb{F}^d$ using randomness r_i , and publishes it on a bulletin board.
3. The data owners, workers, and client use the MPC protocol to do the following:
 - Each data owner provides input \mathbf{a}_i and randomness r_i
 - For each i , compute $c'_i = \text{C}_{\text{ck}_i}(\llbracket \mathbf{a}_i \rrbracket, \llbracket r_i \rrbracket)$; if $c_i \neq c'_i$ then abort
4. The client provides function f to the trusted party. The trusted party determines a QAP Q computing f and a function f' solving Q , and performs key generation of the adaptive zk-SNARK (where one commitment combines the client's input and output). The client gets verification key crsv ; the workers get Q , f' , and the corresponding evaluation key crsp .
5. The data owners, workers, and client continue with the MPC from step 3:
 - Client: provide input \mathbf{a}_c
 - Compute $(\llbracket \mathbf{b} \rrbracket; \llbracket \mathbf{w} \rrbracket) \leftarrow f'(\llbracket \mathbf{a}_1 \rrbracket; \dots; \llbracket \mathbf{a}_n \rrbracket; \llbracket \mathbf{a}_c \rrbracket)$
 - Compute $\llbracket c_c \rrbracket \leftarrow \text{C}_{\text{ck}_c}(\llbracket \mathbf{a}_c \rrbracket, \llbracket \mathbf{b} \rrbracket; \llbracket r_c \rrbracket)$ for random r_c
 - Compute $\llbracket \pi \rrbracket \leftarrow \text{P}(\text{crs}, \{\text{ck}_i\}, \dots, \text{ck}_n, \text{ck}_c, \text{crsp}; c_1, \dots, c_n, \llbracket c_c \rrbracket; \llbracket \mathbf{a}_1 \rrbracket; \dots; \llbracket \mathbf{a}_n \rrbracket; \llbracket \mathbf{a}_c \rrbracket, \llbracket \mathbf{b} \rrbracket; \llbracket r_1 \rrbracket, \dots, \llbracket r_n \rrbracket, \llbracket r_c \rrbracket; \llbracket \mathbf{w} \rrbracket)$
 - Open outputs $\llbracket \mathbf{b} \rrbracket, \llbracket r_c \rrbracket, \llbracket c_c \rrbracket, \llbracket \pi \rrbracket$ to the client
6. The client checks whether $\text{V}(\text{crs}, \text{ck}_1, \dots, \text{ck}_n, \text{ck}_c, \text{crsv}; c_1, \dots, c_n, c_c; \pi) = 1$ and $c_c = \text{C}_{\text{ck}_c}(\mathbf{a}_c, \mathbf{b}; r_c)$ and if so, returns computation result \mathbf{b} .

Fig. 2. The Adaptive Trinocchio protocol

MPC protocol will generally depend on how many workers are corrupted; our construction will realise secure adaptive function evaluation (as opposed to just correct adaptive function evaluation) exactly when the underlying MPC protocol is secure. (As we show below, MPC protocols based on (t, n) -Shamir secret sharing (e.g., [dH12]) between $n = 2t + 1$ workers satisfy these requirements.)

Our protocol is shown in Figure 2. It uses an MPC protocol with the above properties, a trapdoor commitment family, and an adaptive zk-SNARK, instantiated for the function to be computed. The protocol relies on a trusted party that generates the key material of the zk-SNARK, but is otherwise not involved in the computation. Each data owner has an input $\mathbf{a}_i \in \mathbb{F}^d$ and the client has an input $\mathbf{a}_c \in \mathbb{F}^{d'}$ and a function $f : (\mathbb{F}^d)^n \times \mathbb{F}^{d'} \rightarrow \mathbb{F}^{d-d'}$ that it wants to compute on the combined data. Internal variables of the MPC protocol are denoted $\llbracket \cdot \rrbracket$.

In step 1, the trusted party sets up the trapdoor commitment family, generating separate keys for data providers and the client. (This prevents parties from copying each other's input.) In step 2, each data provider publishes a commitment to his input. In step 3, each data providers inputs its data and the randomness used for the commitment to the MPC protocol. The workers re-compute

the commitments based on this opening and abort in case of a mismatch. (This prevents calling P on mismatching inputs in which case it may not be zero-knowledge.) In step 4, the client chooses the function f to be computed, based on which the trusted party performs key generation. (By doing this after the data owners' inputs, we prevent a selective failure attack from their side.) In step 5, the computation is performed. Using MPC, the client's output and witness are computed; a commitment to the client's I/O is produced, and a zk-SNARK proof of correctness with respect to the commitments of the data owners and client is built.⁷ The client learns the output, randomness for its commitment, the commitment itself, and the proof. In step 6, the client re-computes the commitment and verifies the proof; in case of success, it accepts the output.

By sharing commitments between proofs, it is possible to generate key material for a number of small building blocks once, and later flexibly combine them into larger computations without requiring new key material. In particular, as we show in the case study, this enables computations on arbitrary-length data using the same key material (which was impossible in Trinocchio). It is also easy to support multiple clients or multiple commitments per data owner.

In Appendix C, we show that indeed, the above construction achieves the formal definitions of secure adaptive function evaluation (under the same conditions of the corruptions of workers as the underlying MPC protocol) and correct adaptive function evaluation (regardless of corruptions).

5.2 Efficient Instantiation using Secret Sharing and our zk-SNARK

We now show that our zk-SNARKs and MPC based on Shamir secret sharing give a particularly efficient instantiation of the above framework. The idea is the same as for Trinocchio [SVdV16]: our zk-SNARK is essentially an arithmetic circuit of multiplicative depth 1, so given a solution to the QAP, the prover algorithm can be performed under MPC without any communication between the workers.

In more detail, we perform MPC based on Shamir secret sharing between the m workers (e.g., [dH12]). This guarantees privacy as long as at most t workers are passively corrupted, where $m = 2t + 1$. Inputs are provided by the inputters as an additive sharing between all workers: this way actively corrupted inputters cannot provide an inconsistent sharing. The workers Shamir-share and sum up the additive shares to obtain a Shamir sharing of the input. Outputs are provided to recipients either as Shamir shares or as freshly randomised additive shares: the latter allows producing our zk-SNARK proof without any communication.

Either of our zk-SNARK constructions can be used; we provide details for the first one. Below, write $\llbracket \cdot \rrbracket$ for Shamir sharing and $[\cdot]$ for additive sharing. (Note that Shamir sharings can be converted locally to additive sharings at no cost.) In step 3 of the protocol, to open c'_i , the parties apply C_{ck_i} on their additive shares of the input and randomness, add a random additive sharing of zero (which can be generated non-interactively using pseudo-random zero sharing), and reveal

⁷ Equivalently, the workers can open c_c and π and send them to the client in the plain.

the result. In step 5, $[\mathbf{b}]$; $[\mathbf{w}]$ are computed as Shamir secret shares. Next, $[c_c]$ is computed as an additive sharing by applying C_{ck_c} on additive shares and adding a random sharing of zero. Next, P^1 is applied by performing the following steps:

- Generate $\delta_{v,i}, \delta_{w,i}, \delta_{y,i}$ by pseudo-random secret sharing.
- Compute $[\mathbf{h}] = ((\sum_j [\mathbf{x}_j] \cdot v_j(x) + [\delta_v] \cdot t(x)) \cdot (\sum_j [\mathbf{x}_j] \cdot w_j(x) + [\delta_w] \cdot t(x)) - (\sum_j [\mathbf{x}_j] \cdot y_j(x) + [\delta_y] \cdot t(x))) / t(x)$. Essentially this is done by performing the computation straight on Shamir secret shares; because there is only layer of multiplications of shares, this directly gives an additive sharing of the result. Smart use of FFTs gives time complexity $\mathcal{O}(d \cdot \log d)$ [BSCG⁺13,SVdV16].
- All proof elements are now linear combinations of secret-shared data; compute them by taking linear combinations of the (Shamir or additive) shares and adding a random sharing of zero.

What remains is how to compute the solution of the QAP using multi-party computation. Namely, in addition to computing the function result $[\mathbf{b}]$, the MPC also needs to compute witness $[\mathbf{w}]$ to the QAP. Actually, if the function to be computed is described as an arithmetic circuit, this is very easy. Namely, in this case, the witness for the natural QAP for the function is exactly the vector of results of all intermediate multiplications; and these results are already available as Shamir secret shares as a by-product of performing the MPC. Hence, in this case, computing $[\mathbf{w}]$ in addition to $[\mathbf{b}]$ incurs no overhead.

