Efficient Zero Knowledge for Regular Language

Michael Raymond[1], Gillian Evers[1], Jan Ponti[1], Diya Krishnan[2], and Xiang Fu[1(✉)]

[1] Hofstra University, Hempstead, NY, USA
{mraymond2,gevers1,jponti4}@pride.hofstra.edu, Xiang.Fu@hofstra.edu
[2] Carnegie Mellon University, Pittsburgh, PA, USA
diyak@andrew.cmu.edu

Abstract. A succinct zero knowledge proof for regular language membership, i.e., to prove a secret string behind an encryption (hash) belongs to a regular language is useful, e.g., for asserting that an encrypted email is free of malware. The great challenge in practice is that the regular language used is often huge. We present `zkreg`, a distributed commit-and-prove system that handles such complexity. In `zkreg`, cryptographic operations are encoded using arithmetic circuits, and input acceptance is modeled as a zero knowledge subset problem using Σ-protocols. We introduce a Feedback Commit-and-Prove (FB-CP) scheme, which connects Σ-protocols and the Groth16 system with $O(1)$ proof size and verifier cost. We present a close-to-optimal univariate instantiation of zk-VPD, a zero knowledge variation of the KZG polynomial commitment scheme, based on which an efficient zk-subset protocol is developed. We develop a 2-phase proof scheme to further exploit the locality of Aho-Corasick automata. To demonstrate the performance and scalability of `zkreg`, we prove that all ELF files (encrypted and hashed) in a Linux CentOS 7 are malware free. Applying inner pairing product argument, we obtain an aggregated proof of 1.96 MB which can be verified in 6.5 s.

Keywords: zero knowledge proof · zkSNARK · Aho-Corasick Automata · commit-and-prove · commitment schemes

1 Introduction

Consider how email exchange servers are secured for user safety. A common practice is to run anti-virus software on all incoming emails before they are delivered, however, at the cost of user privacy. Instead, a sender can encrypt her email and provide a succinct zero knowledge (zk)-proof certifying its freedom of malware.[1] Given that most prevalent malware scanners use regular languages for signatures, in this paper, we call it the *zero knowledge regular language membership* (zk-Reg) problem. The prover knows a secret string s and a secret key k. The

[1] In this paper, we call a file malware free if it passes the check of a virus scanner.

H. Duan et al. (Eds.): SecureComm 2023, LNICST 567, pp. 369–394, 2025.
https://doi.org/10.1007/978-3-031-64948-6_19

verifier knows a public finite state machine **A** which captures all allowed/benign strings, and is given $\mathsf{hash}(\mathsf{encrypt}(s,k))$. The prover hopes to convince the verifier with a succinct proof π that $s \in L(\mathbf{A})$. In the context of networking and secure communication, zk-Reg has many applications, e.g., to prove that an encrypted DNS request does not contain any forbidden site, and to show an encrypted packet has already passed firewall rules. The non-zk version of the problem can be applied in software distribution where s does not have to be a secret.

It is well known that all NP statements have a zero knowledge proof [26]. There is also a long line of zk-proof research on more expressive machine models such as von Neumann and RAM machines [8,23]. Here, the great challenge in practice is that the size of **A** is out of reach of most (zk)-SNARK provers. Take ClamAV as one example, ignoring all other signature formats, its standard hexadecimal signature database alone generates an Aho-Corasick automaton (AC-DFA) with 19 million states and over 300 million transitions. To directly encode ClamAV's AC-DFA would reach the limit of distributed provers such as DIZK [50]. Thus to prove even a small email message would be prohibitively expensive.

1.1 Contributions

This paper[2] provides a practical solution to the zk-Reg problem. The following is a summary of our contributions. (1) We present a *2-phase proof* solution which exploits the locality of AC-DFA. The scheme brings improvement of prover performance by an order of magnitude. (2) We design a *Feedback Commit-and-Prove (FB-CP)* scheme to connect Σ-protocols and zk-SNARK systems. The scheme provides an alternative to the $\mathtt{CP}_{\mathtt{link}}$ constructions in [4,15,16] with $O(1)$ proof size and verification time. (3) We develop several cryptographic constructions that improve the state of art: (a) a close-to-optimal univariate instantiation of the zk-VPD commitment scheme [54], which only requires 2-pairings for verification (compared with 5-pairings in [54]). It also improves prover cost at a slight increase of proof size; and (b) a zk-subset proof which further improves the concrete efficiency of the zk-batch bilinear membership proof given in [47] (providing reduction of 2 group (field) elements from 5 group (field) elements in the context of [47]). (4) We provide a full implementation and evaluation of the proposed 2-phase proof scheme. We show a distributed Groth16 prover which is 107× faster than the state of the art [50]. To demonstrate the efficiency of our framework, we prove that all ELFs in a Linux CentOS 7 pass the check of ClamAV in zero knowledge, obtaining a 1.96 MB proof that can be verified in 6.5 s. The source code and experimental data are available. [3]

1.2 Technical Overview

Given the problem size, our overall approach is the Commit-and-Prove scheme [1,4,15,16,18], which combines the benefits of zk-SNARK and Σ-protocols.

[2] An extended version of this paper with the full technical details is available at [44].
[3] https://github.com/xfu2006/zkregex.git.

We exploit the locality of AC-DFA. Let *depth* of a state be its shortest distance from the initial state. We notice that for all of the 2479 ELF (object and executable) files in a Linux CentOS 7, after excluding a small frequently visited set (5.5% of all states), over 34.49% (79.95%) have depth no greater than 10 (20). To exploit the locality, we first arithmetize $\mathbf{A}$ by encoding its states and transitions as elements of a large prime field. Let T be the arithmetization of $\mathbf{A}$. Let $\{T_1, ..., T_n\}$ be the subsets of T with T_i containing the elements bounded by depth i. All of these subsets are made public. The prover runs s over $\mathbf{A}$ and let $\mathbf{S}$ be the set of states/transitions along the acceptance path, and m be its max depth. Then our proof scheme consists of essentially two subset proof: (1) $\mathbf{S} \subseteq T_m$ and (2) $T_m \subseteq T$.

Apparently, $\mathbf{S}$ and T_m need to be represented succinctly and in zero knowledge. In [47], a zk-batch accumulator proof is provided by extending bilinear accumulator [42] using Σ-protocol of product proof to "cancel" the additional terms caused by blinding factors. We adapt it for zk-subset proof in our context, and provide a more efficient protocol in Sect. 3.2 by cutting its product proof.

The second component of `zkreg` is a modified Groth16 system (adapting the ccSnark scheme in [16]). It is needed because arithmetic circuit is more convenient than Σ-protocols in encoding encryption and hash operations. In addition, we need to encode AC-DFA, which has more sophisticated machinery than a standard FSA, for its "fail-edge" semantics. In particular, we need to handle a *support-set* (root-set) problem. Recall that the "allowed" transitions/states are published and encoded in bilinear accumulator based zk-set. We need to extract the standard set (i.e., no duplicates) from the multi-set of states and transitions appearing on the acceptance path, before engaging them in the zk-subset proof. This is solved by taking advantage of a well known result for large prime field that: for a multi-set A, the vanishing polynomial of its support-set is $p_A / \gcd(p_A, p'_A)$ where p'_A is the formal derivative of p_A.

We present a Feedback Commit-and-Prove (FB-CP) scheme to connect the two proof components. The system "enforces" the prover to commit *before* she sees the verifier challenge, thus preventing faking a polynomial witness. FB-CP takes an arithmetic circuit, whose secret witness input wires are divided into k segments. The prover first computes the partial proof of the first segment. It is used as a *commitment* to the inputs, so that the Fiat-Shamir heuristics can be applied and some public input wires are re-computed as the hash of the commitment. Then the rest of witness inputs, intermediate and output wires are computed, and the complete Groth16 proof is generated. We show that FB-CP achieves $O(1)$ proof size and verifier cost. Finally, we design an improved univariate instantiation of zk-VPD [54] to provide zero knowledge for FB-CP.

1.3 Related Work

Since its inception [27] zero knowledge proof has generated not only theoretical interests but also numerous applications, e.g., electronic payment systems [38, 39]), anonymous machine learning [36, 40], and verifiable cloud database [54]. See [49] for a complete survey of the recent progress of the field.

We envision our work as an instance of the Commit-and-Prove (CP) proof systems [1,17,18]. In particular, our k-ccGro16 component is a direct generalization of ccGro16 in LegoSnark [17]. In most applications of CP-Snark, a commitment is used to connect heterogeneous proof systems. In our work, it is also used for fixing witness so that Fiat-Shamir is applied. The FB-CP scheme can be regarded as an addition to the $\mathtt{CP}_{\mathtt{link}}$ constructions in [4,15,16]. It essentially asserts the equivalence of two committed polynomials. Compared with the Pedersen $\mathtt{CP}_{\mathtt{link}}$ component of LegoSnark [16], there is no need to create extra prover/verifier keys to apply the linear subspace scheme [34].

In the KZG paper [32], an almost (but incomplete) zk-set solution is provided. Our solution based on pairing based accumulator [47] offers both complete zk and great concrete performance. We note that the performance of lookup arguments has been improved rapidly [21,24,51]. In particular, the most recent construction [21] offers quasi-linear complexity independent of the size of super-set. Integrating lookup arguments in our framework remains as a future direction.

This work is related to the line of research on zk-proofs for state machines [7,8,23]. There are two concurrent works addressing zk-proof of regular expressions. Zombie [53] extends zkMiddleBox [31], and allows one to reason about encrypted DNS requests, using sum-check based SpartanNIZK [46] as the prover. ZK-regex [37] tackles the same problem via MPC-in-the-head, by providing a 2-stage linear scan algorithm for simulating Thompson NFA. Let m be the size of policy specification in extended regular expression that supports intersection and complement, and m' be the size of its equivalent automaton. Let n be the length of the input string. The prover complexity (excluding the cost of finding witness) of Zombie, ZK-regex, and `zkreg` (our work) are $O(mn)$, $O(m'n)$ and $O(m'\mathsf{log}(m')+n\mathsf{log}^2(n))$,[4] respectively. In the worst case, m' can be 2^m. To curb the state explosion problem, in this paper we focus on AC-DFA which is deterministic, and approximate the rest of ClamAV regex patterns (see Sect. 6.1). In ZK-regex, standard Thompson NFA is used. Zombie handles the richest regex among all three. It encodes regex intersection and complement with constant cost. ZK-regex provides an additional secure-regex component where both the policy and the input string are hidden. Compared with Zombie and ZK-regex, the advantage of `zkreg` is the use of zk-subset proof that separates the encoding of automaton from input string. If we replace the zk-subset proof by a pre-processed lookup argument, the prover complexity of `zkreg` can be further improved to be independent of automaton size. As a result, compared with [37,53], in the specific application context where Aho-Corasick is applicable (state explosion avoided), our technique can handle a much larger signature set (policy collection). It is an interesting question if Zombie's technique of encoding regex complement can be integrated with ours.

[4] More precisely, `zkreg` prover cost consists of $O(m'\mathsf{log}(m')+n\mathsf{log}^2(n))$ field operations and $O(m'+n)$ group operations.

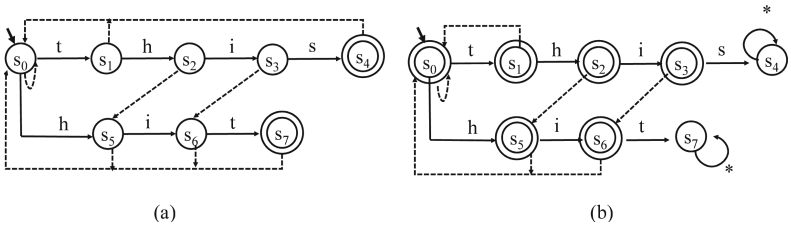

Fig. 1. AC-DFA

2 Preliminaries

2.1 AC-DFA

An Aho-Corasick automaton (AC-DFA) [2] is a deterministic finite state machine built from a set of strings as the virus signature database. Let $V = \{v_0, ..., v_k\}$ and let $m = \sum_{i=0}^{k} |v_i|$. The AC-DFA for V can be constructed in $O(m)$ time. Running a string s over the AC-DFA discovers all contained virus patterns with no more than $2|s|$ transitions. This is achieved by dividing the AC-DFA's transition function into two parts: a regular forward edge relation and a failure edge function. Each state has up to one failure edge, which points to a state that holds the longest suffix of the current input string. Figure 1 shows an example of the AC-DFA for $V = \{\texttt{this}, \texttt{hit}\}$. One can see that the acceptance run of an input string "`thit`" travels through state sequence $s_0, s_1, s_2, s_3, s_6, s_7$ with the transition $s_3 \rightarrow s_6$ as a failure link. Negation of an AC-DFA for a virus signature set, e.g., as shown in Fig. 1(b), can capture the set of "benign" strings.

2.2 Notations and Security Assumptions

Let λ be the security parameter. We denote negligible in λ as $\epsilon(\lambda)$. We write $f = \epsilon(\lambda)$ as $f \approx 0$, and $|f - g| = \epsilon(\lambda)$ as $f \approx g$. Let $\mathcal{G}$ be a generator of bilinear groups, i.e., $(p, \mathbf{g}_1, \mathbf{g}_2, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e) \leftarrow \mathcal{G}(1^\lambda)$. Here $\mathbb{G}_1$, $\mathbb{G}_2$, and $\mathbb{G}_T$ all have prime order p, with $\mathbf{g}_1$ ($\mathbf{g}_2$) as the generator of $\mathbb{G}_1$ ($\mathbb{G}_2$). $e : \mathbb{G}_1 \times \mathbb{G}_2 \rightarrow \mathbb{G}_T$ is the bilinear map s.t. for any $a, b \in \mathbb{Z}_p$: $e({\mathbf{g}_1}^a, {\mathbf{g}_2}^b) = e(\mathbf{g}_1, \mathbf{g}_2)^{ab}$ and $e(\mathbf{g}_1, \mathbf{g}_2)$ is the generator of $\mathbb{G}_T$. Given a field $\mathbb{F}$, $a \xleftarrow{\$} \mathbb{F}$ means to sample a from $\mathbb{F}$ uniformly.

