# SoK: Zero-Knowledge Range Proofs

**Miranda Christ**
Columbia University, New York, NY, USA

**Foteini Baldimtsi**
Mysten Labs, Palo Alto, CA, USA
George Mason University, Fairfax, VA, USA

**Konstantinos Kryptos Chalkias**
Mysten Labs, Palo Alto, CA, USA

**Deepak Maram**
Mysten Labs, Palo Alto, CA, USA

**Arnab Roy**
Mysten Labs, Palo Alto, CA, USA

**Joy Wang**
Mysten Labs, Palo Alto, CA, USA

## Abstract

Zero-knowledge range proofs (ZKRPs) allow a prover to convince a verifier that a secret value lies in a given interval. ZKRPs have numerous applications: from anonymous credentials and auctions, to confidential transactions in cryptocurrencies. At the same time, a plethora of ZKRP constructions exist in the literature, each with its own trade-offs. In this work, we systematize the knowledge around ZKRPs. We create a classification of existing constructions based on the underlying building techniques, and we summarize their properties. We provide comparisons between schemes both in terms of properties as well as efficiency levels, and construct a guideline to assist in the selection of an appropriate ZKRP for different application requirements. Finally, we discuss a number of interesting open research problems.

**2012 ACM Subject Classification** Security and privacy → Cryptography

**Keywords and phrases** Range proofs, zero knowledge

**Digital Object Identifier** 10.4230/LIPIcs.AFT.2024.14

**Related Version** *Full Version*: `https://eprint.iacr.org/2024/430` [27]

## 1 Introduction

Zero-knowledge (ZK) proofs have received much attention in recent years, with an abundance of generic protocols being developed using various assumptions and techniques. Although these generic protocols are becoming very efficient and easier to implement, there are still cases for specific types of statements, where customized ZK protocols are preferable.

Zero-knowledge range proofs (ZKRPs) are a subclass of zero-knowledge proofs that proves a structured kind of set membership. A ZKRP allows a prover to convince a verifier that a secret, committed value lies in a given (integer) interval. Brickell et al. [15] introduced the first type of zero-knowledge range proof as a building block in a protocol for revealing a secret discrete logarithm bit-by-bit. Since their introduction, ZKRPs have been used in various applications such as private e-cash protocols [25] (to verify non-negative transaction amounts), anonymous credentials systems [22, 6, 24] (to prove that a secret credential attribute, i.e. user age, falls in a specific range) as well as private voting [43], auctions [3] and privacy preserving federated learning [7] and so on. Additionally, ZKRPs are often used as building blocks for more complex cryptographic schemes. For instance, they have been used to


construct ZK proofs of non-membership [55] and ZK proofs of certain polynomial relations over the integers [23, 21], and they have also been used to prove well-formedness of RLWE ciphertexts [35, 52] and well-formedness of shares in secret-sharing schemes [45, 42].

At the same time, with the rise of decentralized systems and cryptocurrencies, range proofs have received increased attention due to their use in mechanisms that preserve the privacy of transactions posted on the blockchain. For instance, ZKRPs are a key ingredient in confidential transactions [59, 16, 64] – which hide the amount of each transaction posted on the blockchain. The transaction amounts are stored in a committed fashion, and to ensure validity of the transaction the sender must prove that the sum of the output amounts does not exceed the sum of the input amounts. For this check to be sound, the sender must also prove that all output amounts are positive (else an adversarial sender could commit to negative output amounts and create coins out of thin air). For commitments in a group, such as Pedersen commitments, this positivity check also involves showing that the committed value is much less than the order of the group. This check essentially amounts to showing that the committed value is in some integer range $[0, 2^k - 1]$ and is done via a ZKRP. Additionally, ZKRPs are heavily used in protocols for blockchain auditing and solvency solutions [32, 20, 48, 26] to show that transactions or reserves of an organization satisfy certain policies.

This increased interest in ZKRPs has also resulted in a growing number of proposed constructions with different characteristics and properties. With numerous ZKRP constructions available, selecting the suitable scheme for a specific application can be challenging. The goals of this SoK are to organize the space on the various techniques used to construct range proofs, compare their properties in a systematic way, identify open research questions, and provide a guideline to select the appropriate protocol for each type of application.