If a custom MPC protocol for a particular subtask is used, then it is necessary to devise specific QAP equations and an MPC protocol to compute their witness. As an example, consider the MPC operation $[\mathbf{b}] \leftarrow [a \neq 0]$, i.e., b is assigned 1 if $a \neq 0$ and 0 if $a = 0$. For computing $[\mathbf{b}]$, a fairly complex protocol is needed, cf. [dH12]. However, proving that b is correct using a QAP is simple [PHGR13]: introduce witnesses $c := (a + (1 - b))^{-1}$, $d := 1$ and equations:

$$a \cdot c = b \quad a \cdot (d - b) = 0 \quad d \cdot d = d.$$

Indeed, if $a = 0$ then the first equation implies that $b = 0$; if $a \neq 0$ then the second and third equations imply that $b = 1$. In both cases, the given value for c and $d = 1$ make all three equations hold. In our case study, we show similarly how, for complex MPC protocols for fixed-point arithmetic, simple QAPs proving correctness exist with easily computable witnesses.

6 Prototype and Distributed Medical Research Case

In this section, we present a proof-of-concept implementation of our second zk-SNARK construction and our Adaptive Trinocchio protocol. Computations can be specified at a high level using a Python frontend; executed either locally or in a privacy-preserving way using multi-party computation; and then automatically proven and verified to be correct by a C++ backend. We show how two different computations can be performed on the same committed data coming from multiple data owners (with key material independent from input length, and optionally in a privacy-preserving way): aggregate survival statistics on two patient

populations, and the “logrank test”: a common statistical test whether there is a statistically significant difference survival rate between the populations.

6.1 Prototype of our zk-SNARK and Adaptive Trinocchio

Our prototype, available at <https://github.com/meilof/geppetri>, is built on top of VIFF, a Python implementation of MPC based on Shamir secret sharing. In VIFF, computations on secret shares are specified as normal computations by means of operator overloading, e.g., assigning $c=a*b$ induces a MPC multiplication protocol. We add a new runtime to VIFF that also allows computations to be performed locally without MPC.

To support computation proofs, we developed the `viffvc` library that provides a new data type: `VcShare`, a wrapper around a secret share. Each `VcShare` represents a linear combination of QAP variables. Addition and multiplication by constants of `VcShares` is performed locally by manipulating the linear combination. Constants v are represented as $v \cdot one$, where witness one satisfies $one \cdot one = one$ so $one = 1$. When two `VcShares` $\lambda_1 x_1 + \dots$ and $\mu_1 x_1 + \dots$ are multiplied, a local or MPC multiplication operation is performed on the underlying data, and the result is a new `VcShare` x_k wrapping the result as a new QAP variable. QAP equation $(\lambda_1 x_1 + \dots) \cdot (\mu_1 x_1 + \dots) = 1 \cdot x_k$ is written to a file, and the multiplication result x_k or its secret share, when known, is written to another file. Apart from multiplication, some additional operations are supported. For the $\llbracket b \rrbracket \leftarrow \llbracket a \neq 0 \rrbracket$ operation discussed in Section 5.2, the implementation computes $\llbracket b \rrbracket$ and $\llbracket c \rrbracket = (\llbracket a \rrbracket + (1 - \llbracket b \rrbracket))^{-1}$, and writes these secret shares/values and the equations from Section 5.2 to the respective files. We also support secret indexing (e.g., [dH12]), and fixed-point computations as discussed below.

Computations are performed by this custom VIFF-based system together with an implementation of our zk-SNARK. A first tool, `qapgen`, generates the CRS for our trapdoor commitment scheme. A second tool, `qapinput`, builds a commitment to a given input; and computes secret shares of these inputs that are used for MPC computations. Then, our Python implementation is used to compute the function, either locally or using multi-party computation. At the end of this execution, there is one file with the QAP equations, and one file with values/shares for each QAP variable. Our `qapgenf` tool uses the first file to perform key generation of the QAP (this is done only once and for next executions, previous key material is re-used). Our `qapprove` tool uses the second file to generate the zk-SNARK proof (shares) to be received by the client. Finally, a `qapver` tool verifies the proof based on the committed inputs and outputs.

6.2 Application to Medical Survival Analysis

We have applied our prototype to (adaptively) perform computations on survival data about two patient populations. In medical research, survival data about a population is a set of tuples (n_j, d_j) , where n_j is the number of patients still in the study just before time j and d_j is the number of deaths at time j . We assume

Algorithm 1 Anonymized survival data computation

Require: $[[\mathbf{d}_1], [\mathbf{n}_1], [\mathbf{d}_2], [\mathbf{n}_2]]$: block of survival data points for two populations

Ensure: $([[d'_1], [n'_1], [d'_2], [n'_2]])$ aggregated survival data for the block

1: **function** SUMM($[[\mathbf{d}_{i,1}], [\mathbf{d}_{i,2}], [\mathbf{n}_{i,1}], [\mathbf{n}_{i,2}]]$)

2: **return** $(\sum_i [[\mathbf{d}_{1,i}], [\mathbf{n}_{1,i}], \sum_i [[\mathbf{d}_{2,i}], [\mathbf{n}_{2,i}]]]$

both populations are distributed among multiple hospitals, that each commit to their contributions $(d_{j,1}, n_{j,1}, d_{j,2}, n_{j,2})$ to the two populations at each time.

Aggregate Survival Data Our first computation is to compute an aggregate version of the survival data, where each block $\{d_{j,1}, n_{j,1}, d_{j,2}, n_{j,2}\}_{j=1}^{25}$ of 25 time points is summarised as $(\sum_j d_{j,1}, n_{1,1}, \sum_j d_{j,2}, n_{1,2})$. The function SUMM computing this summary is shown in Algorithm 1. Function SUMM translates into a QAP on 26 commitments: as input, for each time point j , a commitment $\sum_i c_{i,j}$ to the combined survival data $([[\mathbf{d}_{i,1}], [\mathbf{n}_{i,1}], [\mathbf{d}_{i,2}], [\mathbf{n}_{i,2}]])$ from the different hospitals i at that time (using the fact that commitments are homomorphic); as output, a commitment to $([[d'_1], [n'_1], [d'_2], [n'_2]])$.

Logrank test Our second computation is the so-called “Mantel-Haenzel logrank test”, a statistical test to decide whether there is a significant difference in survival rate between the two populations (as implemented, e.g., in R’s `survdif` function). Given the survival data from two populations, define:

$$E_{j,1} = \frac{(d_{j,1} + d_{j,2}) \cdot n_{j,1}}{n_{j,1} + n_{j,2}}; \quad V_j = \frac{n_{j,1}n_{j,2}(d_{j,1} + d_{j,2})(n_{j,1} + n_{j,2} - d_{j,1} - d_{j,2})}{(n_{j,1} + n_{j,2})^2 \cdot (n_{j,1} + n_{j,2} - 1)};$$
$$X = \frac{\sum_j E_{j,1} - \sum_j d_{j,1}}{\sum_j V_j}.$$

The null hypothesis for the logrank test, i.e., the hypothesis that the two curves represent the same underlying “survival function”, corresponds to $X \sim \chi_1^2$. This null hypothesis is rejected (i.e., the curves are different) if $1 - \text{cdf}(X) > \alpha$, where cdf is the cumulative density function of the χ_1^2 distribution and, e.g., $\alpha = 0.05$. We use MPC to compute X , and then apply the cdf in the clear.

Our implementation consists of two different functions: a function BLOCK (Algorithm 2) that computes $(E_{j,1}, V_j, d_{j,1})$ given the survival data at point j ; and a function FIN that, given $\sum E_{j,1}$, $\sum V_j$, and $\sum d_{j,1}$ computes X (Algorithm 3). As above, function BLOCK is applied to commitment $\sum_i c_{i,j}$ to the combined survival data from different hospitals at a particular time, giving output commitment c'_j . Function FIN is applied to commitment $\sum_j c'_j$ to $(\sum E_{j,1}, \sum V_j, \sum d_{j,1})$, again using the fact that commitments are homomorphic; outputting a commitment to X that is output to the client.

Algorithms 2 and 3 use fixed-point numbers $[[x]]^f$, representing value $x \cdot 2^{-k}$ where we use precision $k = 20$. We use the fixed-point multiplication $[[c]]^f \leftarrow [[a]]^f \cdot [[b]]^f$ and division $[[c]]^f \leftarrow [[a]]^f / [[b]]^f$ protocols due to [dH12].