Following Groth16 [30], we write $\mathbb{G}_1$ and $\mathbb{G}_2$ as additive groups. Given $a \in \mathbb{Z}_p$, we denote ${\mathbf{g}_1}^a$ as $[a]_1$, and similar are $\mathbb{G}_2$ and $\mathbb{G}_T$. For instance, ${\mathbf{g}_1}^a {\mathbf{g}_1}^b$ is written as $[a]_1 + [b]_1$ or $[a+b]_1$, $({\mathbf{g}_2}^a)^b$ as $[ab]_2$, and $e({\mathbf{g}_1}^a, {\mathbf{g}_2}^b)$ is denoted as $[a]_1 \cdot [b]_2$ or $[ab]_T$. We use $\vec{a} \in \mathbb{Z}_p^n$ to define a vector of n field elements: $(\vec{a}_0, \ldots, \vec{a}_{n-1})$, where $\vec{a}_0$ is the first element of $\vec{a}$. Similarly, $[\vec{a}]_1$ represents a vector of $\mathbb{G}_1$ elements $([\vec{a}_0]_1, \ldots, [\vec{a}_{n-1}]_1)$. Given $\vec{b} \in \mathbb{Z}_p^n$, $\vec{b}[\vec{a}]_1$ is a vector $([\vec{a}_0\vec{b}_0]_1, \ldots, [\vec{a}_{n-1}\vec{b}_{n-1}]_1)$, and dot product $[\vec{a}]_1 \cdot [\vec{b}]_2$ is defined as $([\vec{a}_0\vec{b}_0]_T, \ldots, [\vec{a}_{n-1}\vec{b}_{n-1}]_T)$. We may also use $\boldsymbol{a}[i]$ to indicate its i'th element (index from 0). For a two dimensional array $\vec{b}$, $\vec{b}[i]_j$ denotes the element at row i column j. Given a multi-set S, its *support set* is $\hat{S} = \{x \mid x \in S\}$. For instance, given $T = \{1, 2, 3, 2, 3\}$, $\hat{T} = \{1, 2, 3\}$.

Our system is based on a number of security assumptions: discrete logarithm assumption (DL) [32], q-Strong Diffie-Hellman (q-SDH) [32], q-Power Knowledge of Exponent (q-PKE) [29], and q-computational power Diffie-Hellman (q-CPDH) [29]. These are frequently used assumptions in commitment schemes and zk-SNARK. Their definitions are given in Appendix A of the full version of this paper [44]. We refer readers to [29,30,32] for the discussion on their computational hardness.

2.3 Σ-Protocols and Zk-SNARK

Given an NP relation R, we say that $(\mathcal{P}, \mathcal{V})$ is an interactive proof system (Σ-protocol) if prover $\mathcal{P}$ demonstrates the knowledge of $(x, w) \in R$ to verifier $\mathcal{V}$, disclosing zero knowledge about the witness w. The Σ-protocols we present in this paper are perfect complete, knowledge sound, and honest verifier zero knowledge (HVZK). Its formal definition is given in [44, Appendix B]. A Σ-protocol can be converted to non-interactive using the Fiat-Shamir heuristic under the random oracle model [22].

We use an adapted notation from [14] to specify zk-protocols. Consider Schnorr's DLOG protocol as an example: $\Sigma_{\mathsf{DLOG}}(\mathbf{h})\{(x) : \mathbf{h} = [x]_1\}$. Here $\mathtt{DLOG}$ in Σ_{DLOG} is a mnemonic. $(\mathbf{h})$ is the public information. The tuple before ":" is the secret known by prover only, i.e., (x). Then the statement inside curly braces states the relation: the prover knows the secret discrete logarithm of $\mathbf{h}$. Applying Fiat-Shamir heuristics to $\Sigma_{\mathsf{DLOG}}(\mathbf{h})$, we get a non-interactive proof and denote it as $\pi_{\mathsf{DLOG}}(\mathbf{h})$. Schnorr proof can be generalized to multiple bases. We use $1/0 \leftarrow \mathtt{CheckDLOG}(\mathbf{C}_x, \pi_x, (\mathbf{g}_0, \ldots, \mathbf{g}_k))$ to denote the verifier function for π_{DLOG}: the prover knows $\vec{x}_0, \vec{x}_2, \ldots, \vec{x}_k$ so that $\mathbf{C}_x = \vec{x}_0\mathbf{g}_0 + \vec{x}_1\mathbf{g}_1 + ... + \vec{x}_k\mathbf{g}_k$.

zkSNARK (Zero Knowledge Succinct Non-interactive ARgument of Knowledge) systems (e.g., [25,30,46]) provide generic specification for arbitrary relation, and succinct proof size and verification time. Recently, many bilinear group friendly encryption and hash algorithms are developed, e.g., Poseidon [28], and MiMC [3]. In our system, we encode the AC-DFA as an arithmetic circuit. This circuit is then converted to a Rank-1 Constraint System (R1CS), and then to a Quadratic Arithmetic Program (QAP). The QAP is then fed to a modified Groth16 system [30] for proof generation.

2.4 Commitment Schemes

A commitment is a cryptographic primitive that allows one to commit to a secret message and later to open it.

Pedersen Vector Commitment. Let $\mathbf{g}, \mathbf{h} \in \mathbb{G}$ and $\mathsf{log}_{\mathbf{g}}(\mathbf{h})$ is unknown to the prover. Given $s \in \mathbb{Z}_p$ and r sampled from $\mathbb{Z}_p^*$, a Pedersen commitment [43] to s is defined as $\mathtt{CommitPed}(s, r) = s\mathbf{g} + r\mathbf{h}$. Given $\vec{\mathbf{g}} \in \mathbb{G}^{n+1}$ where no linear relation is known for $\vec{\mathbf{g}}$, a Pedersen vector commitment [12,47] to $\vec{a} = (a_0, \ldots, a_{n-1})$, using opening r, is $\mathtt{CommitPed}(\vec{a}, r) = \sum_{i=0}^{n-1} \vec{a}_i\vec{\mathbf{g}}_i + r\vec{\mathbf{g}}_n$. Pedersen commitment is perfect hiding and computational binding.

Polynomial Commitment. Given key $\left(\left([s^i]_1\right)_{i=0}^{q}, \left([\alpha s^i]_1\right)_{i=0}^{q}\right)$, the KZG commitment [32] $\mathbf{C}_{p,1} \in \mathbb{G}_1$ to a polynomial $p(X) \in \mathbb{Z}_p[X]$ is defined as $[p(s)]_1$. Let $p(X) = \sum_{i=0}^{d} a_i X^i$. $\mathbf{C}_{p,1}$ is then $\sum_{i=0}^{d} a_i [s^i]_1$, where each $[s^i]_1$ is from the prover key. One can provide $\mathbf{C}_{p,2} = [\alpha p(s)]_1$ as the *proof of knowledge* for $\mathbf{C}_{p,1}$. By the q-PKE assumption (Definition 6 in [44, Appendix A]), the following bilinear pairing check convinces the verifier that the prover knows all coefficients of the hiding $p(X)$: $\mathbf{C}_{p,1} \cdot [\alpha]_2 = \mathbf{C}_{p,2} \cdot [1]_2$. Given a point t, and let $y = p(t)$. There exist a polynomial $w(X) = (p(X) - y)/(X - t)$. Then based on the q-SDH assumption (Definition 5 in [44, Appendix A]), the following check proves that the $p(X)$ behind $\mathbf{C}_{p,1}$ evaluates to y at point t: $(\mathbf{C}_{p,1} - [y]_1) \cdot [1]_2 = [w(s)]_1 \cdot [s - t]_2$.

zk-VPD Scheme. The KZG commitment scheme is not hiding and it also leaks information of a point evaluation. In [54], a zero knowledge ℓ-variate polynomial delegation scheme (zk-VPD) is presented to address the problem. The idea is to hide the polynomial evaluation behind a Pedersen commitment, thus retaining zero knowledge. Concretely, a zk-VPD commitment to a polynomial $p(X)$ is a pair $(\mathbf{C}_{p,1}, \mathbf{C}_{p,2})$ where $\mathbf{C}_{p,2}$ is the proof of knowledge for $\mathbf{C}_{p,1}$, and $\mathbf{C}_{p,1}$ is an extended KZG commitment to $p(X)$ blinded by random factor. Its opening operation: $(\mathbf{C}_y, \pi) \leftarrow \texttt{Open}(p, r_p, t, \sigma_\Sigma)$, produces a proof which asserts that the secret $p(X)$ behind $\mathbf{C}_{p,1}$ evaluates to a secret value y at t and $\mathbf{C}_y$ is a Pedersen commitment to y. The zk-VPD scheme is complete, binding and zero knowledge. We provide its formal definition for univariate polynomials in Appendix A.

3 Σ-Protocols

We provide two Σ-protocols in this section. The first shows a committed polynomial evaluates to a secret value at a given point in zk, and the second proves a subset relation between two zero knowledge sets. We then apply the inner product arguments [13] to both protocols for proof aggregation.

3.1 ZK-Proof for Polynomial Evaluation

The zk-VPD scheme [54] provides a zero-knowledge solution for showing that an ℓ-variate polynomial behind a commitment evaluates to a secret value at a public point.[5] In this section, we provide a drop-in replacement for the construction provided in [54] for univariate polynomials. It provides lower concrete cost (2 pairings vs 5 pairings for verifier work and saving half of the prover cost). In Appendix A, we provide the original univariate zk-VPD construction as a baseline for comparison. The verifier cost of our construction is close to optimal, considering that the standard (non-zk and non-hiding) KZG commitment scheme's evaluation proof costs 2 pairings at verifier.

[5] Many hiding variations of KZG, e.g., [32, Section 4.2] and Marlin [19, Appendix B.2], still leak evaluation values, which needs bounded-zk technique [15,19]. In our case we need the evaluation to be zk, and hence adapting zk-VPD is the best fit.

1 **Trusted Set-up:** $\sigma_\Sigma \leftarrow \texttt{Setup}(1^\lambda, G, q)$
Parse G as $(p, \mathbf{g}_1, \mathbf{g}_2, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e)$. Sample s_1, s_2, α from $\mathbb{Z}_p^*$. Compute $\gamma \leftarrow \left(\left((s_1)^i\right)_{i=0}^{q}, \left(\alpha(s_1)^i\right)_{i=0}^{q}, s_2, \alpha s_2, s_1 s_2\right)$.
Let prover-verifier key $\sigma_\Sigma = ([\gamma]_1, [\gamma]_2)$. Return σ_Σ.

2 **Commit to Polynomial:** $(\mathbf{C}_p, \pi_p) \leftarrow \texttt{CommitPoly}(p(X), r_p, \sigma_\Sigma)$: Return $([p(s_1) + r_p s_2]_1, [\alpha(p(s_1) + r_p s_2)]_1)$.

3 **Validate Polynomial Commitment:** $1/0 \leftarrow \texttt{Check}(\mathbf{C}_p, \pi_p, \sigma_\Sigma)$: Return 1 if and only if $\mathbf{C}_p \cdot [\alpha]_2 = \pi_p \cdot [1]_2$.

4 **Zk-Prove Polynomial Evaluation:** $(\mathbf{C}_y, \pi) \leftarrow \texttt{Open}(p(X), r_p, t, \sigma_\Sigma)$:
Compute: $y \leftarrow p(t)$. $q_1(X) \leftarrow (p(X) - y)/(X - t)$.
Sample r_y, r_q from $\mathbb{Z}_p^*$. Compute: $\mathbf{C}_y \leftarrow [y + r_y s_2]_1$, $\mathbf{C}_1 \leftarrow [q_1(s_1) + r_q s_2]_1$, and $\mathbf{C}_2 \leftarrow [s_2((s_1 - t)r_q + (r_y - r_p))]_1$.
Compute Schnorr proofs:
$\pi_{\texttt{DLOG}y} \leftarrow \pi_{\texttt{DLOG}}(\mathbf{C}_y, ([1]_1, [s_2]_1))\{(y, r_y) : \mathbf{C}_y = y[1]_1 + r_y[s_2]_1\}$, and
$\pi_{\texttt{DLOG}2} \leftarrow \pi_{\texttt{DLOG}}(\mathbf{C}_2, ([s_2(s_1 - t)]_1, [s_2]_1))\{(r_1, r_2) : \mathbf{C}_2 = r_1[s_2(s_1 - t)]_1 + r_2[s_2]_1\}$.
Let $\pi = (\mathbf{C}_1, \mathbf{C}_2, \pi_{\texttt{DLOG}y}, \pi_{\texttt{DLOG}2})$. Return $(\mathbf{C}_y, \pi)$.