**Our contributions and organization.** We start by defining the necessary background on cryptographic schemes and computational assumptions in Section 2. In Section 3, we provide a taxonomy of general approaches used in the construction of zero-knowledge range proofs. Concretely, we identify three underlying methods used in the constructions of known ZKRP schemes: (a) square decomposition, (b) binary/n-ary decomposition and (c) hash-chain approach. We describe each method in detail, and for $n$-ary decomposition we present an abstraction that allows us to synthesize the several techniques used. Our abstraction is of independent interest, and could potentially lead to new insights. Then, in Section 4, we collect the set of properties beyond the standard soundness and zero-knowledge that are desirable in certain application scenarios of ZKRPs, such as aggregation, transparent setup and efficiency considerations. In Sections 5-7 we classify all known (to the best of our knowledge) ZKRP constructions under the three methods we identified in Section 3. For each method, we provide an analytical list of known protocols and we compare all protocols based on the desirable properties listed in Section 4. In Section 8, we provide a guideline for how to select the best type of ZKRP construction based on the desired properties and then in Section 9, we report storage and computation (verifier/prover time) costs of the most popular ZKRP constructions using existing and new benchmarks. (We provide a more detailed list of known ZKRP applications in the full version of the paper [27].) Finally, we identify a series of research gaps relevant to ZKRP which we believe can serve as a starting point for future research works in Section 10.

**Comparison with prior work.** We compare our paper with the previous survey of range proofs by Morais, Koens, van Wijk, and Koren [60]. The technical portion of [60] focuses largely on Boudot's four-square decomposition construction [14], the signature-based construction of

CCs [22], and Bulletproofs [17]. It omits or does not go into detail on many other works, such as the line of code-based constructions, the newer and more efficient square-decomposition constructions, the polynomial commitment-based constructions, the hash chain constructions and lattice based constructions. In particular, many of the most efficient schemes such as Sharp [29] and BFGW [11] are not covered in their survey. Their work also provides a comparison only of the three schemes that it focuses on. Our SoK is significantly more comprehensive, and here is a summary of how our work goes beyond [60]. First, to the best of our knowledge, we provide a complete description of techniques and schemes in the ZKRP category and we extensively compare all such schemes based on their techniques, assumptions, and other properties. Additionally, we observe a useful abstraction for binary decomposition-based range proofs, breaking such proofs into two components, and presenting the techniques used for each of these components. An important aspect for our work, especially for practitioners who will use our SoK to determine the most suitable ZKRP for their application, is that we provide new benchmarks and assemble existing benchmarks for easier comparison. We plan to open-source the code used for our benchmarks. Finally, we include open questions and research gaps, and a flowchart to help identify the most appropriate range proof construction family for various applications.

## 2 Preliminaries

We use boldface, like $\mathbf{a} = (a_1, \ldots, a_n)$, to denote a vector, and we let $\mathrm{wt}(\mathbf{a})$ denote its Hamming weight. We use $\circ$ to denote the Hadamard product, i.e., $\mathbf{a} \circ \mathbf{b} = (a_1 b_1, \ldots, a_n b_n)$. For a nonzero value $a$, we use $\mathbf{a}^n$ to denote the vector $(1, a, a^2, \ldots, a^{n-1})$. We let $\mathbf{0}^n$ denote the length-$n$ vector $(0, \ldots, 0)$. For two vectors $\mathbf{x}, \mathbf{y}$, we let $\mathbf{x}^{\mathbf{y}} = (x_1^{y_1}, \ldots, x_n^{y_n})$ denote element-wise exponentiation. We use $\lambda$ to denote the security parameter, $\mathcal{A}$ to denote an adversary, $\mathbb{Z}$ to denote the integers, and $negl(\cdot)$ to denote a negligible function. We use the word *efficient*, or *p.p.t.*, to mean probabilistic polynomial time.

▶ **Definition 1** (Commitment scheme [50]). *A commitment scheme is a pair of efficiently computable algorithms* $(\mathsf{Gen}, \mathsf{Com})$ *where:*

- $\mathsf{Gen}(1^\lambda)$ *is an efficient randomized algorithm that outputs public parameters* $\mathsf{p}$.
- $\mathsf{Com}(\mathsf{p}, m, r)$ *is an efficient deterministic function that takes as input the public parameters, a message* $m$*, and randomness* $r$*. It outputs a commitment to* $m$.