Algorithm 2 Logrank computation for each time step

Require: $[[\mathbf{d}_{i,1}], [\mathbf{d}_{i,2}], [\mathbf{n}_{i,1}], [\mathbf{n}_{i,2}]]$ survival data at time point i

Ensure: $([e_i]^f, [v_i]^f, [d_i])$ contributions to $\sum_j E_{j,1}$, $\sum_j V_j$, $\sum_j d_{j,1}$ for test statistic

```
1: function BLOCK( $[[\mathbf{d}_{i,1}], [\mathbf{d}_{i,2}], [\mathbf{n}_{i,1}], [\mathbf{n}_{i,2}]]$ )
2:    $[[ac]] \leftarrow [[\mathbf{d}_{i,1}]] + [[\mathbf{d}_{i,2}]]$ 
3:    $[[bd]] \leftarrow [[\mathbf{n}_{i,1}]] + [[\mathbf{n}_{i,2}]]$ 
4:    $[[frc]]^f \leftarrow [[ac]] / [[bd]]$ 
5:    $[e_i]^f \leftarrow [[frc]]^f \cdot [[\mathbf{n}_{i,1}]]$ 
6:    $[[vn]] \leftarrow [[\mathbf{n}_{i,1}]] \cdot [[\mathbf{n}_{i,2}]] \cdot [[ac]] \cdot ([[bd]] - [[ac]])$ 
7:    $[[vd]] \leftarrow [[bd]] \cdot [[bd]] \cdot ([[bd]] - 1)$ 
8:    $[v_i]^f \leftarrow [[vn]] / [[vd]]$ 
9:   return  $([e_i]^f, [v_i]^f, [d_i])$ 
```

Algorithm 3 Logrank final computation

Require: $[[es], [vs], [ds]]$: summed-up values required to compute X

Ensure: $[[chi]]^f$ test statistic comparing two curves; supposedly $chi \sim \chi_1^2$

```
1: function FIN( $[[es], [vs], [ds]]$ )
2:    $[[ds]]^f \leftarrow [ds] \ll \text{PRECISION}$ 
3:    $[[dmi]]^f \leftarrow [ds]^f - [vs]^f$ 
4:    $[[chi0]]^f \leftarrow [dmi]^f / [vs]^f$ 
5:    $[[chi]]^f \leftarrow [chi0]^f \cdot [dmi]^f$ 
6:   return  $[[chi]]^f$ 
```

To prove that $[[c]]^f \leftarrow [[a]]^f \cdot [[b]]^f$ is correct, note that we need to show that $2^k c - a \cdot b \in [-2^k, 2^k]$, or equivalently, that $\alpha := 2^k c - a \cdot b + 2^k \geq 0$ and $\beta := 2^k - (2^k c - a \cdot b) \geq 0$. We prove this by computing, using MPC, bit decompositions [dH12] $\alpha = \alpha_0 + \alpha_1 \cdot 2 + \dots + \alpha_k \cdot 2^k$ and $\beta = \beta_0 + \beta_1 \cdot 2 + \dots + \beta_k \cdot 2^k$ (indeed, α and β are $\leq k + 1$ bits long); these α_i, β_i are the witnesses to QAP equations:

$$\begin{aligned} \forall i : \alpha_i \cdot (1 - \alpha_i) = 0 \quad c - a \cdot b + 2^k = \alpha_0 + \alpha_1 \cdot 2 + \dots + \alpha_k \cdot 2^k \\ \forall i : \beta_i \cdot (1 - \beta_i) = 0 \quad \beta = 2^k - (c - a \cdot b) = \beta_0 + \beta_1 \cdot 2 + \dots + \beta_k \cdot 2^k. \end{aligned}$$

Similarly, note that $[[c]]^f \leftarrow [[a]]^f / [[b]]^f$ is correct if and only if $2^k a - b \cdot c \in [-b, b]$, i.e., $\gamma := b + 2^k a - b \cdot c \geq 0$ and $\delta := b - (2^k a - b \cdot c) \geq 0$. If b has bitlength at most K (i.e., the represented number has absolute value $\leq 2^{K-k}$), then γ and δ have at most $K + 1$ bits. As above, we prove correctness by determining $(K + 1)$ -length bit decompositions of γ and δ and proving them correct. Proving correctness of $[[c]]^f \leftarrow [[a]] / [[b]]$ is analogous.

Performance Table 2 shows the performance of our proof-of-concept implementation for computing aggregate survival data and the logrank test (on a modern laptop). As input, we used the “btrial” data set included in R’s “kmsurv” package (on which we indeed reproduce R’s `survdifff` result) of 175 data points. Apart from having one data point per commitment, we also experiment with having a “block size” of 25 or 175 data points. For the logrank test, we use one

Aggregate	Computation (function): 0.0s (w/o MPC) / 0.1s (w/ MPC)		
	Computation (function+witness): 0.0s (w/o MPC) / 0.1s (w/ MPC)		
	BS=1	BS=25	BS=175
	QAP degree: 3 Prover: 0.3s/0.4s Verifier: 1.2s/1.5s	QAP degree: 3 Prover: 0.1s/0.1s Verifier: 0.2s/0.2s	QAP degree: 57 Prover: 0.0s/0.0s Verifier: 0.0s/0.0s
Logrank	Computation (function): 0.2s (w/o MPC) / 190.5s (w/ MPC)		
	Computation (function+witness): 0.6s (w/o MPC) / 235.2s (w/ MPC)		
	BS=1	BS=25	BS=175
	QAP deg (block): 173 QAP deg (fin): 85 Prover: 13.9s/78.5s Verifier: 3.9s/4.9s	QAP deg (block): 4304 QAP deg (fin): 85 Prover: 16.2s/81.0s Verifier: 0.2s/0.3s	QAP deg (block): 30104 QAP deg (fin): 85 Prover: 9.8s/73.5s Verifier: 0.0s/0.0s

Table 2. Performance: computation/proving/verification; with/without MPC

QAP per block; larger blocks mean less work for the verifier (since there are fewer proofs) but, in theory, more work for the prover (since the proving algorithm is superlinear in the QAP size). For aggregation, we use one QAP per 25 data points or per commitment, whichever is more.

We time the performance of running the computation, producing the proof, and verifying it, with or without MPC. As expected, MPC induces a large overhead for the computation, especially for the logrank test (due to the many fixed-point computations). MPC also incurs an overhead for proving: this is because of the many exponentiations with $|\mathbb{F}|$ -sized secret shares rather than small witnesses. Note that proving is faster than computing with MPC: the underlying operations are slower [SVdV16], but the QAP proof is in effect on a verification circuit that is smaller than the circuit of the computation itself [dHSV16]. Proving is faster for block size 175 than block size 25, which is unexpected; this may be because our FFT subroutine rounds up QAP degrees to the nearest power of two, which is favourable in the 175-sized case but not in the 25-sized case. As expected, verification is faster for larger block sizes. (The overhead of MPC here is due to recombining the proof shares into one overall proof.)

7 Conclusion

In this work, we have given the first practical Pinocchio-based adaptive zk-SNARK; applied it in the privacy-preserving outsourcing setting; and presented a proof-of-concept implementation. We mention a few promising directions for follow-ups. Concerning our construction for making Pinocchio adaptive, it would be interesting to see if it can be applied to make recent, even more efficient zk-SNARKS (e.g., [Gro16] in the generic group model) adaptive as well. Moreover, apart from providing a non-adaptive zk-SNARK, Geppetto also introduces the interesting idea of proof bootstrapping, where the verification procedure of the zk-SNARK itself can be performed by means of a verifiable computation, so

multiple related proofs can be verified in constant time. Applying this technique in our setting should combine our flexibility with their constant-time verification.

Concerning our privacy-preserving outsourcing framework, it is interesting to see if, apart from secret sharing plus our SNARK, other appealing instantiations are possible. Also, the combination of MPC and verifiable computation raises the challenge to construct efficient QAPs for specific operations and build efficient MPC protocols for computing their witnesses. We have presented zero testing and fixed-point computations as examples, but the same idea is applicable to many other operations as well. More generally, extending our zk-SNARK prototype with more basic operations, and improving its user-friendliness, would help bring the techniques closer to practice.

Acknowledgements This work is part of projects that have received funding from the European Union’s Horizon 2020 research and innovation programme under grant agreement No 643964 (SUPERCLOUD) and No 731583 (SODA).