5 **Check Polynomial Evaluation:** $1/0 \leftarrow \texttt{Verify}(\mathbf{C}_p, \mathbf{C}_y, t, \pi, \sigma_\Sigma)$:
Assumption: $\exists \pi_p$ for $\mathbf{C}_p$ s.t. $\texttt{Check}(\mathbf{C}_p, \pi_p, \sigma_\Sigma) = 1$.
Parse π as $(\mathbf{C}_1, \mathbf{C}_2, \pi_{\texttt{DLOG}y}, \pi_{\texttt{DLOG}2})$. Return 1 if and only if:
$\texttt{CheckDLOG}(\mathbf{C}_y, \pi_{\texttt{DLOG}y}, ([1]_1, [s_2]_1)) = 1 \ \wedge \ \texttt{CheckDLOG}(\mathbf{C}_2, \pi_{\texttt{DLOG}2}, ([s_2(s_1 - t)]_1, [s_2]_1)) = 1 \ \wedge \ (\mathbf{C}_p - \mathbf{C}_y + \mathbf{C}_2) \cdot [1]_2 = \mathbf{C}_1 \cdot [s_1 - t]_2$.

Fig. 2. $\Sigma_{\texttt{univar\_zk\_vpd}}$: Univariate zk-VPD Scheme

Figure 2 presents the $\Sigma_{\texttt{univar\_zk\_vpd}}$ construction. For simplicity, we do not distinguish between prover and verifier keys. Its `CommitPoly` and `Check` algorithms are the same as zk-VPD [54]. We thus focus on the zero knowledge polynomial evaluation proof (`Open` and `Verify`). We briefly describe its design idea and compare it with the construction presented in [54].

Recall that in the non-hiding KZG commitment scheme, the evaluation proof is built upon the following observation. If $y = p(t)$ then there exists a polynomial $q_1(X)$ s.t. for any $u \in \mathbb{Z}_p$: $p(u) - y = q_1(u)(u - t)$. This is tested using an equation of two pairings: $([p(s_1)]_1 - [y]_1) \cdot [1]_2 = [q_1(s_1)]_1 \cdot [s_1 - t]_2$. Following the zk-VPD construction [54] all KZG commitments of polynomials are blinded in the `Open`() operation in Fig. 2: $[p(s_1)]_1$ is mapped to $\mathbf{C}_p$, $[q_1(s_1)]_1$ mapped to $\mathbf{C}_1$. Then another commitment $\mathbf{C}_2$ is introduced to "balance off" the terms introduced by the blinding factors. This is verified by the third equality check in the `Verify`() operation. The first two equations in `Verify`() perform the standard Schnorr DLOG verification, which demonstrates the prover's knowledge of multi-base discrete logarithm of $\mathbf{C}_2$ and $\mathbf{C}_y$.

In [54], all zk-VPD commitments come with a proof of knowledge. We observed that: $\mathbf{C}_1$ (for $q_1(X)$) can be used without a proof of knowledge. Then with a slight tweak of the formula of $\mathbf{C}_2$, we can cut from 5 pairings needed at the verifier side to only 2 pairings. In addition, because no proof of knowledge is needed for $\mathbf{C}_1$, the prover cost can be cut in half, because there is no need

1 **Prove Product Relation:** $(\mathbf{C}'_q, \pi) =$
$\texttt{PrvSubset}(\mathbf{C}_p, \mathbf{C}_q, \pi_q, [w(s_1)]_1, [\alpha w(s_1)]_1, [w(s_1)]_2, r_p, r_q, \sigma_\Sigma, \texttt{bMutate})$:
$\#\ p(X) = q(X)w(X)$. r_p is the opening of $\mathbf{C}_p$, and r_q is opening of $\mathbf{C}_q$
Sample r_w from $\mathbb{Z}_p^*$. If bMutate sample r'_q from $\mathbb{Z}_p^*$ otherwise $r'_q \leftarrow 0$. Let
$\mathbf{C}'_q \leftarrow \mathbf{C}_q + r'_q[s_2]_1$ and $\pi'_q \leftarrow \pi_q + r'_q[\alpha s_2]_1$. $r''_q \leftarrow r_q + r'_q$.
Compute: $\mathbf{C}_w \leftarrow [w(s_1)]_1 + r_w[s_2]_1$ and $\pi_w \leftarrow [\alpha w(s_1)]_1 + r_w[\alpha s_2]_1$.
$\mathbf{C}_{w,2} \leftarrow [w(s_1)]_2 + r_w[s_2]_2$. Note $\mathbf{C}_w \in \mathbb{G}_1$ and $\mathbf{C}_{w,2} \in \mathbb{G}_2$.
Compute: $\mathbf{C}_1 \leftarrow r_w\mathbf{C}'_q + r''_q\mathbf{C}_w - r''_q r_w[s_2]_1 - [r_p]_1$.
$\pi_1 \leftarrow \pi_{\texttt{DLOG}}(\mathbf{C}_1, (\mathbf{C}'_q, \mathbf{C}_w, [s_2]_1, [1]_1))\{(r_1, r_2, r_3, r_4) : \mathbf{C}_1 = r_1\mathbf{C}_{q'} + r_2\mathbf{C}_w + r_3[s_2]_1 + r_4[1]_1\}$.
Let $\pi = (\pi'_q, \mathbf{C}_1, \pi_1, \mathbf{C}_w, \pi_w, \mathbf{C}_{w,2})$. Return $(\mathbf{C}_{q'}, \pi)$.

2 **Check Subset Relation:** $1/0 \leftarrow \texttt{VerSubset}(\mathbf{C}_p, \mathbf{C}_q, \pi, \sigma_\Sigma)$:
Assumption: $\exists \pi_p$ s.t. $\mathbf{C}_p \cdot [\alpha]_2 = \pi_p \cdot [1]_2$.
Parse π as $(\pi_q, \mathbf{C}_1, \pi_1, \mathbf{C}_w, \pi_w, \mathbf{C}_{w,2})$. Return 1 if and only if
$\texttt{CheckDLOG}(\mathbf{C}_1, \pi_1, (\mathbf{C}_q, \mathbf{C}_w, [s_2]_1, [1]_1)) = 1 \ \wedge\ \mathbf{C}_w \cdot [\alpha]_2 = \pi_w \cdot [1]_2 \ \wedge\ \mathbf{C}_q \cdot [\alpha]_2 = \pi_q \cdot [1]_2 \ \wedge\ \mathbf{C}_w \cdot [1]_2 = [1]_1 \cdot \mathbf{C}_{w,2} \ \wedge\ \mathbf{C}_p \cdot [1]_2 + \mathbf{C}_1 \cdot [s_2]_2 = \mathbf{C}_q \cdot \mathbf{C}_{w,2}$.

Fig. 3. $\Sigma_{\texttt{subset}}$: Zero Knowledge Subset Protocol

to generate $[\alpha q_1(s_1)]_1$. In Appendix B, we present the proof of our improved construction. The intuition of the proof is that even though an adversary can try to submit $\mathbf{C}_1$ with a "relaxed" requirement of no proof of knowledge needed, creating a fake zk-evaluation proof still breaks the DL and q-SDH assumptions.

Lemma 1. *Under the DL, q-PKE, and q-SDH assumptions, the $\Sigma_{univar\_zk\_vpd}$ construction in Fig. 2 is a univariate zk-VPD scheme defined in [44, Definition 9].*

Efficiency: As each extended Schnorr proof has two bases, each sends 1 $\mathbb{G}_1$ and 2 $\mathbb{Z}_p$ elements. In total, the zk-polynomial evaluation proof costs 4 $\mathbb{G}_1$ and 4 $\mathbb{Z}_p$, and verifier spends 2 pairings plus 6 group operations over $\mathbb{G}_1$.

3.2 ZK Subset Proof

Accumulators such as RSA [6,10] and bilinear pairing based [42] can compress a large set of elements into one succinct representation and provide membership or subset proofs. In this paper, we extend the bilinear accumulator [42,47] for representing (multi)-set of transitions and states of an AC-DFA. We extend and present a more efficient construction than the recent work of zero knowledge batch membership for bilinear accumulator [47]. We are able to cut a Σ-product proof from the protocol of [47], resulting in a reduction of 2 group elements and 2 field elements from 5 group and 5 field elements in proof, had our technique applied in the context of [47].

Note that different from [47] (where the algebraic group model (AGM) is assumed), we assume the plain model in this paper due to recent discussions [33,52] over AGM.

Definition 1. *Given a bilinear group* $(p, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e)$. *For a multi-set* $A = \{a_1, ..., a_n\} \in \mathbb{F}^n$, *define* $p_A(X) = \prod_{i=1}^{n}(X - a_i)$ *be its* vanishing polynomial. *Given a key* $\left(\left([(s_1)^i]_1\right)_{i=0}^{q}, [s_2]_1\right)$, *the bilinear accumulator of* A *is denoted as* $\mathbf{acc}_A = [p_A(s_1)]_1$. *Let* $r \xleftarrow{\$} \mathbb{Z}_p^*$ *be the opening,* $\mathbf{zkset}_A = \mathbf{acc}_A + r[s_2]_1$.

Zero Knowledge Subset Proof. Given $p(X) = q(X)w(X)$ and $\mathbf{acc}_p = [p(s_1)]_1$, $\mathbf{C}_q = [q(s_1)]_1 + r_q[s_2]_1$, $\mathbf{C}_w = [w(s_1)]_1 + r_w[s_2]_1$, and $\mathbf{C}_{w,2} = [w(s_1)]_2 + r_w[s_2]_2$. The basic idea of [47, Section 7.1] is to build a proof for Equation 1 which is eventually reduced to: $[p(s_1)]_1 \cdot [1]_2 = [q(s_1)]_1 \cdot [w(s_1)]_2$ that asserts $p(X) = q(X)w(X)$.

$$\mathbf{acc}_p \cdot [1]_2 + \mathbf{C}_1 \cdot [s_2]_2 = \mathbf{C}_q \cdot \mathbf{C}_{w,2} \tag{1}$$

Here $\mathbf{C}_1$ is introduced to balance off the extra cross-terms caused by the openings of $\mathbf{C}_q$ and $\mathbf{C}_w$. Let $r_1 = r_q$, $r_2 = r_w$, and $r_3 = -r_q r_w$. One can verify that: $\mathbf{C}_1 = r_1\mathbf{C}_w + r_2\mathbf{C}_q + r_3[s_2]_1$.

In [47], Schnorr style zk-proofs are used to establish prover's knowledge of random exponents r_1, r_2, and r_3. In particular, the product relation: $r_3 = -r_1 r_2$. Our observation is that: the zk-proof for the product relation is not needed, given the proof of knowledge for $\mathbf{C}_q$ and $\mathbf{C}_w$.

The `PrvSubset()` operation in Fig. 3 provides the details. First, the $\mathbf{acc}_p$ in [47] is converted to a fully hiding zk-VPD commitment $\mathbf{C}_p$ in our context. The major difference from the proof in [47] is that we just need to provide a DLOG proof for proving the knowledge of the four exponents of $\mathbf{C}_1$, thus, cutting the Σ-product proof used in [47]. Formally, the `VerSubset()` algorithm can be regarded as a Σ-protocol, and we denote it as $\Sigma_{\texttt{subset}}$. It proves that the polynomial hiding behind $\mathbf{C}_q$ is a factor of the polynomial behind $\mathbf{C}_p$. When `bMutate` is enabled, we generate a new proof by mutating an existing one.

Lemma 2. *Under the DL, q-PKE, q-SDH, and q-CPDH assumptions,* $\Sigma_{\texttt{subset}}$ *is perfectly complete, computational sound, and HVZK.*

Appendix C presents the proof for Lemma 2. The following Lemma allows to use the polynomial product relation for proving subsets.

Lemma 3. *Let* $p(X)$ *be the vanishing polynomial for a multi-set* S *of elements in* $\mathbb{Z}_p$, *and let* $\boldsymbol{C}_p$ *be a valid commitment to* $p(X)$ *with proof of knowledge. For any* $\boldsymbol{C}_q$ *if there exists* $\pi_{q\subseteq p}$ *s.t.* $\texttt{VerSubset}(\boldsymbol{C}_p, \boldsymbol{C}_q, \pi_{q\subseteq p}, \sigma_\Sigma) = 1$, *then the polynomial behind* $\boldsymbol{C}_q$ *vanishes at a subset of* S.

Efficiency: The prover has to perform $O(n\log(n))$ field operations (for computing $w(X)$), and $O(n)$ group operations, where n is the degree of $p(X)$ behind $\mathbf{C}_p$. The proof consists of 5 $\mathbb{G}_1$, 1 $\mathbb{G}_2$ and 4 $\mathbb{Z}_p$ elements. The verifier spends 9 pairings and 5 $\mathbb{G}_1$ multiplications. Given n zk-subset claims and proofs, it is possible to aggregate them into one single proof of $\log(n)$ size, using inner pairing product argument [13]. Details are show in [44, Appendix G].

4 Arithmetic Circuit

The arithmetic circuit in `zkreg` takes a *secret* witness stream of input characters (4-bit nibbles) and its acceptance path by the AC-DFA. It enforces the AC-DFA semantics (e.g., the fail-edges), performs encryption and hash operations, generates the vanishing polynomials for states and transitions, and evaluates them at a given random point. We introduce several design decisions that optimize the prover performance.