*A commitment scheme must be* binding *and* hiding*, defined as follows:*
*A commitment scheme is* binding *if for all p.p.t. adversaries* $\mathcal{A}$*, it is infeasible to come up with two different messages corresponding to a given commitment.*

$$\Pr_{\mathsf{p} \leftarrow \mathsf{Gen}(1^\lambda)} \left[ \begin{array}{c} (m_0, r_0), (m_1, r_1) \leftarrow \mathcal{A}(1^\lambda, \mathsf{p}) \wedge \\ (m_0 \neq m_1) \wedge \\ \mathsf{Com}(\mathsf{p}, m_0, r_0) = \mathsf{Com}(\mathsf{p}, m_1, r_1) \end{array} \right] = negl(\lambda)$$

*A commitment scheme is computationally (resp., statistically)* hiding *if for all p.p.t. (resp., unbounded) adversaries* $\mathcal{A}$*, it is infeasible to distinguish whether a commitment corresponds to any* $m_0$ *or* $m_1$ *known to* $\mathcal{A}$*. That is, for all* $m_0, m_1$*:*

$$\Pr_{r \leftarrow \$} \left[ \begin{array}{c} c \leftarrow \mathsf{Com}(\mathsf{p}, m_0, r) \\ \mathcal{A}(1^\lambda, \mathsf{p}, c, m_0, m_1) = 1 \end{array} \right] \approx \Pr_{r \leftarrow \$} \left[ \begin{array}{c} c \leftarrow \mathsf{Com}(\mathsf{p}, m_1, r) \\ \mathcal{A}(1^\lambda, \mathsf{p}, c, m_0, m_1) = 1 \end{array} \right]$$

A commitment scheme is *homomorphic* if

$$\mathsf{Com}(\mathsf{p}, m_0, r_0) + \mathsf{Com}(\mathsf{p}, m_1, r_1) = \mathsf{Com}(\mathsf{p}, m_0 + m_1, r_0 + r_1).$$

Next we define zero-knowledge proof and non-interactive zero-knowledge proof (NIZK). Most of the ZKRPs in this SoK are in fact non-interactive. In the following sections, we will skip mention of the non-interactive aspect, unless not clear from context. We provide informal definitions next, while deferring the formal definition of NIZK and its properties to our full version [27].

▶ **Definition 2** (Zero-knowledge proof). *Let $L$ be a language in NP and $R$ be a polynomially verifiable relation, such that $x \in L \iff \exists w : R(x, w)$. A zero-knowledge proof system for $L$ is a tuple of efficient interactive algorithms (Prover, Verifier, Simulator), such that the following properties hold:*

- *Completeness. Given $(x, w) \in R$, the honest execution of the Prover (given $x$, $w$) and the Verifier (given only $x$) result in the Verifier outputting 1.*
- *Soundness. Given $x \notin L$, a malicious Prover interacting with the Verifier can only make it output 1 with negligible probability.*
- *Zero-Knowledge. Given $x \in L$, the Simulator can produce an interaction transcript of an honest Prover with a (possibly) malicious Verifier, that is computationally indistinguishable from an actual execution transcript of the Prover with the Verifier. Note that the Simulator doesn't get $w$, while the Prover gets $w$.*

A *non-interactive zero-knowledge (NIZK)* proof system is a zero-knowledge proof system, where the Prover, given $(x, w)$ just sends one message $\pi$ to the Verifier and the Verifier outputs 0/1 based on $(x, \pi)$. A NIZK has an additional setup algorithm $CRSGen$, which outputs a common reference string (CRS) used by all the proofs and verifications. Instead of a CRS, some NIZKs can also specify a random oracle. The Simulator algorithm is allowed to keep trapdoors about the CRS, or be able to simulate the random oracle.

A zero-knowledge *proof of knowledge* requires that an adversary which produces a valid proof for a statement also knows a valid witness. This is formally captured by requiring the existence of an extractor, which can run the adversary's code and produce the witness.

▶ **Definition 3** (ZKRP). *A* zero-knowledge range proof (ZKRP) *is a zero-knowledge proof of knowledge for the following relation:*

$$RP_{\mathsf{p}} = \{((y, u, v), (m, r)) : y = \mathsf{Com}(\mathsf{p}, m, r) \wedge u \leq m \leq v\}$$

*where $\mathsf{p}, y, u,$ and $v$ are known to the verifier, and* Com *is some particular commitment scheme.*

A question may arise since $\mathsf{p}$ is hard-coded in the language definition: what if a malicious prover samples $\mathsf{p}$ badly and thus renders the NIZK-soundness property vacuous? We note that most applications require both commitment security and NIZK-soundness. These requirements enforce that the attacker of the application's security cannot badly sample $\mathsf{p}$.