References

- [BBFR15] M. Backes, M. Barbosa, D. Fiore, and R. M. Reischuk. ADSNARK: Nearly Practical and Privacy-Preserving Proofs on Authenticated Data. In *Proceedings S&P*, 2015.
- [BSCG⁺13] E. Ben-Sasson, A. Chiesa, D. Genkin, E. Tromer, and M. Virza. SNARKs for C: Verifying Program Executions Succinctly and in Zero Knowledge. In *Proceedings CRYPTO*. 2013.
- [Can98] R. Canetti. Security and Composition of Multi-party Cryptographic Protocols. *Journal of Cryptology*, 13:2000, 1998.
- [CFH⁺15] C. Costello, C. Fournet, J. Howell, M. Kohlweiss, B. Kreuter, M. Naehrig, B. Parno, and S. Zahur. Geppetto: Versatile Verifiable Computation. In *Proceedings S&P*, 2015.
- [DFGK14] G. Danezis, C. Fournet, J. Groth, and M. Kohlweiss. Square Span Programs with Applications to Succinct NIZK Arguments. In *Proceedings ASIACRYPT*, 2014.
- [dH12] S. de Hoogh. *Design of large scale applications of secure multiparty computation: secure linear programming*. PhD thesis, Eindhoven University of Technology, 2012.
- [dHSV16] S. de Hoogh, B. Schoenmakers, and M. Veeningen. Certificate Validation in Secure Computation and Its Use in Verifiable Linear Programming. In *Proceedings AFRICACRYPT*, 2016.
- [FFG⁺16] D. Fiore, C. Fournet, E. Ghosh, M. Kohlweiss, O. Ohrimenko, and B. Parno. Hash First, Argue Later: Adaptive Verifiable Computations on Outsourced Data. In *Proceedings CCS*, 2016.
- [GGPR13] R. Gennaro, C. Gentry, B. Parno, and M. Raykova. Quadratic Span Programs and Succinct NIZKs without PCPs. In *Proceedings EUROCRYPT*. 2013.
- [Gro10] J. Groth. Short Pairing-Based Non-interactive Zero-Knowledge Arguments. In *Proceedings ASIACRYPT*, 2010.
- [Gro16] J. Groth. On the Size of Pairing-Based Non-interactive Arguments. In *Proceedings EUROCRYPT*, 2016.

- [Lip16] H. Lipmaa. Prover-Efficient Commit-and-Prove Zero-Knowledge SNARKs. In *Proceedings AFRICACRYPT*, 2016.
- [Par15] B. Parno. A Note on the Unsoundness of vnTinyRAM’s SNARK. Cryptology ePrint Archive, Report 2015/437, 2015. <http://eprint.iacr.org/>.
- [PHGR13] B. Parno, J. Howell, C. Gentry, and M. Raykova. Pinocchio: Nearly Practical Verifiable Computation. In *Proceedings S&P*, 2013.
- [SVdV16] B. Schoenmakers, M. Veeningen, and N. de Vreede. Trinocchio: Privacy-Preserving Outsourcing by Distributed Verifiable Computation. In *Proceedings ACNS*, 2016.
- [Vee17] M. Veeningen. Pinocchio-Based Adaptive zk-SNARKs and Secure/Correct Adaptive Function Evaluation. In *Proceedings AFRICACRYPT*, 2017.

A Security Proof of our Adaptive zk-SNARK

Proof (of Theorem 1). We show each of the properties in turn:

Perfect completeness By inspection.

Argument of knowledge Our proof follows the structure of [PHGR13]. Let \mathcal{A} be a NUPPT algorithm. We construct extractor $\mathcal{E}_{\mathcal{A}}$ as follows. Suppose \mathcal{A} returns a verifying proof, consisting of a sequence of tuples:

$$\langle\langle C_i \rangle_1, \langle \alpha_i C_i \rangle_2, \langle V_i \rangle_1, \langle \alpha_v V_i \rangle_2, \langle W_i \rangle_2, \langle \alpha_w W_i \rangle_1, \langle Y_i \rangle_1, \langle \alpha_y Y_i \rangle_2, \langle Z_i \rangle_1$$

and a proof element $\langle H \rangle_1$.

Extractor $\mathcal{E}_{\mathcal{A}}$ runs the extractor of the commitment family (i.e., the d -PKE extractor) to obtain, for each i , openings \mathbf{v}_i, r_i of commitment $\langle C_i \rangle_1, \langle \alpha_i C_i \rangle_2$. (Indeed, set-up of the extractable commitment scheme family is exactly the set-up of the d -PKE assumption; and key generation of the commitment scheme family and zk-SNARK together form the auxiliary input generator.) Note that, letting $c_i(x) := r_i + \mathbf{v}_{i,1}x + \dots + \mathbf{v}_{i,d}x^d$, we have $\langle C_i \rangle_1 = \langle c_i(x) \rangle_1$.

At the same time, we can use the d -PKE extractor to extract polynomials $v_i(x)$, $w_i(x)$, and $y_i(x)$ of degree at most d such that $\langle V_i \rangle_1 = \langle r_v \cdot v_i(x) \rangle_1$, $\langle W_i \rangle_2 = \langle r_w \cdot w_i(x) \rangle_2$, $\langle Y_i \rangle_1 = \langle r_y \cdot y_i(x) \rangle_1$. (E.g., for $\langle V_i \rangle_1$, the q -PKE set-up generates $\langle r_v \rangle_1$ and random s ; the auxiliary input generator uses this to perform set-up and key generation of the commitment family, and key generation not involving $\langle r_v x^i \rangle_1$; the d -PKE adversary uses $\{\langle r_v x^i \rangle_1, \langle \alpha_v r_v x^i \rangle_2\}$ to generate the appropriate evaluation key elements, runs \mathcal{A} , and returns $\langle V_i \rangle_1, \langle \alpha_v V_i \rangle_2$.)

Next, the extractor computes $r'_v := r_v(s^{d+1})^{-1}$, $r'_w := r_w(s^{2d+2})^{-1}$, $r'_y := r_y(s^{3d+3})^{-1}$, and determines $P_i(x) = c_i(x) + r'_v x^{d+1} v_i(x) + r'_w x^{2d+2} w_i(x) + r'_y x^{3d+3} y_i(x)$. In the QAP solution, let $W+1, \dots, W+N$ be the indices of the witness. Let $p_j(x) = r'_v x^{d+1} v_j(x) + r'_w x^{2d+2} w_j(x) + r'_y x^{3d+3} y_j(x)$. The extractor finds $\mathbf{w}_1, \dots, \mathbf{w}_N, \delta_{v,1}, \delta_{w,1}, \delta_{y,1}$ such that

$$P_1(x) = r_1 + \sum_{j=1}^d \mathbf{v}_{1,j} \cdot (x^j + p_j(x)) + \sum_{j=1}^N \mathbf{w}_j p_{W+j}(x) + \delta_{v,1} r'_v x^{d+1} t + \delta_{w,1} r'_w x^{2d+2} t + \delta_{y,1} r'_y x^{3d+3} t,$$

otherwise it fails. For all other i , the extractor finds $\delta_{v,i}, \delta_{w,i}, \delta_{y,i}$ such that

$$P_i(x) = r_i + \sum_{j=1}^d \mathbf{v}_{i,j} \cdot (x^j + p_{(i-1)d+j}(x)),$$

otherwise it fails. Finally, the extractor returns $\mathbf{v}_1, \dots, \mathbf{v}_n; r_1, \dots, r_n; \mathbf{w}$.

We need to show that, if \mathcal{A} produces a verifying proof, the extractor fails only with negligible probability and the returned values are in \mathcal{R} .

So suppose the extractor fails, i.e., it cannot write polynomial $P_i(x)$ as the linear combination specified above, with non-negligible probability. We construct an adversary \mathcal{B} for the $(4d+3)$ -PDH assumption as follows. \mathcal{B} receives a $(4d+3)$ -PDH challenge, i.e., values $\{\langle x^i \rangle_1, \langle x^i \rangle_2\}_{i=0,1,\dots,4d+3,4d+5,\dots,8d+6}$. \mathcal{B} performs $\mathsf{G0}, \mathsf{Gc}$ as usual using $\{\langle x^i \rangle_1, \langle x^i \rangle_2\}_{i=0,\dots,d}$. For G , rather than picking r_v, r_w at random, it picks $r'_v, r'_w, r'_y = r'_v r'_w$ at random and defines $r_v(x) = r'_v x^{d+1}$, $r_w(x) = r'_w x^{2d+2}$, $r_y(x) = r'_y x^{3d+3}$ implicitly. Instead of generating β_i at random, it generates it as follows. Let \mathcal{P} be the collection of polynomials that P_i should be a linear combination of. That is, for $i = 1$, take polynomials $1, x^j + p_j(x), p_{W+j}(x), r'_v x^{d+1} t(x), r'_w x^{2d+2} t(x),$ and $r'_y x^{3d+3} t(x)$; for other j , take $1, x^j + p_j(x), r'_v x^{d+1} t(x), r'_w x^{2d+2} t(x),$ and $r'_y x^{3d+3} t(x)$. Choose polynomial $\beta'_i(x)$ of degree $\leq 4d+3$ uniformly at random such that, for every $p(x) \in \mathcal{P}$, $\beta'_i(x) \cdot p(x)$ has zero coefficient at $4d+4$. Let $\beta_i = \beta'_i(s)$. Observe that, by construction, the PDH adversary can generate all needed terms without ever knowing s .