An AC-DFA is a tuple (Σ, S, s_0, F, T) where $\Sigma = \{i\}_{i=0}^{|\Sigma|-1}$ is the alphabet, $S = \{i\}_{i=0}^{|S|-1}$ and $F = \{i\}_{i=0}^{|F|-1}$ are the set of states and final states. $s_0 \in S$ (its encoded value being 0) is the initial state and each transition in T is a tuple (s, c, t, b) where $s, t \in S$ are the source/destination states, $c \in \Sigma$ is the input character, and b is a Boolean flag indicating if the transition is a fail-edge. Given $\mathbb{Z}_p$ of a bilinear group (usually $p > 2^{252}$), we define an encoding function $\rho : T \to \mathbb{Z}_p$, letting $m = \lceil \mathsf{log}(|S|) \rceil$ and $n = \lceil \mathsf{log}(|\Sigma|) \rceil$:

$$\rho((s, c, t, b)) = b + 2c + 2^{n+1}s + 2^{m+n+1}t + 2^{2\,m+n+2} \tag{2}$$

In practice for ClamAV, we take $m = 25$ and $n = 4$. Note that the circuit has to perform a range check on all elements of a transition, e.g., b is really a Boolean flag. The last item 2^{2m+n+2} in Eq. 2 is used to separate the sets of transitions and states, thus saving the range proof cost for states.

4.1 Support Set

Let U denote the set of transitions along an acceptance path. It is a multi-set, i.e., one element may appear multiple times. When the prover needs to argue in the Σ-protocol that it is a subset of all possible transitions of the AC-DFA, its *support set* $\hat{U}$ (i.e., the set of the distinct elements in U) is needed by the Σ-protocol. It is possible to encode the relation between U and $\hat{U}$ using the switch-network technique in TinyRAM [9], however, the circuit size is $|U|\mathsf{log}(|U|)$. In the following, we introduce a technique that incurs $O(|U|\mathsf{log}^2(|U|))$ field operations (due to half-GCD) for supplying the circuit witness, but it results in $O(|U|)$ circuit size. The scheme runs faster in concrete time. For instance, for a malware-free file of size 1 MB and depth 10, the half-GCD field operation costs 662 s of CPU time, which is less than 2.7% of the total CPU time spent on the proof.

Given U and $\hat{U}$, let p_U and $p_{\hat{U}}$ be the vanishing polynomials. Let p'_U be the derivative of p_U (thus monic and non-zero). It is known that these polynomials satisfy the following [20, Ch. 9] for fields of a large prime order.

$$\frac{p_U}{\gcd(p_U, p'_U)} \equiv p_{\hat{U}} \tag{3}$$

Example 41. Let $U = \{1, 1, 2\}$ and $\hat{U} = \{1, 2\}$. Clearly, $p_U(X) = (X-1)^2(X-2)$, and $p'_U(X) = (X-1)(2(X-2)+(X-1))$. Then, $\gcd(p_U, p'_U) = (X-1)$ and $\frac{p_U}{\gcd(p_U, p'_U)}$ is $(X-1)(X-2)$, which is identical to $p_{\hat{U}}$. □

Circuit Specification: $H \leftarrow \mathbb{CIRC}_{n,\mathbf{A}}(\vec{I}, \overrightarrow{w_1}, \overrightarrow{w_2})$
Parse $\vec{I}$ as $\{r\}$. Parse $\overrightarrow{w_1}$ as $(\vec{T}, \texttt{Evi}(p_{\vec{u}}), k)$ where for each $i \in [0, n)$:
$\vec{T}_i = (\vec{s}_i, \vec{c}_i, \vec{s}_{i+1}, \vec{b}_i)$, and $\vec{u} = \vec{s} \cup \rho(\vec{T})$. Parse $\overrightarrow{w_2} = (v, r_2, \vec{s}_1, \vec{s}_n)$.
Abort if any of the following assertion fails.

1. Assert $\forall i \in [0, n) : \vec{b}_i \in \{0, 1\}$.
2. Assert $\forall i \in [0, n) : \vec{c}_i \in \Sigma$.
3. Assert $\vec{s}_n < |F|$.
4. Fail-edge semantics: $\forall i \in [0, n-1) : \vec{b}_i = 1 \Rightarrow \vec{c}_{i+1} = \vec{c}_i$.
5. $\vec{t} \leftarrow \left(\rho(\vec{T}_i)\right)_{i=0}^{n-1}$. $v_1 \leftarrow (\prod_{i=0}^{n-1}(\vec{t}_i - r))(\prod_{i=0}^{n}(\vec{s}_i - r))$. Parse $\texttt{Evi}(p_{\vec{u}}) = \left(p_i\right)_{i=1}^{7}$. Assert $p_1(r) = v_1$, and verify $\texttt{Evi}(p_{\vec{u}})$ is valid: $p_2 = p_1' \ \wedge \ p_1(r) = p_3(r)p_4(r) \ \wedge \ p_2(r) = p_3(r)p_5(r) \ \wedge \ p_6(r)p_4(r) + p_7(r)p_5(r) = 1$.
6. Assert $v = p_4(r)$.

Recollect input chars from $\vec{c}$ by skipping those for fail edges. Let them be $\vec{d}$.
Return $\mathsf{hash}(\mathsf{encrypt}(\vec{d}, k))$.

Fig. 4. $\mathbb{CIRC}_{n,\mathbf{A}}$ for AC-DFA $\mathbf{A}$ with Input Size n

According to Bézout's Lemma, for any $f, g \in \mathbb{Z}_p[X]$: $\gcd(f, g) = 1$ if and only if there are $a, b \in \mathbb{Z}_p[X]$ s.t. $af + bg = 1$. We denote a and b as $\texttt{BZ}_1(f, g)$, and $\texttt{BZ}_2(f, g)$.

Definition 2. *Let A be a multi-set of field elements of $\mathbb{Z}_p$, let p_A be its vanishing polynomial, and p'_A the formal derivative of p_A. The* evidence polynomials *for p_A, denoted as $\texttt{Evi}(p_A)$, is defined as the following tuple of seven polynomials:*

$$\left(\begin{array}{l} p_A, p'_A, \gcd(p_A, p'_A), \frac{p_A}{\gcd(p_A, p'_A)}, \frac{p'_A}{\gcd(p_A, p'_A)}, \\ \texttt{BZ}_1\left(\frac{p_A}{\gcd(p_A, p'_A)}, \frac{p'_A}{\gcd(p_A, p'_A)}\right), \texttt{BZ}_2\left(\frac{p_A}{\gcd(p_A, p'_A)}, \frac{p'_A}{\gcd(p_A, p'_A)}\right) \end{array}\right)$$

Note that in Definition 2, the 4th polynomial, i.e., $\frac{p_A}{\gcd(p_A, p'_A)}$ is the vanishing polynomial for the support set, i.e., $p_{\hat{A}}$. Given the coefficients of the polynomials in $\texttt{Evi}(p_A)$, one can use arithmetic circuit to efficiently check p'_A is the derivative of p_A. For the rest of the check on $\texttt{Evi}(p_A)$, we present a Feedback Commit-and-Prove (FB-CP) scheme in Sect. 5.2 to take advantage of Schwartz-Zippel.

4.2 Circuit Specification

We present the specification of $\mathbb{CIRC}_{n,\mathbf{A}}$ in Fig. 4. Here, n is the input length and $\mathbf{A}$ is the AC-DFA that it encodes. $\mathbb{CIRC}_{n,\mathbf{A}}$ is defined as a function that takes public input wires $\vec{I}$ and witness input wires that are split into two segments: $\overrightarrow{w_1}$ and $\overrightarrow{w_2}$. All wire values are elements of $\mathbb{Z}_p$. Intuitively, the circuit performs consistency checks on the input wires, and produces an output wire H where $H = \mathsf{hash}(\mathsf{encrypt}(\vec{c}, k))$ for a secret key k and input string $\vec{c}$. The commitment to $\overrightarrow{w_2}$ will be taken from the Groth16 system and used to connect with Σ protocols.

As shown in Fig. 4, there is only one public input wire: r. It is used as the random input point for evaluating polynomials. The first witness segment $\overrightarrow{w_1}$

consists of the acceptance path of the secret input string, evidence polynomials, and a symmetric encryption key k. $\vec{T}$ is a vector of transitions, where each element is written as a tuple $(\vec{s}_i, \vec{c}_i, \vec{s}_{i+1}, \vec{b}_i)$. Let $\vec{u} = \vec{s} \cup \rho(\vec{T})$, i.e., the encoding of states and transitions. Let $\hat{\vec{u}}$ be its support set. The second component of $\vec{w_1}$ is $\texttt{Evi}(p_{\vec{u}})$, and its correctness will be validated in the circuit.

Let $v \leftarrow p_{\hat{\vec{u}}}(r)$, and $r_2 \xleftarrow{\$} \mathbb{Z}_p^*$. Witness segment 2 consists of $(v, r_2, \vec{s}_1, \vec{s}_n)$. Its Pedersen commitment is used to bridge with Σ-protocols. $\vec{s}_1$ and $\vec{s}_n$ are the beginning and last state of the acceptance path, which is used to connect the proofs for consecutive chunks of a big file. r_2 is the opening (blinding factor) of the Pedersen commitment.

In Fig. 4, actions (1)–(4) encode the AC-DFA machinery, mainly the logic of fail-edges. Note that the range proof for states can be saved due to encoding tricks introduced earlier. In action (5), the multi-set of transition encoding and states are built. The circuit computes its evaluation at random point r, and then use it to compare with the $p_1(X)$ in $\texttt{Evi}(p_{\vec{u}})$, thus establishing that the p_1 in $\texttt{Evi}(p_{\vec{u}})$ is a vanishing polynomial of the multi-set of transitions and states. Then the rest of the check verifies Bézout's identity relation and others for the validity of $\texttt{Evi}(p_{\vec{u}})$. This establishes that the $p_4(X)$ in $\texttt{Evi}(p_{\vec{u}})$ is the vanishing polynomial for the support set, i.e., $p_{\hat{\vec{u}}}$. Then it verifies the v in segment 2 is equal to $p_4(r)$. The circuit assumes that all witness wires are committed before the random input r is supplied as input. This is addressed in Section 5.

5 2-Phase Proof Scheme

In this section, we provide a modified Groth16 [30] zkSNARK system, based on which, we develop a 2-phase proof strategy to exploit the locality of AC-DFA. We design a Feedback Commit-and-Prove (FB-CP) scheme for realizing the 2-phase strategy, by connecting Σ-protocols with Groth16.

5.1 *k*-ccGro16

The k-segment *commitment-carrying* Groth16 scheme (k-ccGro16) is a generalization of the ccGro16 proof system in LegoSNARK [17, Appendix H.5], by extending it from one committed segment to multiple.

Definition 3. *A k-segment quadratic arithmetic program (k-QAP) is a tuple*

$$R = \left(p, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, \ell, \{u_i(X), v_i(X), w_i(X)\}_{i=0}^{m}, t(X), \{b_j\}_{j=0}^{k}\right)$$

where its statement is $(a_0, \ldots, a_\ell) \in \mathbb{Z}_p^\ell$ with $a_0 = 1$, and witness $(a_{\ell+1}, \ldots, a_m) \in \mathbb{Z}_p^{m-\ell}$. Let n be the degree of $t(X)$, the witness is accepted if and only if there exists a $n-2$ degree polynomial $h(X)$ s.t. the following holds.

$$\sum_0^m a_i u_i(X) \sum_0^m a_i v_i(X) = \sum_0^m a_i w_i(X) + h(X)t(X)$$

1 Trusted Set-up: $(\sigma_G, \tau) \leftarrow \mathrm{Setup}(R)$

Parse $R = \left(p, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, \ell, \{u_i(X), v_i(X), w_i(X)\}_{i=0}^{m}, t(X), \{b_j\}_{j=0}^{k}\right)$. Sample α, β, γ, $\{\delta_i\}_{i=1}^{k}$, x from $\mathbb{Z}_p^*$.

For $1 \leq j \leq k$, compute $\kappa_j \leftarrow \left\{\frac{\beta u_i(x)+\alpha v_i(x)+w_i(x)}{\delta_j}\right\}_{i=b_{j-1}}^{b_j-1}$. Define simulator trapdoor $\tau = \left(\alpha, \beta, \gamma, \{\delta_i\}_{i=1}^{k}, x\right)$. Generate σ_1, σ_2 as below:

$$\sigma_1 \leftarrow \left(\alpha, \beta, \{\delta_i\}_{i=1}^{k}, \{x^i\}_{i=0}^{n-1}, \{\kappa_j\}_{j=1}^{k}, \left\{\frac{\beta u_i(x)+\alpha v_i(x)+w_i(x)}{\gamma}\right\}_{i=0}^{\ell}, \left\{\frac{x^i t(x)}{\delta_k}\right\}_{i=0}^{n-2}\right), \text{ and } \sigma_2 \leftarrow \left(\beta, \gamma, \{\delta_i\}_{i=1}^{k}, \{x^i\}_{i=0}^{n-1}\right)$$

Compute $\sigma_G \leftarrow ([\sigma_1]_1, [\sigma_2]_2)$. Return (σ_G, τ).

2 Prove: $([A]_1, [B]_2, \{[C_i]_1\}_{i=1}^{k}) \leftarrow \mathrm{Prove}((a_0, \ldots, a_m), R, ([\sigma_1]_1, [\sigma_2]_2))$.

Sample r, s, and $\{r_j\}_{j=1}^{k}$ from $\mathbb{Z}_p^*$. Compute $\pi = \left([A]_1, [B]_2, \{[C_j]_1\}_{j=1}^{k}\right)$ where

$$[A]_1 \leftarrow \left[\alpha + \sum_{i=0}^{m} a_i u_i(x) + r\delta_k\right]_1 \qquad [B]_2 \leftarrow \left[\beta + \sum_{i=0}^{m} a_i v_i(x) + s\delta_k\right]_2$$

For each $1 \leq j \leq k-1$: $[C_j]_1 \leftarrow \left[\left(\sum_{i=b_{j-1}}^{b_j-1} a_i \left(\frac{\beta u_i(x)+\alpha v_i(x)+w_i(x)}{\delta_j}\right)\right) + r_j\delta_k\right]_1$.