**Pedersen commitments.** Most range proofs use *Pedersen commitments* [63] as the underlying commitment scheme. Let $\mathbb{G}$ be a cyclic group of prime order and $g$ and $h$ be generators of that group, where the discrete logarithm relationship between $g$ and $h$ is not known. The Pedersen commitment $\mathsf{Com}(x, r)$ for a value $x \in \mathbb{G}$ with randomness $r$ is $g^x h^r$.

Pedersen commitments are statistically hiding, and their binding property is based on the hardness of the discrete logarithm assumption.

▶ **Definition 4** (Discrete Logarithm Assumption). *Let $\mathbb{G}$ be a group of order $p$ and let $g$ be a generator of $\mathbb{G}$. A challenger samples a random $x \leftarrow \mathbb{Z}_p$ and sends $g^x$ to an adversary. The Discrete Logarithm Assumption states that it is infeasible for the adversary to output $x$, given $(\mathbb{G}, g, g^x)$.*

Apart from the Discrete Logarithm setting, we will also describe schemes based on the hardness of the RSA problem, as well as lattices.

▶ **Definition 5** (RSA Assumption). *A challenger samples primes $p$ and $q$ and sets $N = pq$. It picks a quantity $e$ co-prime to $\phi(N)$, where $\phi(N) = (p-1)(q-1)$ is Euler's totient function. Then it randomly samples $z \leftarrow [1, N]$ and sends $(N, e, z)$ to the adversary. The adversary outputs $y$. The RSA Assumption states that the probability of $y^e = z \pmod N$ is negligible.*

▶ **Definition 6** (Strong RSA Assumption). *The Strong RSA Assumption states that the RSA problem is intractable even when the adversary is allowed to choose the public exponent $e$ (for $e \geq 3$).*

▶ **Definition 7** (SIS Assumption). *Let $q, n, m \in \mathbb{Z}^+$, $\beta \in \mathbb{R}^+$ be given, where $\beta \ll q$. A challenger samples a random matrix $\mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m}$. The SIS Assumption states that it is infeasible for an adversary to find a nonzero $m$-vector $\mathbf{e}$, such that $\mathbf{Ae} = \mathbf{0} \mod q$ and $||\mathbf{e}||_2 \leq \beta$.*

## 3 General Approaches

Efficient zero-knowledge range proofs typically use three classes of approaches: square decomposition, $n$-ary decomposition, and hash chains. We present these approaches below, then explore specific instantiations of these approaches in more detail in their respective sections. We also mention the approach of using generic zero-knowledge proofs.

We describe these approaches for proving that a committed value lies in a range of the form $[0, n^k - 1]$, or that a committed value is positive in the case of square decomposition. Most works consider ranges of this form, which may seem at a first glance to be a relaxed version of the problem. However, when the commitments used are homomorphic, it turns out to be sufficient for constructing more general range proofs with only a small amount of work to translate.

Assume that we have the ability to prove that any committed value is in the interval $[0, n^k - 1]$. To prove that $z$ is in some interval $[u, v]$, one can show first that $(z-u) \in [0, n^k - 1]$ and then that $(v-z) \in [0, n^k - 1]$. Thus, $z \geq u$ and $z \leq v$. Certain constructions from integer commitments (e.g., CKLR [30]) can combine these checks into proving a single equation: $(z-u)(v-z) \geq 0$. It is easy to obtain commitments for $(z-u)$ and $(v-z)$ homomorphically, given a commitment to $z$. For non-homomorphic commitments, one can do this translation by creating a commitment $c$ to $z-u$, proving in zero knowledge that $c$ indeed commits to $z-u$, and performing this range proof with respect to $c$.

### 3.1 Square decomposition

The square decomposition method involves writing the committed integer as a sum of squares in order to prove that it is positive. A common version of this method, the four-square decomposition method, uses Lagrange's four-square theorem. This theorem states that for any integer $z \in \mathbb{Z}_{\geq 0}$, there exist $x_1, x_2, x_3, x_4 \in \mathbb{Z}$ such that

$$z = x_1^2 + x_2^2 + x_3^2 + x_4^2 \tag{1}$$

Thus, to prove that a committed value $z$ is non-negative, it suffices to prove knowledge of $x_1, \ldots, x_4$ such that Equation 1 holds. However, it is crucial that the relation of Equation 1 holds *over the integers* since it may hold for a negative $z$ if we are working in some group rather than over $\mathbb{Z}$. For example, in $\mathbb{Z}_5$ it is possible that $z = -1$ and $0^2 + 1^2 + 2^2 + 2^2 = 9 = z \pmod 5$. To avoid such problems, this approach requires a special type of commitment called an integer commitment.