Now, \mathcal{B} performs the argument of knowledge experiment with extractor $\mathcal{E}_{\mathcal{A}}$. Using the d -PKE extractor, it obtains polynomial $P_i(x)$ satisfying the (\mathbf{Z}) -check, so $\langle P(x) \cdot \beta'_i(x) \rangle_1 = \langle Z_i \rangle_1$, but with non-negligible probability, $P_i \notin \text{span}(\mathcal{P})$. In looking at how \mathcal{B} selected $\beta'_i(x) = \beta^0 + \beta^1 x + \dots + \beta^{4d+3} x^{4d+3}$, observe that for every non-constant $p \in \mathcal{P}$, the zeroness requirement on the $(4d+4)$ th coefficient of $\beta'_i(x)$ induces a linear constraint on coefficients $\beta^1, \dots, \beta^{4d+3}$. Let B be the vector space of $\beta^1, \dots, \beta^{4d+3}$ satisfying these constraints. A priori this vector space has dimension $4d+3$, and every linearly independent non-constant polynomial in \mathcal{P} reduces this dimension by one. Now consider the subspace $B' \subset B$ of vectors $\beta^1, \dots, \beta^{4d+3}$ for which additionally $P_i(x) \cdot \beta_i(x)$ has zero coefficient at $4d+4$. Since $P_i \notin \text{span}(\mathcal{P})$, this induces a new constraint, so $\dim B' = \dim B - 1$. Now, a random polynomial $\beta'_i(x)$ can be chosen by choosing β^0 uniformly random from \mathbb{F} and $\beta = \beta^1, \dots, \beta^{4d+3}$ uniformly random from B . Note that, from the point of view of the \mathcal{A} , every choice from B is equally likely. This is because \mathcal{A} only sees $\beta_i(s)$ which includes the uniformly random constant term β^0 . So, the chance that $\beta'_i(x)$ lies in the set B' defined by the adversary's $p(x)$ is equal to the chance that any random \mathbf{b} lies in B' , i.e., $1/|B'|$. So, with probability $1 - 1/|\mathbb{F}|$, the coefficient of $P_i(x) \cdot \beta'_i(x)$ at $4d+4$ is nonzero. But in this case, from $\langle Z_j \rangle_1 = \langle P_i(x) \cdot \beta_i(x) \rangle_1$ and $\langle x^j \rangle_1$ ($j \neq 4d+4$), \mathcal{B} can compute $\langle x^{4d+4} \rangle_1$. This contradicts the $(4d+3)$ -PDH assumption, so in fact, the extractor fails only with negligible probability.

It remains to show that the values returned by the extractor are in \mathcal{R} . By the properties of the PKE-extractor we know that $\mathbf{v}_1, \dots, \mathbf{v}_n, r_1, \dots, r_n$ correctly open the given commitments; it remains to show that $\mathbf{x} := (\mathbf{v}_1; \dots; \mathbf{v}_n; \mathbf{w}) \in \mathcal{R}$. Suppose this is not the case, then we build an adversary to the $(8d+6)$ -SDH

assumption as follows. Given a $(8d+6)$ -SDH challenge $\{\langle x^i \rangle_1, \langle x^i \rangle_2\}_{i=0,1,\dots,8d+6}$ it uses this challenge to perform key generation (again generating r_v, r_w, r_y implicitly as polynomials), and with these keys performs the argument of knowledge experiment. Suppose the experiment returns values such that \mathbf{x} is not in \mathcal{R} . Let $V'(x) = \sum_{j=1}^{W+N} \mathbf{x}_i \cdot v_j(x) + (\sum_i \delta_{v,i}) \cdot t(x)$, $W'(x) = \sum_{j=1}^{W+N} \mathbf{x}_i \cdot w_j(x) + (\sum_i \delta_{w,i}) \cdot t(x)$, $Y'(x) = \sum_{j=1}^{W+N} \mathbf{x}_i \cdot y_j(x) + (\sum_i \delta_{y,i}) \cdot t(x)$. Since $\mathbf{x} \notin \mathcal{R}$, by definition of $v_j(x)$, $w_j(x)$, and $y_j(x)$, we have that $t(x)$ does not divide $p(x) = V'(x)W'(x) - Y'(x)$. Let $(x-r)$ be a monomial dividing $t(x)$ but not $p(x)$. The adversary uses the extended Euclidean algorithm to compute $d(x) = \gcd(p(x), t(x))$, $a(x)$ of degree $\leq 2d-1$ and $b(x)$ of degree $\leq d-1$ such that $a(x)t(x) + b(x)p(x) = d(x)$. Multiplying the left-hand side and right-hand side by $1/((x-r)d(x))$ and re-arranging, we have that

$$a(x) \cdot \frac{t(x)}{(x-r)d(x)} + b(x) \cdot \frac{t(x)}{(x-r)d(x)} \cdot \frac{p(x)}{t(x)} = \frac{1}{x-r}.$$

Note that $(x-r)d(x)$ divides $t(x)$, so $t(x)/((x-r)d(x))$ is a polynomial. Hence, the adversary can evaluate $\langle a(x) \cdot \frac{t(x)}{(x-r)d(x)} \rangle_1$ and hence $e(g, g)^{a(s) \cdot \frac{t(s)}{(s-r)d(s)}}$. For the same reason, the adversary can compute $\langle b(x) \cdot \frac{t(x)}{(x-r)d(x)} \rangle_1$. Translated to the present situation, the proof check **(H)** states that:

$$e(\langle r_v V'(x) \rangle_1, \langle r_w W'(x) \rangle_2) = e(\langle H \rangle_1, \langle r_v r_w t \rangle_2) \cdot e(\langle r_v r_w Y'(x) \rangle_1, \langle 1 \rangle_2).$$

By the properties of the pairing, this implies that $g^{r_v V'(s) \cdot r_w W'(s) - r_v r_w Y'(s)} = \langle H \rangle_1^{r_v r_w t(s)}$, so $\langle H \rangle_1 = g^{(V'(s)W'(s) - Y'(s))t(s)^{-1}} = \langle p(x)/t(x) \rangle_1$. Hence, the adversary can also evaluate $e(g, g)^{b(s) \cdot \frac{t(s)}{(s-r)d(s)} \cdot \frac{p(s)}{t(s)}}$. But then, as established above, the adversary can evaluate $e(g, g)^{\frac{1}{s-r}}$. This contradicts the $(8d+6)$ -SDH assumption, so in fact the extractor returns values in \mathcal{R} .

Perfectly composable zero-knowledgeness The existence and correctness of our simulator \mathcal{S} follow the analogous result from [GGPR13].

Suppose the simulator is given commitments $\mathbf{u}_i = (\langle C_i \rangle_1, \langle \alpha_c C_i \rangle_2)$. The simulator generates, for each i , random $\delta_{v,i}, \delta_{w,i}, \delta_{y,i}$, and sets $\langle V_i \rangle_1 = \langle r_v \delta_{v,i} t \rangle_1$, $\langle W_i \rangle_2 = \langle r_w \delta_{w,i} \rangle_2$, $\langle Y_i \rangle_1 = \langle r_y \delta_{y,i} t \rangle_1$, which is statistically identical to the real $\langle V_i \rangle_1, \langle W_i \rangle_2, \langle Y_i \rangle_1$ as long as $t(s) \neq 0$. Note that verification relations **(V)**–**(H)** now fix the remaining proof elements, and moreover, the simulator can generate remaining proof elements satisfying these relations. Namely, it computes $\langle \alpha_v V_i \rangle_2 = \langle \alpha_v r_v \delta_{v,i} t \rangle_2$, $\langle \alpha_w W_i \rangle_1 = \langle \alpha_w r_w \delta_{w,i} \rangle_1$, $\langle \alpha_y Y_i \rangle_2 = \langle \alpha_y r_y \delta_{y,i} t \rangle_2$; $\langle Z_i \rangle_1 = \langle \beta_i \cdot (r_v \delta_{v,i} t + r_w \delta_{w,i} + r_y \delta_{y,i} t) \rangle_1 + \beta_i \cdot \langle C_i \rangle_1$, and $\langle H \rangle_1 = \langle \delta_{v,i} \delta_{w,i} - \delta_{y,i} \rangle_1$. One verifies that these elements indeed satisfy the verification relations. Since $\langle V_i \rangle_1, \langle W_i \rangle_2, \langle Y_i \rangle_1$ are distributed as in the real proof and the other proof elements follow from them (so also have the same distribution), the simulator has produced a proof that is indistinguishable from real, as required. \square