For the last segment, $[C_k]_1 \leftarrow \left[\left(\sum_{i=b_{k-1}}^{b_k-1} a_i \left(\frac{\beta u_i(x)+\alpha v_i(x)+w_i(x)}{\delta_k}\right)\right) + \frac{h(x)t(x)}{\delta_k} + As + Br - rs\delta_k - \sum_{i=1}^{k-1} r_i\delta_i\right]_1$.

3 Verify: $0/1 \leftarrow \mathrm{verify}((a_0, \ldots, a_\ell), \pi, R, ([\sigma_1]_1, [\sigma_2]_2))$

Parse π as $\left([A]_1, [B]_2, \{[C_j]_1\}_{j=1}^{k}\right)$. Return 1 if and only if:

$$[A]_1 \cdot [B]_2 = [\alpha]_1 \cdot [\beta]_2 + \sum_{i=0}^{\ell} a_i \left[\frac{\beta u_i(x)+\alpha v_i(x)+w_i(x)}{\gamma}\right]_1 \cdot [\gamma]_2 + \sum_{i=1}^{k} [C_i]_1 \cdot [\delta_i]_2.$$

4 Simulation: $\pi \leftarrow \mathrm{simulate}((a_0, \ldots, a_\ell), \tau, R)$

Sample A, B, $\{C_i\}_{i=1}^{k-1}$ from $\mathbb{Z}_p^*$. Compute $C_k \leftarrow \frac{AB-\alpha\beta-\left(\sum_{i=0}^{\ell} a_i(\beta u_i(x)+\alpha v_i(x)+w_i(x))\right)-\sum_{i=1}^{k-1}\delta_i C_i}{\delta_k}$. Return $\pi = \left([A]_1, [B]_2, \{[C_i]_1\}_{i=1}^{k}\right)$.

Fig. 5. k-Segment ccGro16 (k-ccGro16) Protocol

A witness tuple is divided into k segments, and their scope is defined by a boundary vector $\{b_j\}_{j=0}^{k}$ *in ascending order, where* $b_0 = \ell + 1$ *and* $b_k = m + 1$. *Each segment i corresponds to a slice of the witness, i.e.,* $(a_{b_{i-1}}, \ldots, a_{b_i-1})$. *It is required that for each segment* $j \in [1, k-1]$*:* $\{u_i(x)\}_{i=b_{j-1}}^{b_j-1}$ *are* linearly independent.[6]

Algorithm 5 presents the details of k-ccGro16. We briefly explain the design idea here and refer readers to [17, Appendix H.5] for details, from which k-ccGro16 is derived. The trusted setup generates prover keys $([\sigma_1]_1, [\sigma_2]_2)$ and simulator trap-door τ. Compared with the standard Groth16 [30], for each wit-

[6] Linear independence can be ensured by adding dummy constraints when compiling QAP [8, Lemma 2.4], or directly reasoning about relation between R1CS variables.

ness segment j, there is an additional trapdoor parameter δ_j (in contrast to one δ in Groth16). The $\left\{\frac{\beta u_i(x)+\alpha v_i(x)+w_i(x)}{\delta}\right\}$ component in Groth16 is similarly "segmented" into k segments (denoted as κ_j for each $j \in [1,k]$) in the new scheme. k-ccGro16 generates a proof $\pi = \left([A]_1, [B]_2, \{[C_j]_1\}_{j=1}^k\right)$, which splits the $[C]_1$ in Groth16 into k pieces, i.e., $\{[C_j]_1\}_{j=1}^k$. A blinding component $r_j\delta_k$ is added to each C_j for zero knowledge. δ_j ensures knowledge extraction and avoids inter-mix of values from different segments. Its correctness is stated in Theorem 1.

Theorem 1. *The protocol given in Algorithm 5 is perfectly complete, perfectly zero knowledge, and statistical knowledge sound against adversaries using a polynomial number of generic bilinear group operations.*

5.2 Feedback Commit-and-Prove (FB-CP)

In Fig. 6, we present FB-CP for realizing the 2-phase proof strategy. The basic idea is to fix part of the arithmetic circuit inputs first by committing to them, recompute the polynomial input point by applying Fiat-Shamir, and then compute the rest of Groth16 proof. Note that the $\texttt{CP}_{\texttt{link}}$ in LegoSnark [16] cannot replace FB-CP as it is needed by $\mathbb{CIRC}_{n,\mathbf{A}}$ for validating evidence polynomials.

Its one-time setup has three steps, where step (2) needs to be carried out by a trusted party. The setup first generates σ_Σ and σ_c, the prover/verifier keys for Σ protocols and k-ccGro16. Then for the given AC-DFA, it publishes $\mathbf{S}$, the encoded states and transitions, and then the corresponding subsets bounded by depth, denoted as $\mathbf{S}_i$ for depth i. For each subset, its bilinear accumulator $\mathbf{A}_i$ and the corresponding proof $\mathbf{W}_i$ are pre-computed, to speed up zk-subset proofs later. For $\mathbf{A}_i$, its knowledge proof is $\pi_{\mathbf{A}_i}$, and likewise $\pi_{\mathbf{W}_i}$ for $\mathbf{W}_i$.

The prover is given an input string s, an encryption key k, and her job is to prove that $s \in L(\mathbf{A})$. Note that s will not be visible to verifier, who can only see $\mathsf{hash}(\mathsf{encrypt}(s,k))$. The prover proceeds in three steps.

In the first step, the prover runs s over the AC-DFA, generates its acceptance path and let $\vec{T}$ be the encoded states and transitions and $\hat{\mathbf{T}}$ its support set. She then generates $\texttt{Evi}(p_{\vec{T}})$, as the proof for $\hat{\mathbf{T}}$ being the support set of $\vec{T}$, using these polynomials she prepares $\overrightarrow{w_1} = (\vec{T}, \texttt{Evi}_{\vec{T}}, k)$ as the first segment of witness for $\mathbb{CIRC}_{n,\mathbf{A}}$. She then uses the Groth16 prover algorithm to compute $\mathbf{C}_1$ for the first segment (but not all of the Groth16 proof). Then, she computes $\mathbf{C}_{\hat{\mathbf{T}}}$, the zk-VPD commitment to $\hat{\mathbf{T}}$. These two commitments are first presented to the verifier to fix these inputs.

In the second step, the verifier generates a random nonce r. This is simulated using Fiat-Shamir by setting r as a hash on $\mathbf{C}_{\hat{\mathbf{T}}}$ and $\mathbf{C}_1$. Then the prover takes r, feed it to $\mathbb{CIRC}_{n,\mathbf{A}}$, and generates the rest of the intermediate and output wires. Write $\texttt{Evi}(p_{\vec{T}})$ as $\texttt{Evi}_{\vec{T}}$. Recall that the circuit evaluates all the polynomials in $\texttt{Evi}_{\vec{T}}$ and checks their relations (e.g., Bézout's identity). As all polynomial coefficients are already fixed in $\mathbf{C}_1$, given that r is random, the scheme is sound

1 Setup: $(\mathbb{R}, \sigma_\Sigma, \sigma_c, \{(\mathbf{A}_i, \pi_{\mathbf{A}_i})\}_{i=0}^{u}, \{(\mathbf{W}_i, \pi_{\mathbf{W}_i}, \mathbf{W}_{i,2})\}_{i=0}^{u}) \leftarrow \mathtt{Setup}(\lambda, \mathbf{A})$

1. Parse AC-DFA $\mathbf{A}$ as (Σ, S, s_0, F, T). Compute $\mathbf{S} = S \cup \rho(T)$. Let u be the max depth of $\mathbf{S}$. Compute $\{\mathbf{S}_i\}_{i=1}^{u}$, where $\mathbf{S}_i$ is the subset of $\mathbf{S}$ bounded by depth i. Build $\mathbb{CIRC}_{n,\mathbf{A}}$ from $\mathbf{A}$, and let $\mathbb{R}$ be its QAP. Run $G \leftarrow \mathcal{G}(1^\lambda)$.
2. $(\sigma_c, \tau_c) \leftarrow$ k-ccGro16.$\mathtt{Setup}(\mathbb{R})$ and $\sigma_\Sigma \leftarrow \Sigma_{\mathtt{univar\_zk\_vpd}}.\mathtt{Setup}(1^\lambda, G, |\mathbf{S}|)$.
3. For each $i \in [1, u]$: compute $p_{\mathbf{S}_i}$, and let $w_i(x) = p_{\mathbf{S}}/p_{\mathbf{S}_i}$. Compute: $(\mathbf{A}_i, \pi_{\mathbf{A}_i}) \leftarrow \Sigma_{\mathtt{univar\_zk\_vpd}}.\mathtt{CommitPoly}(p_{\mathbf{S}_i}, 0, \sigma_\Sigma)$, and $(\mathbf{W}_i, \pi_{\mathbf{W}_i}) \leftarrow \Sigma_{\mathtt{univar\_zk\_vpd}}.\mathtt{CommitPoly}(w_i, 0, \sigma_\Sigma)$. $\mathbf{W}_{i,2} = [w_i(s_1)]_2$. Note that $\mathbf{S}_u = \mathbf{S}$ and $\mathbf{A}_u$ is the accumulator for $\mathbf{S}$. Return $(\mathbb{R}, \sigma_\Sigma, \sigma_c, \{(\mathbf{A}_i, \pi_{\mathbf{A}_i})\}_{i=0}^{u}, \{(\mathbf{W}_i, \pi_{\mathbf{W}_i}, \mathbf{W}_{i,2})\}_{i=0}^{u})$.

2 Prove: $\pi \leftarrow \mathtt{Prove}(\mathbf{A}, s, k, \sigma)$

1. Parse σ as $(\mathbb{R}, \sigma_\Sigma, \sigma_c, \{(\mathbf{A}_i, \pi_{\mathbf{A}_i})\}_{i=0}^{u}, \{(\mathbf{W}_i, \pi_{\mathbf{W}_i}, \mathbf{W}_{i,2})\}_{i=0}^{u})$. Run input s over $\mathbf{A}$. Extract multi-set of transitions and states $\vec{T}$, and compute its vanishing poly $p_{\hat{\mathbf{T}}}$. Sample $r_{\hat{\mathbf{T}}}$ from $\mathbb{Z}_p^*$ and compute: $(\mathbf{C}_{\hat{\mathbf{T}}}, \pi_{\mathbf{C}_{\hat{\mathbf{T}}}}) \leftarrow \Sigma_{\mathtt{univar\_zk\_vpd}}.\mathtt{CommitPoly}(p_{\hat{\mathbf{T}}}, r_{\hat{\mathbf{T}}}, \sigma_\Sigma)$; $\overrightarrow{w_1} = \{\vec{T}, \mathtt{Evi}(p_{\vec{T}}), k\}$.
   Let $b_0 = 3$ and $b_1 = |\overrightarrow{w_1}| + 3$ (because QAP public inputs are $\{1, r, H\}$).
   Use k-ccGro16.$\mathtt{Prove}$ in Figure 5 to compute the following:

$$\mathbf{C}_1 \leftarrow \left[\left(\sum_{i=b_0}^{b_1-1} \overrightarrow{w_1}[i-3] \left(\frac{\beta u_i(x) + \alpha v_i(x) + w_i(x)}{\delta_1} \right) \right) + r_1 \delta_3 \right]_1 .$$

2. Sample r_2 from $\mathbb{Z}_p^*$. Apply Fiat-Shamir: $r \leftarrow \mathsf{hash}(H, \mathbf{C}_1, \mathbf{C}_{\hat{\mathbf{T}}})$.
   Let s_1 and s_n be the first and last state of $\vec{T}$. Define $\vec{I} = \{r\}$, $\overrightarrow{w_2} = \{p_{\vec{\mathbf{T}}}(r), r_2, s_1, s_n\}$.
   Compute $H = \mathbb{CIRC}_{n,\mathbf{A}}(\vec{I}, \overrightarrow{w_1}, \overrightarrow{w_2})$. Convert inputs of $\mathbb{CIRC}_{n,\mathbf{A}}$ to R1CS, and then QAP witness, letting it be $\vec{a}$. Now apply the full k-ccGro16 (where $\mathbf{C}_1' = \mathbf{C}_1$ for honest prover).

$$\big(\mathbf{A}, \mathbf{B}, \mathbf{C}_1', \mathbf{C}_2, \mathbf{C}_3\big) \leftarrow \text{k-ccGro16.}\mathtt{Prove}(\vec{a}, \mathbb{R}, \sigma_c).$$