## Integer commitments

An *integer commitment scheme* is a commitment scheme where binding holds *over* $\mathbb{Z}$. That is, for all p.p.t. adversaries $\mathcal{A}$,

$$\Pr_{\mathsf{p}\leftarrow\mathsf{Gen}(1^\lambda)}\left[\begin{array}{c}(m_0,r_0),(m_1,r_1)\leftarrow\mathcal{A}(1^\lambda,\mathsf{p})\\ \wedge(m_0\neq_{\mathbb{Z}} m_1)\\ \wedge\mathsf{Com}(\mathsf{p},m_0,r_0)=\mathsf{Com}(\mathsf{p},m_1,r_1)\end{array}\right]=negl(\lambda)$$

where $m_0 \neq_{\mathbb{Z}} m_1$ denotes that $m_0$ and $m_1$ are not equal *over the integers*. Bounded integer commitments (used in [30]) satisfy the same binding property, but are weaker in that the message space is restricted to some bounded interval, e.g., $\{x \in \mathbb{Z} : |x| \leq B\}$. For constructing range proofs, this boundedness is not an issue as long as the ranges in question are well within the bounds.

Pedersen commitments, for example, are not integer commitments as their message space is $\mathbb{Z}_p$, and any messages that are equivalent $\pmod p$ result in the same commitment given the same randomness: $g^m h^r = g^{m+p}h^r$ over a cyclic group of order $p$. This attack against binding fails if the order of the group is unknown, and indeed many integer commitment schemes (e.g., Fujisaki-Okamoto commitments, and constructions of [30, 29]) operate in groups of unknown order.

**Fujisaki-Okamoto commitments [40].** We recall an overview of FO commitments but refer the reader to [40] for details. FO commitments operate over a group of unknown order $\mathbb{Z}_N^*$. $g$ and $h$ are generators of large subgroups of $\mathbb{Z}_N^*$, whose relation is unknown. The commitment to $x \in \mathbb{Z}$ is

$$\mathsf{Com}_{FO}(\mathsf{p},x,r) := g^x h^r$$

This commitment is computationally hiding when $r$ is chosen uniformly in the interval $[2^{-\lambda}\cdot N+1,\ldots,2^\lambda\cdot N-1]$. Fujisaki-Okamoto commitments are computationally binding under the factoring assumption.

## 3.2 $n$-ary decomposition

The $n$-ary decomposition method involves committing to the digits of the committed value $z$ in some base $n$. For simplicity, assume for this explanation that we use base 2, although certain approaches can be generalized to other bases. Thus, if the prover wishes to show that $z \in [0, 2^k - 1]$, the prover writes $z = z_0 \cdot 2^0 + z_1 \cdot 2^1 + \ldots + z_{k-1} \cdot 2^{k-1}$ and generates commitments to $z_0, \ldots, z_{k-1}$. The prover then shows that both of the following properties hold, which we present as predicates:

**Digit validity ($\mathsf{DV}(\mathbf{z})$):** $\mathsf{DV}(\mathbf{z}) = 1$ if and only if $z_i \in \{0, 1\}$ for all $i \in [0, k-1]$.

**Representativeness ($\mathsf{Rep}(\mathbf{z}, z)$):** $\mathsf{Rep}(\mathbf{z}, z) = 1$ if and only if $z = \sum_{i=0}^{k-1} z_i \cdot 2^i$.

In terms of these predicates, the $n$-ary decomposition method proves membership in the following relation:

$$R_{\mathsf{decomp}} = \{(\mathsf{p},(c_1,c_2,n,k),(z,\mathbf{z},r,\mathbf{r})) : c_1 = \mathsf{Com}(\mathsf{p},z,r) \\ \wedge\, c_2 = \mathsf{Com}(\mathsf{p},\mathbf{z},\mathbf{r}) \wedge \mathsf{DV}(\mathbf{z}) \wedge \mathsf{Rep}(\mathbf{z},z)\}$$

We note that here, we slightly abuse notation and use $\mathsf{Com}$ to commit to a vector $\mathbf{z}$ with a vector of randomness $\mathbf{r}$.

There are (at least) four common tools used to show that the digits are valid for the desired base; i.e., for binary decomposition they all lie in $\{0, 1\}$. These tools include zero-knowledge *set membership arguments*, *product arguments*, *inner product arguments*, and *polynomial commitments*. These strategies are primarily applicable for base 2, with the exception of set membership, which easily extends to any arbitrary base.