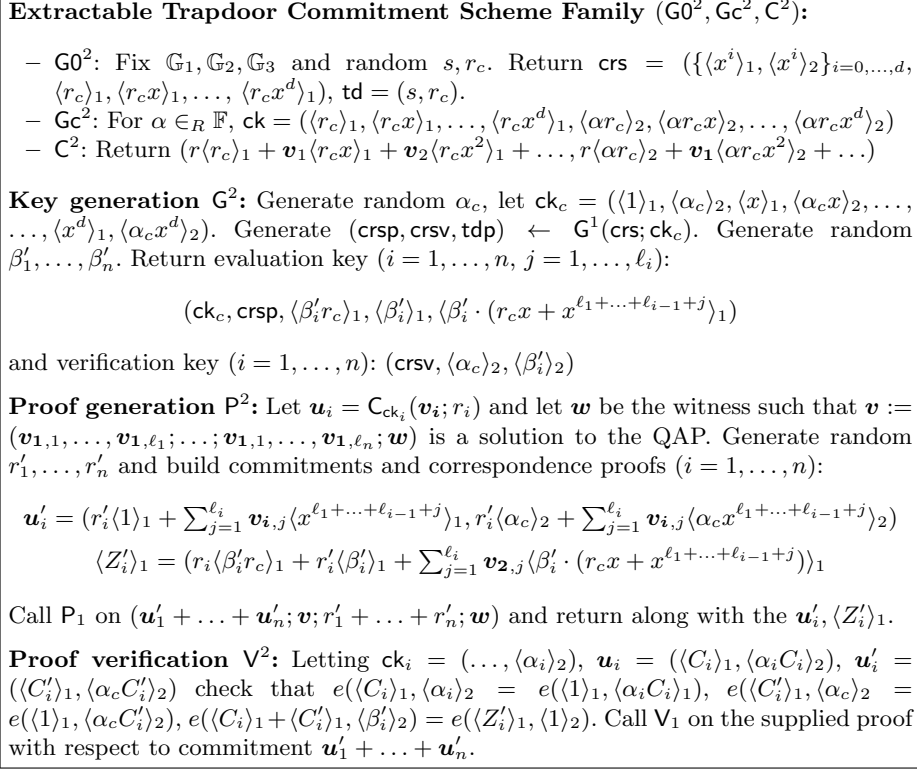


Fig. 3. Pinocchio-Based Adaptive zk-SNARK (G^2, P^2, V^2)

B Details and Proof of Optimized Adaptive zk-SNARK

Figure 3 provides the details of zk-SNARK (G^2, P^2, V^2), the second optimization presented in from Section 4.

Theorem 2. *Under the $(4d+3)$ -PDH, d -PKE, and $(8d+6)$ -SDH assumptions, with d the max degree of the QAPs used, (G^2, P^2, V^2) is an adaptive zk-SNARK.*

Proof. We show each of the properties in turn:

Perfect completeness By inspection.

Argument of knowledge Suppose \mathcal{A} returns commitments $\mathbf{u}_1, \dots, \mathbf{u}_n$, where $\mathbf{u}_i = (\langle C_i \rangle_1, \langle \alpha_i C_i \rangle_2)$, and verifying proof $\pi = (\langle C'_i \rangle_1, \langle \alpha_c C'_i \rangle_2, \langle Z'_i \rangle_1, \pi')$. Extractor $\mathcal{E}_{\mathcal{A}}$ works as follows. Analogously to above, it uses the d -PKE extractor, it extracts openings \mathbf{v}_i, r_i to commitments \mathbf{u}_i (so $\langle C_i \rangle_1 = \langle r_c \cdot c_i(x) \rangle_1$, where $c_i(x) = r_i + \mathbf{v}_{i,1} \cdot x + \dots$). It also extracts openings \mathbf{v}'_i, r'_i to intermediate commitments $\mathbf{u}'_i = (\langle C'_i \rangle_1, \langle \alpha_c C'_i \rangle_2)$ (so $\langle C'_i \rangle_1 = \langle c'_i(x) \rangle_1$, where $c'_i(x) = r'_i + \mathbf{v}_{i,1}' x + \dots$). Using the

extractor of the zk-SNARK from the preceding construction, it further obtains opening \mathbf{v}', r' of the intermediate commitment $\mathbf{u}' := \mathbf{u}'_1 + \dots + \mathbf{u}'_n$ and witness \mathbf{w} . (Note that $\mathsf{G0}^2$ consists of performing setup $\mathsf{G0}^1$ and generating additional CRS elements based on its output.) If any \mathbf{v}'_i is incorrect, i.e., it does not consist of $\ell_1 + \dots + \ell_{i-1}$ zeros, followed by $\mathbf{v}_{i,1}, \dots, \mathbf{v}_{i,\ell_i}$, followed by zeros, it fails. If \mathbf{v}' is not equal to $\mathbf{v}'_1 + \dots + \mathbf{v}'_n$, it also fails. Otherwise it returns $(\mathbf{v}_1, \dots, \mathbf{v}_n; r_1, \dots, r_n; \mathbf{w})$.

We need to show that, if \mathcal{A} produces a verifying proof, the extractor fails only with negligible probability and the returned values are in \mathcal{R} . By the argument of knowledge property of the used zk-SNARK, we know that except with negligible probability, \mathbf{v}', r' correctly open \mathbf{u}' , and $(\mathbf{v}', \mathbf{w})$ are in the relation of the QAP. By the properties of the d -PKE extractor, the returned $\mathbf{v}_1, \dots, \mathbf{v}_n; r_1, \dots, r_n$ are also openings to the commitments $\mathbf{u}_1, \dots, \mathbf{u}_n$. Moreover, if the extractor does not fail, $\mathbf{v}_1, \dots, \mathbf{v}_n$ and $\mathbf{v}'_1 + \dots + \mathbf{v}'_n$ relate to each other in the way intended by the computation, hence the returned values are in \mathcal{R} . So we are only left to show that the extractor fails with negligible probability.

First, suppose the extractor fails because \mathbf{v}' is not equal to $\mathbf{v}'_1 + \dots + \mathbf{v}'_n$. Note that (\mathbf{v}', r') and $(\mathbf{v}'_1 + \dots + \mathbf{v}'_n, r'_1 + \dots + r'_n)$ are two different openings to commitment \mathbf{u}' . But this contradicts the binding property of the intermediate commitment scheme. Indeed, given a $(4d + 3)$ -PDH challenge, we can run the above extractor to obtain, with non-negligible probability, two openings to commitment \mathbf{u}' , from which by the argument of [Gro10] we can compute $\langle x^{4d+3} \rangle_1$, contradicting the $(4d + 3)$ -PDH assumption. The other possibility is that the extractor fails because any \mathbf{v}'_i was incorrectly formed. If this is the case, then we use this to construct an adversary \mathcal{A} to the $(4d + 3)$ -PDH assumption. Given a $(4d + 3)$ -PDH challenge $\{\langle x^i \rangle_1, \langle x^i \rangle_2\}_{i=0,1,\dots,4d+3,4d+5,\dots,8d+6}$, \mathcal{A} runs the above extractor, but in doing so chooses β'_i implicitly as a polynomial such that, for every polynomial $p_j(x) := r'_c x^{d+1+j} + x^{\ell_1+\dots+\ell_{i-1}+j}$, $\beta'_i(x) \cdot p_j(x)$ has zero coefficient at $4d+4$. Since the extractor fails, by assumption the polynomial $P_i(x) = r_c c_i(x) + c'_i(x)$ does not lie in the span of the $\{p_j(x)\}$. Following exactly the same reasoning as in the proof for the argument of knowledge property of our previous construction, with probability $1 - 1/\mathbb{F}$, this implies that $P_i(x) \cdot \beta'_i(x)$ has nonzero coefficient at $4d+4$, enabling \mathcal{A} to compute $\langle x^{4d+4} \rangle_1$ and contradict $(4d + 3)$ -PDH.

Perfectly composable zero-knowledgeness Our simulator samples random r'_i , lets $\langle C'_i \rangle_1 = \langle r'_i \rangle_1$, $\langle \alpha_c C'_i \rangle_2 = \langle \alpha_c r'_i \rangle_2$, and $\langle Z_i \rangle_1 = \beta'_i \cdot (\langle C_i \rangle_1 + \langle C'_i \rangle_1)$, and then uses the simulator for the above zk-SNARK with commitment $(\langle C'_1 \rangle_1 + \dots + \langle C'_n \rangle_1, \langle \alpha_c C'_1 \rangle_2 + \dots + \langle \alpha_c C'_n \rangle_2)$. These proofs clearly have the correct distribution. \square

C Security Model and Proofs for Section 5

We model “secure adaptive function evaluation” and “correct adaptive function evaluation” using the ideal/real paradigm, i.e., by specifying an “ideal world” in which the task at hand is carried out by an incorruptible trusted party,

and demanding that real-world executions give the same result as executions in this ideal world (analogously to [SVdV16]). We then prove that our “Adaptive Trinocchio” construction is secure in this model.