3. Compute $w_{\hat{\mathbf{T}}}(X) \leftarrow p_{\mathbf{S}_d}(X)/p_{\hat{\mathbf{T}}}(X)$; $(\mathbf{W}_{\hat{\mathbf{T}}}, \pi_{\mathbf{W}_{\hat{\mathbf{T}}}}) \leftarrow \Sigma_{\mathtt{univar\_zk\_vpd}}.\mathtt{CommitPoly}(w_{\hat{\mathbf{T}}}, 0, \sigma_\Sigma)$, and $\mathbf{W}_{\hat{\mathbf{T}},2} = [w_{\hat{\mathbf{T}}}(s_1)]_2$. Generate subset and zk-kzg proofs:

$$\begin{aligned}
&(\mathbf{C}'_{\mathbf{A}_d}, \pi_{\mathbf{A}_d \subset \mathbf{A}_u}) \leftarrow \mathtt{PrvSubset}(\mathbf{A}_u, \mathbf{A}_d, \pi_{\mathbf{A}_d}, \mathbf{W}_d, \pi_{\mathbf{W}_d}, \mathbf{W}_{d,2}, 0, 0, \sigma_\Sigma, \mathtt{T});\\
&(\mathbf{C}_{\hat{\mathbf{T}}}, \pi_{\hat{\mathbf{T}} \subset \mathbf{A}_d}) \leftarrow \mathtt{PrvSubset}(\mathbf{C}'_{\mathbf{A}_d}, \mathbf{C}_{\hat{\mathbf{T}}}, \pi_{\hat{\mathbf{T}}}, \mathbf{W}_{\hat{\mathbf{T}}}, \pi_{\mathbf{W}_{\hat{\mathbf{T}}}}, \mathbf{W}_{\hat{\mathbf{T}},2}, r'_{\mathbf{A}_d}, r_{\hat{\mathbf{T}}}, \sigma_\Sigma, \mathtt{F}),\\
&\qquad \text{where } r'_{\mathbf{A}_d} \text{ is the opening of } \mathbf{C}'_{\mathbf{A}_d}\\
&(\mathbf{C}_y, \pi_y) \leftarrow \Sigma_{\mathtt{univar_zk_vpd}}.\mathtt{Open}(p_{\hat{\mathbf{T}}}, r_{\hat{\mathbf{T}}}, r, \sigma_\Sigma)\\
&\pi_{y,2} \leftarrow \pi_{\mathsf{SAME}}(\mathbf{C}_y, \mathbf{C}_2)\{(y, r_y, r_1, r_2, r_3, r_4) : \mathbf{C}_y = y[1]_1 + r_y[s_2]_2 \ \wedge\\
&\qquad \mathbf{C}_2 = y[\kappa_2[0]]_1 + r_1[\kappa_2[1]]_1 + r_2[\kappa_2[2]]_1 + r_3[\kappa_2[3]]_1 + r_4[\delta_3]_1\}
\end{aligned}$$

Return $\big(r, (\mathbf{A}, \mathbf{B}, \mathbf{C}_1, \mathbf{C}_2, \mathbf{C}_3), (\mathbf{C}'_{\mathbf{A}_d}, \pi_{\mathbf{A}_d \subset \mathbf{A}_u}), (\mathbf{C}_{\hat{\mathbf{T}}}, \pi_{\hat{\mathbf{T}} \subset \mathbf{A}_d}), (\mathbf{C}_y, \pi_y, \pi_{y,2})\big)$.

3 Verify: $1/0 \leftarrow \mathtt{Verify}(\mathbb{R}, H, \pi, \sigma)$
Parse π as $\big(r, (\mathbf{A}, \mathbf{B}, \mathbf{C}_1, \mathbf{C}_2, \mathbf{C}_3), (\mathbf{C}'_{\mathbf{A}_d}, \pi_{\mathbf{A}_d \subset \mathbf{A}_u}), (\mathbf{C}_{\hat{\mathbf{T}}}, \pi_{\hat{\mathbf{T}} \subset \mathbf{A}_d}), (\mathbf{C}_y, \pi_y, \pi_{y,2})\big)$. Return 1 iff. all of the following checks pass:

1. k-ccGro16.$\mathtt{Verify}((1, r, H), (\mathbf{A}, \mathbf{B}, \mathbf{C}_1, \mathbf{C}_2, \mathbf{C}_3), \mathbb{R}, \sigma_c) \ \wedge \ r = \mathsf{hash}(H, \mathbf{C}_1, \mathbf{C}_{\hat{\mathbf{T}}})$.
2. $\Sigma_{\mathtt{subset}}.\mathtt{VerSubset}(\mathbf{A}_u, \mathbf{C}_{\mathbf{A}'_d}, \pi_{\mathbf{A}_d \subset \mathbf{A}_u}, \sigma_\Sigma) \ \wedge \ \Sigma_{\mathtt{subset}}.\mathtt{VerSubset}(\mathbf{C}_{\mathbf{A}'_d}, \mathbf{C}_{\hat{\mathbf{T}}}, \pi_{\hat{\mathbf{T}} \subset \mathbf{A}_d}, \sigma_\Sigma)$.
3. $\Sigma_{\mathtt{univar\_zk\_vpd}}.\mathtt{Verify}(\mathbf{C}_{\hat{\mathbf{T}}}, \mathbf{C}_y, r, \pi_y, \sigma_\Sigma) \ \wedge \ \mathtt{CheckSAME}(\mathbf{C}_y, \mathbf{C}_2, \pi_{y,2}, \sigma_\Sigma)$.

Fig. 6. Feedback Commit-and-Prove (FB-CP)

with overwhelming probability by Schwartz-Zippel. For the second segment $\overrightarrow{w_2} = (v, r_2, s_1, s_n)$, the circuit asserts that $v = \mathtt{Evi}_{\vec{T}}[3](r)$ (i.e., the $p_{\hat{\mathbf{T}}}$ for honest prover) and s_1 and s_n are the first and last states in the acceptance path of s.

In the third step, the prover computes Σ-proofs. The π_{SAME} proof $(\pi_{y,2})$ establishes that $\mathbf{C}_2$ and $\mathbf{C}_y$ as Pedersen commitment hide the same value y

(this is verified using `CheckSAME` in step (3) of `Verify`, which is derived from `CheckDLOG`). Then the $\Sigma_{\mathsf{univar\_zk\_vpd}}$ proof establishes that $y = p_{\hat{\mathbf{T}}}(r)$, where $\mathbf{C}_{\hat{\mathbf{T}}}$ is the commitment to $p_{\hat{\mathbf{T}}}$. This now links the Σ-protocols with k-ccGro16, i.e., the $p_{\hat{\mathbf{T}}}$ behind $\mathbf{C}_{\hat{\mathbf{T}}}$ is indeed the $\mathtt{Evi}_{\vec{T}}[3]$ in $\overrightarrow{w_1}$ of $\mathbb{CIRC}_{n,\mathbf{A}}$.

Then the two zk-subset proofs establishes that $\hat{\mathbf{T}}$ is indeed a subset of allowed state/transition set $\mathbf{S}$. This is accomplished by proving that $\hat{\mathbf{T}}$ (letting its depth be d) is a subset of $\mathbf{S}_d$, and then proving $\mathbf{S}_d$ is a subset of $\mathbf{S}$. Due to the use of $\Sigma_{\mathtt{subset}}.\mathtt{PrvSubset}$, $\mathbf{C}'_{\mathbf{A}_d}$ is used (instead of $\mathbf{A}_d$), thus retaining zero knowledge.

The verifier is given a hash H, the prover/verifier key σ, and a proof π. By running the `Verify`() algorithm, she can be convinced that there exists a string s and a key k s.t. $H = \mathsf{hash}(\mathsf{encrypt}(s, k))$ and $s \in L(\mathbf{A})$) where $\mathbf{A}$ is the AC-DFA that the prover and verifier agrees upon (for which σ is the prover/verifier key).

Theorem 2. *The protocol given in Algorithm 6 is perfectly complete, perfectly zero knowledge, and computational sound in the random oracle model, under the assumption of DL, q-PKE, q-SDH, and q-CPDH and adversaries perform polynomial number of generic bilinear group operations.*

Efficiency: Let $|s|$ be the input string length, d the depth of its acceptance path, and $|\mathbf{S}_d|$ the size of the corresponding subset bounded by depth d. The verifier cost is apparently $O(1)$. The prover has to pay $O(|s|\mathsf{log}^2(|s|) + |\mathbf{S}_d|\mathsf{log}(|\mathbf{S}_d|))$ field operations because of half-GCD algorithm and polynomial division, and $O(|s|\mathsf{log}(|s|))$ for R1CS witness to QAP and $O(|\mathbf{S}_d|)$ group operations for Σ-proofs, and the Groth16 itself costs $O(|s|)$ group operations.

6 Implementation and Evaluation

6.1 System Architecture and Data-Set

We provide a full implementation of `zkreg`, based upon `Arkworks` [5] for field/group arithmetic and Rust OpenMPI [41] for distributed processing. We use an instrumented JSnark [35] for converting arithmetic circuit to R1CS, which is fed to a distributed QAP and Groth16 system implemented in Rust. The `zkreg` implementation consists of 33k LOC of Rust and 15k LOC of Java. Our evaluation uses a cluster of 4 CentOS 7.1 servers (each with 112 vCPU and 448 GB of RAM). The actual GCP network bandwidth is 85Gbps, with average ping time 0.1ms. We use BLS12-381 curve, which provides 128-bit security.

The AC-DFA is generated from the hex-signature database of ClamAV, which consists of 101634 signature strings, with 89% being fixed string patterns that can be directly handled by AC-DFA.[7] There are 11040 regex patterns, which will cause state explosion when compiled to DFA. We take a conservative approximation approach, where we treat regex operators as separators and split each regex signature as a collection of fixed pattern strings. A file is regarded as a virus if it contains any of these patterns. The approximation is "conservative" in

[7] Retrieved 03/28/2022. All hex signatures are contained in `main.ndb`.

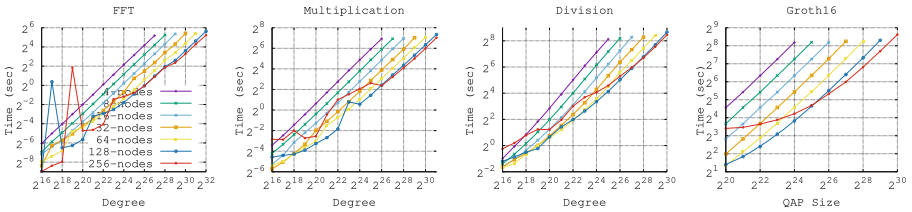

Fig. 7. Distributed System Performance over BLS12-381

the sense that it never reports false-negative for a real virus, but it might report false positives. This is mainly caused by very short patterns resulted from regex signature sets. We removed 8553 such pattern strings. This results in a pattern string collection of 110692, based on which the AC-DFA is built. The AC-DFA contains 19 million states and 342 million transitions. We then run the AC-DFA over all the 2479 ELF files (object and executable files) in a Linux CentOS 7.1, and then there is no false positive reported. The ELF data set has a total size of $747MB$. It exhibits good locality. 96% files have depth less than 40 (which accounts for 22% of the AC-DFA states), with max-depth 252.

6.2 Distributed Processing

We first implement a distributed vector structure using Rust OpenMPI [41], using which coefficients of a polynomial can be stored over a cluster of computer nodes. Then we take the algorithm presented in DIZK [50] and implement distributed FFT operations. Compared with DIZK, we made the following additions, mainly for generating Σ-protocol proofs. We implement polynomial division using Hensel lifting [48] with $O(n\mathsf{log}(n))$ complexity, and also the half-GCD algorithm [45] with $O(n\mathsf{log}^2(n))$ complexity.

Figure 7 presents the performance and scalability of `zkreg`. Distributed FFT is the basis of all. Compared with DIZK, our FFT operation over BN-254 achieves 34× speed-up at degree 2^{28} (2.6 s vs 90.5 s by DIZK). For curve BLS12-381, it costs 37 s for FFT of 2^{32} degree. Half-GCD is only needed at lower degree, However, it does not scale well with MPI. At the degree of 2^{20}, single thread half-GCD costs 450 s, running with 32 (256) MPI-nodes costs 310 (401) seconds. We mainly achieve parallelism of GCD via batch processing proofs. We provide a distributed implementation of the Groth16 system similarly. Compared with DIZK [50], with 256-executors for QAP size of 2^{30} over BN-254, our system needs 266 s of prover time (107× faster than $2^{14.8}$ s needed by DIZK [50, Fig. 5(b)]), and for BLS12-381 it needs 393 s.

Arithmetic circuit is encoded using the JSnark library [35]. Let n be the length of the input string in terms of bytes, the total cost of the circuit is $124n$ R1CS constraints ($30n$ for AC-DFA, $34n$ for encryption/hash, $60n$ for support-set). For instance, an input file of 1008 bytes needs 125263 R1CS constraints. The circuit cost is high, even when bilinear friendly hash and encryption such as

Table 1. Cost Breakdown of Prover Time

	File1 Depth: 10 Size: 2^{16}	File2 Depth: 10 Size: 2^{20}	File3 Depth: 300 Size: 2^{20}
	R1CS: 8.3×10^6 Poly: 2.9×10^7	R1CS: 1.3×10^8 Poly: 2.9×10^7	R1CS: 1.3×10^8 Poly: 3.5×10^8
Step	Cost (sec)	Cost (sec)	Cost (sec)
(1) 1-thread half-GCD.	69.6	662.9	700.6
(2) Load Witness.	1.0	9.3	10.1
(3) Groth16 Step 1.	0.3	2.1	2.1
(4) Apply Fiat-Shamir.	0.2	0.3	0.4
(5) QAP Witness.	2.6	17.1	15.8
(6) Groth16 Step 2.	2.5	21.4	21.1
(7) Σ_{subset} Proof 1.	0.1	0.01	0.01
(8) Σ_{subset} Proof 2.	30.3	31.6	242.9
(9) $\Sigma_{\mathsf{univar\_zk\_vpd}}$ Proof.	0.2	3.0	3.4
Turnaround Time	106.8	747.7	996.4
Adjusted Total	37.6	94.5	314.2

All file size in bytes. R1CS performance "faster" than reported in Fig. 7 because many witness wire inputs (e.g., states and transitions) are 56-bit numbers.