### 3.2.1 Set membership

A set membership proof shows that a committed value lies in some publicly known set $\Phi$; that is, it is a proof of knowledge for the following relation:

$$SM = \{(\mathsf{p}, (\Phi, y), (m, r)) : y = \mathsf{Com}(\mathsf{p}, m, r) \wedge m \in \Phi\}$$

Although one could define a set membership proof with respect to a *private* committed set, in our application the set is determined by the publicly known base.

**Digit validity.** Set membership arguments are useful for instances of $R_{\mathsf{decomp}}$ where the commitment scheme used for $\mathbf{z}$ commits to its components individually; that is,

$$c_2 = (\mathsf{Com}(\mathsf{p}, z_0, r_0), \ldots, \mathsf{Com}(\mathsf{p}, z_{k-1}, r_{k-1}))$$

for some scheme $\mathsf{Com}$. Then, one can show digit validity by providing a set membership proof for each element of $c_2$, with respect to the set $\Phi = \{0, 1, \ldots, n-1\}$. However, such protocols require commitments and range proofs of length at least linear in $k$.

**Representativeness.** There is no general way to show representativeness using set membership proofs; schemes using this construction (e.g., [22]) rely on properties of the specific commitment scheme used.

### 3.2.2 Product arguments

A product argument is a proof system for showing that the product of two committed values $a$ and $b$ is some value $c$. Typically, this equality holds in the group underlying the commitment scheme. For example, for Pedersen commitments in a group of prime order $p$, this argument shows that $ab \equiv c \pmod{p}$. For integer commitments, we have the stronger property that this equality holds over the integers: $ab = c$.

**Digit validity.** Product arguments are useful for proving digit validity base 2, if as with set membership $c_2$ consists of individual bit commitments. To show that a committed bit $b$ is in $\{0, 1\}$, the prover can commit to a value $a$ and prove that $ab = 0$ and $a + b = 1$. Observe that if $b \neq 0$, $a$ must be 0 to satisfy the first equation. Then the second equation implies that $b = 1$. Thus, $b$ must be 0 or 1. Furthermore, the prover can always find a satisfying $a$; if $b = 0$, $a = 1$, and if $b = 1$, $a = 0$. Inner product arguments, which we present next, allow the prover to simultaneously show many product relations more efficiently.

**Representativeness.** As is the case with set membership proofs, product arguments are primarily useful for showing digital validity rather than representativeness.

### 3.2.3 Inner product arguments

An inner product argument (IPA) is a proof system for showing that the inner product of two committed vectors is some value. The inner product used in Bulletproofs [17] shows the following relation, using Pedersen commitments, where $\mathbb{G}$ denotes a group of prime order:

$$\left\{\left(\mathbf{g}, \mathbf{h} \in \mathbb{G}^k, P \in \mathbb{G}, z \in \mathbb{Z}_p \ ; \ \mathbf{a}, \mathbf{b} \in \mathbb{Z}_p^n\right) \ : \ P = \mathbf{g}^{\mathbf{a}}\mathbf{h}^{\mathbf{b}} \wedge z = \langle \mathbf{a}, \mathbf{b} \rangle\right\}$$

Here, $P$ is a binding (but not hiding) commitment to the vectors $\mathbf{a}$ and $\mathbf{b}$. Therefore, Bulletproofs introduces blinding factors to make this argument zero-knowledge. Bulletproofs also constructs an argument for the Hadamard product relation (i.e., $\mathbf{c} = \mathbf{a} \circ \mathbf{b}$) from their inner product argument, though we do not present the details here.

**Digit validity.** A useful fact used when constructing zero-knowledge range proofs from inner product arguments is that with overwhelming probability, the inner product of a nonzero vector $\mathbf{a}$ and a random vector $\mathbf{b}$ is nonzero. Thus, the prover can convince the verifier that $\mathbf{a}$ is $\mathbf{0}^k$ by showing that its inner product with a random challenge vector is 0. Using the same idea as for product arguments, the prover can commit to the binary representation of the given value as a vector $\mathbf{z}$, then use an inner product argument to show simultaneously that *all* components of this vector are indeed bits. That is, the prover commits to a vector $\mathbf{z}' = \mathbf{1}^k - \mathbf{z}$, and shows for a random $\mathbf{x}$ that:

$$\left\langle \mathbf{z}' - (\mathbf{1}^k - \mathbf{z}), \mathbf{x} \right\rangle = 0 \text{ and } \mathbf{z}' \circ \mathbf{z} = \mathbf{0}^k$$

The lattice-based scheme [4] uses this approach as well.