C.1 Execution Model

Our execution model is based on Canetti’s traditional standalone execution model [Can98], of which we now give an overview. We consider a fixed set of parties (in our case: a number of data owners, a client, and a number of workers), and assume static, non-uniform probabilistic polynomial time (NUPPT) adversaries.

A *distribution ensemble* $X = \{X(\lambda, a)\}_{\lambda \in \mathbb{N}, a \in D}$ is an infinite sequence of probability distributions, indexed over a security parameter λ and inputs a from a given domain D .

Definition 6 (Computational indistinguishability ([Can98])). *Let $\delta : \mathbb{N} \rightarrow \{0, 1\}$. We say that distribution ensembles X and Y have computational distance at most δ if for every algorithm \mathcal{D} that is probabilistic polynomial-time in its first input, for all sufficiently large λ , all a and all auxiliary information $w \in \{0, 1\}^*$ we have:*

$$|\text{Prob}(\mathcal{D}(1^\lambda, a, w, x) = 1) - \text{Prob}(\mathcal{D}(1^\lambda, a, w, y) = 1)| < \delta(\lambda),$$

where x is chosen from distribution $X(\lambda, a)$, y is chosen from distribution $Y(\lambda, a)$, and the probabilities are taken over the choices of x , y , and the random choices of \mathcal{D} . If ensembles X and Y have computational distance at most λ^{-c} for all $c > 0$ then we say that X and Y are computationally indistinguishable.

At a high level, a protocol is said to “implement” an ideal functionality if real-life executions of the protocol and ideal-world executions of the ideal functionality are computationally indistinguishable.

In the real-life model, parties P_i perform protocol π together with adversary \mathcal{A} . Parties are NUPPT Turing machines that receive input x_i . The adversary is a NUPPT machine that receives as input the identities of the corrupted parties and their inputs; an auxiliary input; and a value λ for the security parameter. Computation takes place in a fully connected, ideally authenticated, synchronous network with rushing. Moreover, we assume access to a secure bulletin board, i.e., with guaranteed and authenticated delivery. The adversary sees all communication of passively corrupted parties and controls all communication of actively corrupted parties. In the (g_1, \dots, g_k) -hybrid model, parties have access to an incorruptible trusted party that evaluates the functions g_1, \dots, g_k . At the end of the computation, uncorrupted parties output whatever is specified in the protocol, corrupted parties output \perp , and the adversary outputs an arbitrary function over its view. We write $(y_1, \dots, y_m, a) = \text{EXEC}_{\pi, \mathcal{A}}(\lambda; x_1, \dots, x_m; z)$ for the random variable consisting of the outputs y_i of parties P_i and output a of the adversary, and $\text{EXEC}_{\pi, \mathcal{A}}$ for the distribution ensemble $\{\text{EXEC}_{\pi, \mathcal{A}}(\lambda; x_1, \dots, x_m; z)\}_{\lambda \in \mathbb{N}; x_1, \dots, x_m, z \in \{0, 1\}^*}$. When working

Secure adaptive function evaluation	Correct adaptive function evaluation
– Honest data owners send inputs \mathbf{x}_i to trusted party	– Honest data owners send inputs \mathbf{x}_i to trusted party
– Adversary sends inputs \mathbf{x}_i of corrupted data owners to trusted party (active adversary may modify them)	– Adversary sends inputs \mathbf{x}_i of corrupted data owners to trusted party (active adversary may modify them)
– Client sends function f , input \mathbf{x} to trusted party (active: may modify)	– Client sends function f , input \mathbf{x} to trusted party (active: may modify)
– Trusted party computes $\mathbf{y} = f(\mathbf{x}_1, \dots, \mathbf{x}_m, \mathbf{x})$ (where $y_1 = \dots = \perp$ if any $x_i = \perp$)	– Trusted party computes $\mathbf{y} = f(\mathbf{x}_1, \dots, \mathbf{x}_m, \mathbf{x})$ (where $y_1 = \dots = \perp$ if any $x_i = \perp$)
– Trusted party sends \mathbf{y} to client, f to adversary	– Trusted party sends $\mathbf{x}_1, \dots, \mathbf{x}_m, \mathbf{x}, f$ to adversary, receives r
– Client outputs received value (if honest) or \perp (if corrupted); adversary chooses own output	– Trusted party sends \mathbf{y} to client, or \perp if $r = \perp$
	– Client outputs received value (if honest) or \perp (if corrupted); adversary chooses own output

Fig. 4. Ideal-world executions of secure (left) and correct (right) adaptive function evaluation. The highlighted text indicates where the two differ.

in the (g_1, \dots, g_k) -hybrid model, we write $\text{EXEC}_{\pi, \mathcal{A}}^{(g_1, \dots, g_k)}(\lambda; x_1, \dots, x_m; z)$ and $\text{EXEC}_{\pi, \mathcal{A}}^{(g_1, \dots, g_k)}$, respectively.

In the ideal-world model, parties P_i and adversary \mathcal{S} interact with an incorruptible trusted party. As above, P_i and adversary \mathcal{A} are NUPPT machines; P_i have input x_i and the adversary has input the identities of the corrupted parties and their inputs; an auxiliary input; and the security parameter λ . The ideal process then proceeds as indicated by the ideal functionality \mathcal{F} . We write $(y_1, \dots, y_m, a) = \text{IDEAL}_{\mathcal{F}, \mathcal{S}}(\lambda; x_1, \dots, x_m; z)$ for the random variable consisting of the outputs y_i of parties P_i and output a of the adversary, and $\text{IDEAL}_{\mathcal{F}, \mathcal{A}}$ for the distribution ensemble $\{\text{IDEAL}_{\mathcal{F}, \mathcal{A}}(\lambda; x_1, \dots, x_m; z)_{\lambda \in \mathbb{N}; x_1, \dots, x_m, z \in \{0, 1\}^*}\}$.

Definition 7. *Protocol π implements ideal functionality \mathcal{F} (in the (g_1, \dots, g_l) -hybrid model) if for every NUPPT adversary \mathcal{A} there exists a NUPPT simulator \mathcal{S} such that the distribution ensembles $\text{EXEC}_{\pi, \mathcal{A}}$ (resp. $\text{Exec}_{\pi, \mathcal{A}}^{(g_1, \dots, g_l)}$) and $\text{IDEAL}_{\mathcal{F}, \mathcal{S}}$ are computationally indistinguishable.*

C.2 Functionalities: Secure/Correct Adaptive Function Evaluation

“Secure adaptive function evaluation” (Figure 4, left) guarantees privacy *and* correctness; we will realise this ideal functionality if at most a threshold of workers are passively corrupted (but all other parties can be actively corrupted). “Correct adaptive function evaluation” (Figure 4, right) guarantees only correctness; we will realise this ideal functionality regardless of corruptions. In both

models, data owners provide their inputs independently from each other and (unlike Trinocchio) from the function at hand; and corrupted inputters cannot fail selectively based on the input.

Definition 8. *A protocol securely adaptively evaluates function f if it implements the ideal functionality shown in Figure 4. A protocol correctly adaptively evaluates function f if it implements the ideal functionality shown in Figure 4 (right).*

Theorem 3. *If the MPC protocol, trapdoor commitment family, and adaptive zk-SNARK are secure, then “Adaptive Trinocchio” (Figure 2) correctly adaptively evaluates f in the (CommGen, CompGen)-hybrid model. If at most one worker is corrupted, then “Adaptive Trinocchio” securely adaptively evaluates f in the (CommGen, CompGen)-hybrid model.*

C.3 Security Proof: Correct Adaptive Function Evaluation

Figure 5 shows the simulator used to prove that our protocol implements correct function evaluation.

Proof (Theorem 3, “correct” case). We need to show that, for every NUPPT adversary \mathcal{A} , there exists a simulator \mathcal{S} such that, for all inputs $\mathbf{a}_1, \dots, \mathbf{a}_n, \mathbf{a}$ and functions f , real-world execution of the protocol in Figure 2 and ideal-world execution of \mathcal{S} give indistinguishable results. We claim that the simulator in Figure 5 does the job. We show this by, starting from the the ideal-world execution (i.e., the algorithm consisting of the code of the ideal-world trusted party, honest parties, simulator and their orchestration), changing the execution in indistinguishable steps until we arrive at the real-world execution (i.e., the algorithm consisting of the code of the real-world trusted party, honest parties, adversary, and their orchestration):

From Ideal to YAD_1 In step 2, instead of using trapdoor commitments for the honest parties’ inputs, we use the actual inputs given as argument to YAD_1 . In step 4, we use the actual openings of these commitments.