Poseidon [28] and MiMC [3] are used. Also addressing support-set is expensive. An alternative solution is to apply the technique in [23]. However our estimate is that it costs greater ($448n$ R1CS constraints), as the comparison to verify sorted list is expensive. Its advantage is the better scalability than the half-GCD approach. To improve performance in practice, all of circuit generation, conversion to R1CS and QAP are distributed in `zkreg`.

6.3 Proving Linux CentOS 7 Malware Free

To demonstrate the scalability of `zkreg`, we prove all ELFs in a Linux CentOS 7.1 malware free. We need to run a one-time set-up for keys. We generate the subset for depth $D = \{10, 15, 20, 30, 40, 50, 300\}$. The entire set-up takes less than 1 h on the HPC cluster. Due to limit of RAM resources, we chunk larger files into 1 MB pieces, and provide additional Schnorr style DLOG proofs to connect the last/first states of consecutive chunks. We thus have 2954 chunks resulted from 2479 ELF files.

Table 1 shows the prover cost break-down for several sample files (each is padded to a size of closest power of 2 and zk-subset proofs are generated for the closest depth in D). We also show the problem size such as the number of R1CS constraints and the highest degree of polynomials in Σ-protocols. Except for step (1) which is executed with a single thread, all others are run with 256 OpenMPI processes. In practice, step (1) is concurrently processed for multiple

jobs of the same size configuration to cut down cost. We thus have two ways to compute the sum of total cost. The turnaround time stands for the duration from the moment a proof job is submitted to the moment the proof is generated. The adjusted total reflects the "actual" cost of a job by dividing the HGCD cost by the number of concurrent jobs.

According to Lifewire,[8] the average email size is $75kb$. For $64kb$ files, including the single-thread half-GCD cost, its turn-around proof time is 106 s. This implies that the `zkreg` scheme is practical for proving encrypted emails are malware free. The system also demonstrates good performance for $1MB$ files (with depth 10). Its adjusted total is 94 s. For files with depth 300, the prover time is significantly higher, however, such files are rare. The peak system memory usage is $880GB$.

For 2479 ELF files (747 MB) collected, we generate 2954 proofs (2016 bytes each) using 54 h on the HPC cluster, which costs 1350 dollars on GCP. Each proof can be verified by a single-thread verifier in 36 ms. We further apply the inner product pairing product technique [13] to aggregating all proofs. The details are presented in [44, Appendix G]. All 2954 proofs are aggregated into one single proof of 1.96 MB in 727 s, which can be verified in 6.5 (16) seconds with 8 (1) threads.

7 Conclusion

We present an efficient zero knowledge proof system for regular language and demonstrate its performance by proving CentOS 7 malware free.

Acknowledgment. This work is supported by a gift from the Chan Zuckerberg Initiative. Michael Raymond is supported by the Hofstra University SEAS Aspire'22 Summer Research Program and the Stuart and Nancy Rabinowitz Honors College Research Assistant Program. The views and opinions expressed in this work are those of the authors and do not reflect the position of the sponsors of the work. We thank Zachary DeStefano for helpful comments on the paper.

A Univariate Instantiation of zk-VPD [54]

In the following, we present the univariate instantiation of zk-VPD construction in [54, Section 3], as a baseline of comparison.

Let prover key be $\left(\left([s_1^i]_1\right)_{i=0}^{q}, \left([\alpha s_1^i]_1\right)_{i=0}^{q}, [s_2]_1, [\beta s_2]_1, [\alpha s_2]_1, [\alpha]_2, [\beta]_2\right)$, and $(s_1, s_2, \alpha, \beta)$ is the trap-door. We do not distinguish between prover and verifier key for convenience of presentation. Given $p(X)$ and blinding randomness r_p, its zk-VPD commitment $(\mathbf{C}_{p,1}, \mathbf{C}_{p,2})$ is defined as: $([p(s_1) + r_p s_2]_1, [\alpha(p(s_1) + r_p s_2)]_1)$. A commitment can be validated via checking $\mathbf{C}_{p,1} \cdot [\alpha]_2 = \mathbf{C}_{p,2} \cdot [1]_2$.

Given a validated commitment $\mathbf{C}_{p,1}$ and assume that the prover wants to prove that $p(t) = y$ for a given point t. The proof is produced as follows. The prover samples r_y, r_1, r_2 from $\mathbb{Z}_p^*$, and produces the following:

[8] https://www.lifewire.com/what-is-the-average-size-of-an-email-message-1171208.

1. $(\mathbf{C}_{y,1}, \mathbf{C}_{y,2}) = ([y + r_y s_2]_1, [\beta(y + r_y s_2)]_1)$.
2. Compute $q_1(X) = (p(X) - y)/(x - t)$, and $(\mathbf{C}_{1,1}, \mathbf{C}_{1,2}) = ([q_1(s_1) + r_1 s_2]_1, [\alpha(q_1(s_1) + r_1 s_2)]_1)$.
3. $(\mathbf{C}_{2,1}, \mathbf{C}_{2,2}) = ([(r_p - r_y) - r_1(s_1 - t)]_1, [\alpha((r_p - r_y) - r_1(s_1 - t))]_1)$

Intuitively, the claim is that given $(\mathbf{C}_{p,1}, \mathbf{C}_{y,1}, t)$, the prover knows the y behind Pedersen Commitment $\mathbf{C}_{y,1}$ s.t. $y = p(t)$ where $p(X)$ is the polynomial behind $\mathbf{C}_{p,1}$. Here $\mathbf{C}_{p,1}$ is assumed to be already validated (proof of knowledge provided somewhere else), to make a fair comparison with Sect. 3.1. The proof consists of $\mathbf{C}_{y,2}$, $\mathbf{C}_{1,1}$, $\mathbf{C}_{1,2}$, $\mathbf{C}_{2,1}$, $\mathbf{C}_{2,2}$. The verification consists of the following pairing checks:

1. $\mathbf{C}_{y,1} \cdot [\beta]_2 = \mathbf{C}_{y,2} \cdot [1]_2$.
2. $\mathbf{C}_{2,1} \cdot [\alpha]_2 = \mathbf{C}_{2,2} \cdot [1]_2$.
3. $\mathbf{C}_{1,1} \cdot [\alpha]_2 = \mathbf{C}_{1,2} \cdot [1]_2$.
4. $(\mathbf{C}_{p,1} - \mathbf{C}_{y,1}) \cdot [1]_2 = \mathbf{C}_{2,1} \cdot [s_2]_2 + \mathbf{C}_{1,1} \cdot [s_1 - t]_2$.

The verification takes 9 pairings, but the first 4 pairings (i.e., for proof of knowledge of $\mathbf{C}_{y,1}$ and $\mathbf{C}_{2,1}$ can be replaced by Schnorr-style DLOG proofs with minor cost in proof size in exchange for verification speed as operations over $\mathbb{G}_1$ are much faster than pairings), as a fair comparison with Section 3.1. Thus, the univariate zk-VPD scheme needs 5 pairings for verification. The difference in our $\Sigma_{\texttt{univar\_zk\_vpd}}$ scheme in Sect. 3.1 is that the proof of knowledge for $\mathbf{C}_{1,1}$ can be saved, and we also change the last equation to save one more pairing, thus cutting the number of pairings to 2.

B Proof for Lemma 1

Proof. The proof for completeness is apparent. For zero knowledge: the simulator, with the trap-door information s_1, can sample $\mathbf{C}_y, \mathbf{C}_1$ from $\mathbb{G}_1$ and then compute $\mathbf{C}_2 \leftarrow (s_1 - t)\mathbf{C}_1 + \mathbf{C}_y - \mathbf{C}_p$. Then run simulators for the DLOG proof for $\mathbf{C}_y$ and $\mathbf{C}_2$. It thus generates a transcript indistinguishable from ideal ones.

For binding, we will show that if a PPT adversary $\mathcal{A}$ is able to craft a fake evaluation proof with non-negligible probability, then one can build a PPT $\mathcal{B}$ that breaks the q-SDH assumption (Definition 5 in [44]).

Let $\mathcal{B}$ be given a q-SDH challenge [11]: given $\left([s^i]_1\right)_{i=0}^{q}$, $[1]_2$, and $[s]_2$, she needs to present a tuple $(c, [\frac{1}{c+s}]_1)$. From the q-SDH instance, $\mathcal{B}$ builds an instance of $\Sigma_{\texttt{univar\_zk\_vpd}}$ for $\mathcal{A}$ first. Write $s_1 = s$. $\mathcal{B}$ samples s_2 and α, and computes the rest of the keys needed for $\Sigma_{\texttt{univar\_zk\_vpd}}$. For instance, $[s_1 s_2]_1$ is computed as $s_2[s_1]_1$ (as she does not know s_1 but knows s_2). $\mathcal{B}$ then samples t and generates a random polynomial $p(X)$, and passes the $\Sigma_{\texttt{univar\_zk\_vpd}}$ instance to $\mathcal{A}$.

Let $y = p(t)$. Assume $\mathcal{A}$ with non-negligible probability can create a fake evaluation proof $(\mathbf{C}'_1, \mathbf{C}'_2, \pi_{\mathsf{DLOG}_{y'}}, \pi_{\mathsf{DLOG}_{2'}})$ for some $y' \neq p(t)$. By completeness $\mathcal{A}$ can produce a valid proof, and let it be: $(\mathbf{C}_1, \mathbf{C}_2, \pi_{\mathsf{DLOG}_y}, \pi_{\mathsf{DLOG}_2})$. As both pass $\Sigma_{\texttt{univar\_zk\_vpd}}.\texttt{Verify}$, we have the following (using its third equation): (1)

$(\mathbf{C}_p - \mathbf{C}_y + \mathbf{C}_2)\cdot[1]_2 = \mathbf{C}_1\cdot[s_1 - t]_2$; and (2) $\left(\mathbf{C}_p - \mathbf{C}_{y'} + \mathbf{C}'_2\right)\cdot[1]_2 = \mathbf{C}'_1\cdot[s_1 - t]_2$. Subtracting them leads to the following:

$$(\mathbf{C}'_y - \mathbf{C}_y + \mathbf{C}_2 - \mathbf{C}'_2)\cdot[1]_2 = (\mathbf{C}_1 - \mathbf{C}'_1)\cdot[s_1 - t]_2 \tag{4}$$

Apply the knowledge extractors in π_{DLOG_y}, $\pi_{\mathsf{DLOG}_{y'}}$, π_{DLOG_2}, $\pi_{\mathsf{DLOG}_{2'}}$, and write $\mathbf{C}_y = [y + s_2 r_y]_1$, $\mathbf{C}_{y'} = [y' + s_2 r'_y]_1$, $\mathbf{C}_2 = [s_2((s-t)r_1 + r_2)]_1$, $\mathbf{C}_{2'} = [s_2((s-t)r'_1 + r'_2)]_1$, where $\mathcal{A}$ knows y, y', r_y, r'_y, r_1, r'_1, r_2, r'_2 (but not s_1 and s_2). Rewrite Eq. 4 with the known values by $\mathcal{A}$. We have:

$$\begin{bmatrix} (y' - y) + \\ s_2(r'_y - r_y + r_2 - r'_2) \end{bmatrix}_1 \cdot [1]_1 = \begin{pmatrix} \mathbf{C}_1 - \mathbf{C}'_1 \\ +(r'_1 - r_1)[s_2]_1 \end{pmatrix} \cdot [s_1 - t]_2 \tag{5}$$

Write $\Delta y = y' - y$, $A = r'_y - r_y + r_2 - r'_2$, and $\mathbf{h}_1 = \mathbf{C}_1 - \mathbf{C}'_1 + (r'_1 - r_1)[s_2]_1$. These are all known to $\mathcal{A}$. Then Eq. 5 is simplified to the following:

$$[\Delta y + s_2 A]_1 \cdot [1]_2 = \mathbf{h}_1 \cdot [s_1 - t]_2 \tag{6}$$

First of all, note that $\Delta y \neq 0$ by the assumption ($y' \neq y$). This leads to $\Pr[\Delta y + s_2 A = 0] \approx 0$. To see it, assume that $\Delta y + s_2 A = 0$. Since $\Delta y \neq 0$, we have $s_2 A \neq 0$ and thus $A \neq 0$. Then $\mathcal{A}$ can extract $s_2 = -\Delta y / A$, thus breaking the DL assumption.

Now since $\Delta y + s_2 A \neq 0$, and $\mathcal{B}$ knows s_2. $\mathcal{B}$ can compute $\Delta y + s_2 A$ and find its inverse $\frac{1}{\Delta y + s_2 A}$. Plug it into Eq. 6, we have $\frac{1}{\Delta y + s_2 A}\mathbf{h}_1 = [\frac{1}{s_1 - t}]_1$. $\mathcal{B}$ can submit $(-t, \frac{1}{\Delta y + s_2 A}\mathbf{h}_1)$ to break q-SDH. This concludes the proof.

C Proof for Lemma 2

Proof. The completeness is apparent. The HVZK proof is also straight-forward. For soundness (binding), since $\mathbf{C}_p$, $\mathbf{C}_q$, and $\mathbf{C}_w$ all have proof of knowledge, apply their knowledge extractors and re-write them as: $\mathbf{C}_p = [p(s_1)]_1 + r_p[s_2]_1$, $\mathbf{C}_q = [q(s_1)]_1 + r_q[s_2]_1$, $\mathbf{C}_w = [w(s_1)]_1 + r_w[s_2]_1$. Similarly, write $\mathbf{C}_1 = r_1\mathbf{C}_q + r_2\mathbf{C}_w + r_3[s_2]_1 + r_4[1]_1$. Here, the prover knows: $p(X)$, $q(X)$, $w(X)$, r_p, r_q, r_w, $(r_i)_{i=1}^4$.