**Representativeness.** Although we presented an inner product relation where the value $z$ is a public input, many inner product arguments, such as that of Bulletproofs, work also when $z$ is secret and the public input includes only a commitment to $z$. One shows representativeness by a single application of this inner product argument, showing $\langle \mathbf{z}, \mathbf{2}^k \rangle = z$.

Bulletproofs combines some of these checks for greater efficiency and uses blinding factors to make their argument zero-knowledge.

### 3.2.4 Polynomial commitments

A polynomial commitment scheme allows a prover to commit to a polynomial $p(\cdot)$ over a finite field $\mathbb{F}_p$, such that a verifier can query a point $x$ to the prover, which can respond with $p(x)$ and a proof $\pi$ that this evaluation is correct. The scheme should be hiding in that the commitment reveals nothing about the polynomial, and the evaluation proofs reveal no extra information beyond the evaluations themselves. Polynomial commitments are binding in that it is computationally infeasible to produce a verifying proof for an incorrect evaluation of the committed polynomial. A useful property of polynomial commitments is that it is easy for a prover to show that a committed polynomial is identically zero, by providing a proof that its evaluation at a random point is zero. By binding and the Schwartz-Zippel lemma, this occurs with only negligible probability if the polynomial is nonzero.

The following approach, which we describe at a high level, was introduced in BFGW [11] and is detailed nicely in [67]. Suppose that we are given a commitment to $z$ in the form of a polynomial commitment to $f$ such that $f(1) = z$. In constructing a range proof for $z \in [0, 2^k - 1]$, it is useful to work over a subgroup $H = \{1, \omega, \omega^2, \ldots, \omega^{k-1}\}$ and use polynomials whose evaluations over $H$ encode the binary representation of $z$. That is, the prover computes a polynomial $g$ such that:

$$g(\omega^{k-1}) = z_{k-1}$$
$$g(\omega^i) = 2g(\omega^{i+1}) + z_i \ \forall i \in \{0, \ldots, k-2\}$$

Another useful property of polynomial commitments is that one can show that a polynomial $g(X)$ is zero on all of $H$ by committing to a related polynomial $g'(X)$ and proving that $g'(X)$ is identically zero over $\mathbb{F}_p$.

**Digit validity.** The prover shows that the following two polynomials are zero over all of $H$:

$$w_2 = g \cdot (1-g)(X-1)(X-\omega)\cdots(X-\omega^{k-2})$$
$$w_3 = [g(X) - 2g(X\omega)] \cdot [1 - g(X) + 2g(X\omega)] \cdot (X-\omega^{k-1})$$

$w_2$ has zeros at $1, \omega, \ldots, \omega^{k-2}$ by construction. It is zero at $\omega^{k-1}$ if and only if $g(\omega^{k-1}) \in \{0,1\}$. For $w_3$, observe that $g(X) - 2g(X\omega)$ is exactly $z_i$ when evaluated at $\omega^i$. Therefore, $w_3$ is zero at $\{1, \ldots, \omega^{k-2}\}$ if and only if $z_i \in \{0,1\}$.

**Representativeness.** The prover shows that the following polynomial is zero over all of $H$:

$$w_1 = (g-f)(X-\omega)(X-\omega^2)\cdots(X-\omega^{k-1})$$

As [11] notes, this approach can be instantiated with any polynomial commitment scheme that is hiding, binding, and additively homomorphic.

## 3.3 Hash chains

Hash chains can be used to prove that a committed value is at least some threshold. In the hash chain approach, a commitment to a value $z$ is $C_z = H^z(r)$, the output of a hash function applied $z$ times to a random $r$. The proof that $z$ exceeds some threshold $t$ is $\pi = H^{z-t}(r)$. A verifier can check that $H^t(\pi) = C_z$; if $z < t$, then $z - t$ is negative and it is infeasible for a cheating prover to compute a preimage of $r$ under $H$.

This simple hash chain requires prover and verifier time that is exponential in $k$ for ranges $[0, 2^k - 1]$. However, using decomposition techniques, HashWires [24] constructs a hash chain-based range proof requiring only $O(k)$ work.

## 3.4 Generic zero-knowledge

There are many efficient generic zk-SNARKs, such as [44, 41, 9, 18]. These proof systems can be used to construct range proofs. However, because they are generic and do not leverage the structure of the range proof relation, they are less efficient than the tailored range proofs we explore. In Section 9, we include efficiency benchmarks for Groth16 [44], one of the most popular generic zk-SNARKs used in practice.