Indistinguishability of Ideal and YAD_1 Follows directly from the trapdoor property of the commitment family.

From YAD_1 to YAD_2 If the computation has succeeded, instead of letting the trusted party evaluate f and using that as the client’s output, use the result \mathbf{b} of the MPC protocol.

Indistinguishability of YAD_1 and YAD_2 The computation only succeeds if the extracted values are in \mathcal{R} and match the inputs $\mathbf{a}_1, \dots, \mathbf{a}_n, \mathbf{a}_c$ to f . By definition of \mathcal{R} this implies that evaluation of f gives \mathbf{b} .

From YAD_2 to YAD_3 In step 6, skip check (2).

Simulator: Adaptive Trinocchio (“correct” case)

(Given NUPPT adversary \mathcal{A} actively corrupting all workers and subset $\mathcal{C} \cap \mathcal{I}$ of inputters \mathcal{I} .)

1. Perform CRS and key generation of the commitment scheme, keeping trapdoors $\text{crs}, \text{ctd}_1, \dots, \text{ctd}_n, \text{ctd}_o$. Simulate step 1 of the protocol with respect to \mathcal{A} by sending this material.
2. Simulate step 2 of the protocol: provide trapdoor commitments c_i for the honest parties, and use the extractor of the commitment scheme family to extract openings (\mathbf{a}_i, r_i) of the provided commitments c_i . Send these \mathbf{a}_i to the trusted party, and get $\mathbf{a}_1, \dots, \mathbf{a}_n, \mathbf{a}_c, f$ back.
3. Simulate step 3 of the protocol: using the trapdoors of the commitment scheme family, compute randomness r_i that opens the honest parties’ trapdoor commitments c_i to the received values \mathbf{a}_i . Simulate a regular execution of the MPC protocol with respect to the adversary.
4. Simulate step 4 of the protocol: handle calls to **CompGen** by performing the key generation of the adaptive zk-SNARK for QAP Q and keeping the trapdoor tdp .
5. Simulate step 5 of the protocol: take the algorithm that simulates a regular execution of the MPC protocol with respect to the adversary and returns the reconstructed computation result $\mathbf{b}, r_c, c_c, \pi$ for the client. Use the argument of knowledge extractor of the adaptive zk-SNARK to extract $\mathbf{a}'_1; \dots; \mathbf{a}'_n; \mathbf{a}'_c, \mathbf{b}'$ and the randomness, proof, and witness.
6. Simulate step 6 of the protocol: send \top to the trusted party if: (1) the received commitment c_c matches the received opening \mathbf{b}, r_c and the received proof verifies; (2) the values \mathbf{a}'_1, \dots and randomness extracted in the previous step match those extracted in step 2, and $\mathbf{a}'_c, \mathbf{b}', r_c$ match the MPC result; (3) the commitments and extracted values are in $\mathcal{R}_{\text{ck}_1, \dots, \text{ck}_n, \text{ck}_c}$; otherwise send \perp .

Fig. 5. Simulator for Adaptive Trinocchio (“correct” case)

Indistinguishability of YAD_2 and YAD_3 The only difference is if checks (1) and (3) are satisfied, but check (2) is not, i.e., there is a mismatch in extracted values. Suppose $\mathbf{a}_i \neq \mathbf{a}'_i$. We know that $c_i = \mathbf{C}_{\text{ck}_i}(\mathbf{a}_i; r_i)$ (because we extracted those values earlier). However, we also know that $(c_1, \dots, c_n, c_c; \mathbf{a}'_1, \dots, \mathbf{w}') \in \mathcal{R}_{\text{ck}_1, \dots, \text{ck}_n, \text{ck}_c}$, so also $c_i = \mathbf{C}_{\text{ck}_i}(\mathbf{a}'_i; r'_i)$. Note that in YAD_3 we are not using the commitment trapdoors anywhere. So YAD_3 can be used to construct an adversary that opens commitments in two different ways, contradicting the computational binding property of the commitment family. This is possible only with negligible probability, proving indistinguishability in this case.

The cases $\mathbf{a}_c \neq \mathbf{a}'_c$ or $\mathbf{b} \neq \mathbf{b}'$ are analogous, except we know the openings not because they were extracted but because they were output by the MPC protocol.

From YAD_3 to Real In step 5, directly execute the MPC protocol instead of using the zk-SNARK knowledge extractor. In step 6, skip check (3). Observe that what is left is simply a regular execution of the protocol.

Simulator: Adaptive Trinocchio (“private” case)

(Given NUPPT adversary \mathcal{A} passively corrupting at most one worker, and actively corrupting a subset of inputters and possibly the client.)

1. Perform CRS and key generation of the commitment scheme, keeping trapdoors $\text{crs}, \text{ctd}_1, \dots, \text{ctd}_n, \text{ctd}_o$. Simulate step 1 of the protocol with respect to \mathcal{A} by sending this material.
2. Simulate step 2 of the protocol: send commitments c_i to zero on behalf of the honest parties, keeping the randomness r_i used. Obtain commitments c_i provided by the adversary on behalf of corrupted inputters.
3. Simulate step 3 of the protocol using the MPC simulator. For the c'_i output on behalf of the honest inputters, output c_i . For the c'_i output on behalf of the corrupted inputters, output $\text{C}_{\text{ck}_i}(\mathbf{a}'_i, r'_i)$, where \mathbf{a}'_i, r'_i are the values input to the MPC protocol by the inputter. Whenever $c_i \neq c'_i$, set $\mathbf{a}'_i = \perp$ and stop after this step. Send \mathbf{a}'_i of corrupted inputters to the trusted party.
4. Simulate step 4 of the protocol: if the client is corrupted, learn f from its call to CompGen ; otherwise, receive f from the ideal-world trusted party. Handle the adversary’s calls to CompGen by performing key generation on this function f .
5. Simulate step 5 of the protocol. If the client is not corrupted, then the adversary does not learn any outcome of the MPC, so no information is needed to perform the simulation. If the client is corrupted, then get value \mathbf{a}'_c input by the client from the simulator. Send f, \mathbf{a}'_c to the trusted party, and receive computation result \mathbf{b} . Generate random r_c and compute $c_c = \text{C}_{\text{ck}_c}(\mathbf{a}'_c, \mathbf{b}; r_c)$. Apply the zero-knowledge simulator of the zk-SNARK on the CRS, commitment keys, c_1, \dots, c_n, c_c , and the trapdoors, to get a proof π . Simulate the remainder of the MPC such that $\mathbf{b}, r_c, c_c, \pi$ is output to the corrupted client.
6. Simulate step 6 of the protocol: let a corrupted client handle its received result.

Fig. 6. Simulator for Adaptive Trinocchio (“private” case)

Indistinguishability of YAD₃ and Real The only difference is if check (1) is satisfied but check (3) is not. The argument of knowledge property of the zk-SNARK states that this probability is negligible. \square

C.4 Security Proof: Private Adaptive Function Evaluation

Proof (Theorem 3, “private” case). Analogously to the above proof for correct function evaluation, we claim that the simulator from Figure 6 gives indistinguishable results to a real protocol execution, and prove this via a sequence of intermediate probability distributions:

From Ideal to YAD₁ In step 2, for c_i instead of using commitments to zero, we use commitments to the actual inputs of the honest parties. (Note that we can do this without changing anything else in the simulator; in particular, in step 3 the MPC simulator uses c_i without requiring its opening.)

Indistinguishability of Ideal and YAD_1 Follows directly from the perfect hiding property of the commitment family.

From YAD_1 to YAD_2 In step 5, instead of using the zk-SNARK simulator, use P on the commitments c_1, \dots, c_n, c_c , inputs $\mathbf{a}_1, \dots, \mathbf{a}_n, \mathbf{a}'_c, \mathbf{b}$, randomness r_1, \dots, r_n, r_c , and witness \mathbf{w} available from the MPC simulator.

Indistinguishability of YAD_1 and YAD_2 Follows directly by the perfectly composable zero-knowledge property of the zk-SNARK.

Indistinguishability of YAD_2 and Real Note that the only remaining difference between the YAD_2 and a real-world execution is that the MPC protocol is simulated in YAD_2 and executed in Real. Both for the MPC execution in step 3 and in step 5, the output of the computation is the same: in step 3 it is the commitments of the values provided by the honest and corrupted inputters; in step 5, it is the computation result and proof based on those values and the client's input. Hence, by the security of the MPC protocol, the two distributions are indistinguishable. \square