Now rewrite the last equation in the `VerSubset` algorithm: $\mathbf{C}_p \cdot [1]_2 + \mathbf{C}_1 \cdot [s_2]_2 = \mathbf{C}_q \cdot \mathbf{C}_{w,2}$. This leads to:

$$\begin{pmatrix} (s_2)^2\big(r_1 r_q + r_2 r_w + r_3 - r_q r_w\big) + \\ s_2\big((r_p + r_4) + (r_1 - r_w)q(s_1) + (r_2 - r_q)w(s_1)\big) \\ \big(p(s_1) - q(s_1)w(s_1)\big) \end{pmatrix} = 0 \tag{7}$$

Consider the LHS of Eq. 7 as a Laurent polynomial with s_1 and s_2 as variables (controlled by the set-up), and all the others as known coefficients (controlled by the adversary prover). All coefficients of each term $(s_2)^i$ have to be 0, by Schwartz-Zippel. Then for $(s_2)^0$:

$$p(s_1) - q(s_1)w(s_1) = 0$$

If $p(X) - q(X)w(X)$ is not a zero polynomial, it implies that the prover knows a non-trivial polynomial which evaluates to 0 at secret point s_1. It breaks q-CPDH according to Lemma 1 in [29]. The DL and q-SDH are needed for zk-VPD used.

References

1. Agrawal, S., Ganesh, C., Mohassel, P.: Non-interactive zero-knowledge proofs for composite statements. In: Shacham, H., Boldyreva, A. (eds.) CRYPTO 2018. LNCS, vol. 10993, pp. 643–673. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-96878-0_22
2. Aho, A.V., Corasick, M.J.: Efficient string hatching: an aid to bibliographic search. Commun. ACM **18**, 333–340 (1975)
3. Albrecht, M., Grassi, L., Rechberger, C., Roy, A., Tiessen, T.: MiMC: efficient encryption and cryptographic hashing with minimal multiplicative complexity. In: Cheon, J.H., Takagi, T. (eds.) ASIACRYPT 2016. LNCS, vol. 10031, pp. 191–219. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-53887-6_7
4. Aranha, D.F., Bennedsen, E.M., Campanelli, M., Ganesh, C., Orlandi, C., Takahashi, A.: ECLIPSE: enhanced compiling method for pedersen-committed zkSNARK engines. In: Hanaoka, G., Shikata, J., Watanabe, Y. (eds.) PKC 2022. LNCS, vol. 13177, pp. 584–614. Springer, Cham (2022). https://doi.org/10.1007/978-3-030-97121-2_21
5. arkworks contributors. `arkworks` zksnark ecosystem (2022)
6. Barić, N., Pfitzmann, B.: Collision-free accumulators and fail-stop signature schemes without trees. In: Fumy, W. (ed.) EUROCRYPT 1997. LNCS, vol. 1233, pp. 480–494. Springer, Heidelberg (1997). https://doi.org/10.1007/3-540-69053-0_33
7. Ben-Sasson, E., Chiesa, A., Genkin, D., Tromer, E., Virza, M.: SNARKs for C: verifying program executions succinctly and in zero knowledge. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013. LNCS, vol. 8043, pp. 90–108. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-40084-1_6
8. Ben-Sasson, E., Chiesa, A., Tromer, E., Virza, M.: Succinct non-interactive zero knowledge for a von Neumann architecture. In: USENIX, pp. 781–796 (2014)
9. Ben-Sasson, E., Chiesa, A., Genkin, D., Tromer, E., Virza, M.: Tinyram architecture specification. http://www.scipr-lab.org/doc/TinyRAM-spec-0.991.pdf
10. Benaloh, J., de Mare, M.: One-way accumulators: a decentralized alternative to digital signatures. In: Helleseth, T. (ed.) EUROCRYPT 1993. LNCS, vol. 765, pp. 274–285. Springer, Heidelberg (1994). https://doi.org/10.1007/3-540-48285-7_24
11. Boneh, D., Boyen, X.: Short signatures without random oracles and the SDH assumption in bilinear groups. J. Cryptol. **21**, 149–177 (2008)
12. Bünz, B., Bootle, J., Boneh, D., Poelstra, A., Wuille, P., Maxwell, G.: BulletProofs: short proofs for confidential transactions and more. In: SSP, pp. 315–334 (2018)
13. Bünz, B., Maller, M., Mishra, P., Tyagi, N., Vesely, P.: Proofs for inner pairing products and applications. In: Tibouchi, M., Wang, H. (eds.) ASIACRYPT 2021. LNCS, vol. 13092, pp. 65–97. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-92078-4_3
14. Camenisch, J., Stadler, M.: Efficient group signature schemes for large groups. In: Kaliski, B.S. (ed.) CRYPTO 1997. LNCS, vol. 1294, pp. 410–424. Springer, Heidelberg (1997). https://doi.org/10.1007/BFb0052252

15. Campanelli, M., Faonio, A., Fiore, D., Querol, A., Rodríguez, H.: Lunar: a toolbox for more efficient universal and updatable zkSNARKs and commit-and-prove extensions. In: Tibouchi, M., Wang, H. (eds.) ASIACRYPT 2021. LNCS, vol. 13092, pp. 3–33. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-92078-4_1
16. Campanelli, M., Fiore, D., Querol, A.: LegoSNARK: modular design and composition of succinct zero-knowledge proofs. In: CCS, pp. 2075–2092 (2019)
17. Campanelli, M., Fiore, D., Querol, A.: LegoSNARK: modular design and composition of succinct zero-knowledge proofs. IACR Cryptol. ePrint Arch. **2019**, 142 (2019)
18. Chase, M., Ganesh, C., Mohassel, P.: Efficient zero-knowledge proof of algebraic and non-algebraic statements with applications to privacy preserving credentials. In: Robshaw, M., Katz, J. (eds.) CRYPTO 2016. LNCS, vol. 9816, pp. 499–530. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-53015-3_18
19. Chiesa, A., Hu, Y., Maller, M., Mishra, P., Vesely, N., Ward, N.: Marlin: preprocessing zkSNARKs with universal and updatable SRS. In: Canteaut, A., Ishai, Y. (eds.) EUROCRYPT 2020. LNCS, vol. 12105, pp. 738–768. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-45721-1_26
20. Cohen, J.S.: Computer Algebra and Symbolic Computation. A K Peters (2003)
21. Eagen, L., Fiore, D., Gabizon, A.: cq: cached quotients for fast lookups. IACR Cryptol. ePrint Arch. (2022)
22. Fiat, A., Shamir, A.: How to prove yourself: practical solutions to identification and signature problems. In: Odlyzko, A.M. (ed.) CRYPTO 1986. LNCS, vol. 263, pp. 186–194. Springer, Heidelberg (1987). https://doi.org/10.1007/3-540-47721-7_12
23. Franzese, N., Katz, J., Lu, S., Ostrovsky, R. Wang, X., Weng, C.: Constant-overhead zero-knowledge for RAM programs. In: CCS, pp. 178–191 (2021)
24. Gabizon, A., Williamson, Z.J.: Plookup: a simplified polynomial protocol for lookup tables. IACR Cryptol. ePrint Arch. (2020)
25. Gennaro, R., Gentry, C., Parno, B., Raykova, M.: Quadratic span programs and succinct NIZKs without PCPs. In: Johansson, T., Nguyen, P.Q. (eds.) EUROCRYPT 2013. LNCS, vol. 7881, pp. 626–645. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-38348-9_37
26. Goldreich, O., Micali, S., Wigderson, A.: How to prove all NP statements in zero-knowledge and a methodology of cryptographic protocol design (extended abstract). In: Odlyzko, A.M. (ed.) CRYPTO 1986. LNCS, vol. 263, pp. 171–185. Springer, Heidelberg (1987). https://doi.org/10.1007/3-540-47721-7_11
27. Goldwasser, S., Micali, S., Rackoff, C.: The knowledge complexity of interactive proof systems (extended abstract). In: STOC, pp. 291–304 (1985)
28. Grassi, L., Kales, D., Khovratovich, D., Roy, A., Rechberger, C., Schofnegger, M.: Poseidon: a new hash function for zero-knowledge proof systems. In: USENIX Security (2021)
29. Groth, J.: Short pairing-based non-interactive zero-knowledge arguments. In: Abe, M. (ed.) ASIACRYPT 2010. LNCS, vol. 6477, pp. 321–340. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-17373-8_19
30. Groth, J.: On the size of pairing-based non-interactive arguments. In: Fischlin, M., Coron, J.-S. (eds.) EUROCRYPT 2016. LNCS, vol. 9666, pp. 305–326. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-49896-5_11
31. Grubbs, P., Arun, A., Zhang, Y., Bonneau, J., Walfish, M.: Zero-Knowledge Middleboxes. In: 2022 USENIX Security Symposium, pp. 4255–4272 (2022)

32. Kate, A., Zaverucha, G.M., Goldberg, I.: Constant-size commitments to polynomials and their applications. In: Abe, M. (ed.) ASIACRYPT 2010. LNCS, vol. 6477, pp. 177–194. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-17373-8_11
33. Zhang, C., Zhou, H.S., Katz, J.: An analysis of the algebraic group model. In: Agrawal, S., Lin, D. (eds.) ASIACRYPT 2022. LNCS, vol. 13794, pp. 310–322. Springer, Cham (2022). https://doi.org/10.1007/978-3-031-22972-5_11
34. Kiltz, E., Wee, H.: Quasi-adaptive NIZK for linear subspaces revisited. In: Oswald, E., Fischlin, M. (eds.) EUROCRYPT 2015. LNCS, vol. 9057, pp. 101–128. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-46803-6_4
35. Kosba, A., et al.: *cøø*: a framework for building composable zero-knowledge proofs. Cryptology ePrint Archive. 2015/109 (2015)
36. Liu, T., Xie, X., Zhang, Y.: zkCNN: zero knowledge proofs for convolutional neural network predictions and accuracy. In: CCS, pp. 2268–2985 (2021)
37. Luo, N., Weng, C., Singh, J., Tan, G., Piskac, R., Raykova, M.: Privacy-preserving regular expression matching using nondeterministic finite automata. IACR Cryptol. ePrint Arch. (2023). https://eprint.iacr.org/2023/643.pdf
38. Luthra, A., Cavanaugh, J., Olcese, H.R., Hirsch, R.M., Fu, X.: Zeroaudit. In: ACSAC, pp. 798–812 (2020)
39. Miers, I., Garman, C., Green, M., Rubin, A.D.: Zerocoin: anonymous distributed e-cash from bitcoin. In: SSP, pp. 397–411 (2013)
40. Mohassel, P., Zhang, Y.: SecureML: a system for scalable privacy-preserving machine learning. In: SSP, pp. 19–38 (2017)
41. Rust MPI. MPI binding for Rust (2022)
42. Nguyen, L.: Accumulators from bilinear pairings and applications. In: CT-RSA, pp. 275–292 (2005)
43. Pedersen, T.P.: Non-interactive and information-theoretic secure verifiable secret sharing. In: Feigenbaum, J. (ed.) CRYPTO 1991. LNCS, vol. 576, pp. 129–140. Springer, Heidelberg (1992). https://doi.org/10.1007/3-540-46766-1_9
44. Raymond, M., Evers, G., Ponti, J., Krishnan, D., Fu, X.: Efficient zero knowledge for regular language (extended version). IACR Cryptol. ePrint Arch. (2023). https://eprint.iacr.org/2023/907
45. Schönhage, A.: Schnelle berechnung von kettenbruchentwicklugen. Acta Inform. **1**, 139–144 (1971)
46. Setty, S.: Spartan: efficient and general-purpose zkSNARKs without trusted setup. In: Micciancio, D., Ristenpart, T. (eds.) CRYPTO 2020. LNCS, vol. 12172, pp. 704–737. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-56877-1_25
47. Srinivasan, S., Karantaidou, I., Baldimtsi, F., Papamanthou, C.: Batching, aggregation, and zero-knowledge proofs in bilinear accumulators. In: CCS, pp. 2719–2733 (2022)
48. Sudan, M.: 6.S897 algebra and computation. Polynomial Division (2012). http://people.csail.mit.edu/madhu/ST12/scribe/lect06.pdf
49. Thaler, J.: Proofs, arguments, and zero-knowledge (2022). https://people.cs.georgetown.edu/jthaler/ProofsArgsAndZK.pdf
50. Wu, H., Zheng, W., Chiesa, Z., Popa, R.A., Stoica, I.: DIZK: a distributed zero knowledge proof system. In: USENIX Security, pp. 675–692 (2018)
51. Zapico, A., Buterin, V., Khovratovich, D., Maller, M., Nitulescu, A., Simkin, M.: Caulk: lookup arguments in sublinear time. In: CCS, pp. 3121–3134 (2022)
52. Zhandry, M.: To label, or not to label (in generic groups). In: Dodis, Y., Shrimpton, T. (eds.) CRYPTO 2022. LNCS, vol. 13509, pp. 66–96. Springer, Cham (2022). https://doi.org/10.1007/978-3-031-15982-4_3

53. Zhang, C., DeStefano, Z., Arun, A., Bonneau, J., Grubbs, P., Walfish, M.: Zombie: middleboxes that don't snoop. IACR Cryptol. ePrint Arch. (2023) https://eprint.iacr.org/2023/1022
54. Zhang, Y., Genkin, D., Katz, J., Papadopoulos, D., Papamanthou, C.: A zero-knowledge version of vSQL. IACR Cryptol. ePrint Arch. **2017**, 1146 (2017)