It is worth noting that practical benefits may outweigh these efficiency losses. In particular, because of their wide-ranging applications, generic zk-SNARKs offer convenient, well-engineered, and optimized libraries. For example, we used Circom [8] and rapidsnark [47] for our Groth16 benchmarks. Even so, the prover and verifier times for Groth16 are roughly an order of magnitude larger than the more tailored range proofs. Furthermore, if range proofs are required in a larger system that already uses a generic zk-SNARK elsewhere, using this zk-SNARK for the range proof as well may be practically convenient.

# 4 Desirable properties

All zero-knowledge range proofs must satisfy the standard notions of soundness, completeness, and zero knowledge. All ZKRPs that we cover in this SoK are non-interactive. In this section, we discuss some additional nice features that might be desirable in some settings.

**Efficiency.** Unsurprisingly, it is desirable for ZKRPs to be efficient. In blockchain applications, where a transactor must pay for the storage cost and the amount of computation done by validators, it is especially important to minimize proof size and verifier time. The proof size should be at most linear in $k$ for intervals $[0, 2^k - 1]$, and several schemes offer even constant-sized proofs. Though proof size and verifier time are often priorities, prover time also should not be prohibitively large. Since it is hard to directly compare efficiency of the constructions we discuss in Sections 5 - 7 (even in the asymptotic setting the different parameters make one-to-one comparison very hard), we instead opt to provide a *concrete* comparison of some of the most popular ZKRPs in Section 9.

**Transparent setup.** Some range proofs require public parameters that are generated using secret randomness. It is crucial for the security of the proofs that this randomness is not known to the prover. For example, several square decomposition range proofs use RSA-based integer commitments, which require an RSA modulus. Importantly, this modulus $N$ must be generated in such a way that no party know the factorization of $N = pq$. Similarly, BFGW [11] instantiated with KZG commitments [49] requires a powers-of-tau common reference string, which consists of a series of values $g^{\tau^i}$, where no party knows $\tau$. Protocols that require secrecy of the randomness used in parameter generation are said to require *trusted setup*. Trusted setup does not necessarily require a trusted party, as many trusted setup procedures can be conducted by distributed multi-party protocols. Such protocols (often called ceremonies) exist for many common trusted setup procedures, such as generation of RSA modului and powers-of-tau [39, 12, 62].

Ideally, protocols should have a *transparent* setup procedure that does not require secret randomness. For example, the parameters could be generated by applying a hash function to some public randomness, e.g., to generate a random group element or random matrix.

Note that trusted setup is different from having a trusted issuer responsible for distributing the proper commitments to users, e.g., a Pedersen commitment corresponding to that user's account balance. Any protocol needs to assume that the prover and verifier agree on the commitment at hand.

**Aggregation.** Aggregation allows multiple range proofs to be compressed into a single succinct proof. That is, a single prover holding $m$ commitments to values in the same range $[0, 2^k - 1]$ can efficiently generate a short aggregate proof $\pi$ proving all of these range statements simultaneously. For this aggregation property to be nontrivial, $\pi$ should be shorter than the concatenation of $\pi_1, \ldots, \pi_m$. For example, for Bulletproofs, Bulletproofs+, and Bulletproofs++ [17, 28, 36], the aggregate proof for $m$ values in $[0, 2^k - 1]$ consists of only $O(\log(m \cdot k))$ group elements. As the concatenation of $m$ proofs would require $O(m \cdot \log(k))$ group elements, aggregation results in considerable space savings.

In the notion of aggregation considered so far, a single prover knows the openings of all commitments that are being aggregated. A stronger notion of *multi-prover* aggregation allows one to combine range proofs generated by multiple provers, who wish to hide their openings from one another. Bulletproofs enables such aggregation via an MPC protocol run by the parties holding the commitments [17]. Multi-prover aggregation is harder to achieve, and is less well studied than single-prover aggregation.

Aggregation is especially useful for confidential transactions, where minimizing the amount of space used on-chain decreases gas costs. Since range proofs are used to show non-negativity, all range proofs typically prove membership in the same interval.

**Table 1** Properties of square decomposition-based range proofs.

| Square Decomposition-Based Range Proofs | | | | | |
|---|---|---|---|---|---|
| Scheme | Commitment Scheme | Assumptions | Transp. Setup | Proof Aggregation | Batched Ver. |
| Boudot [14] | F-O [40] | Strong RSA | N | N | N |
| Lipmaa [56] | RDF integer comm.* | Strong RSA | N | N | N |
| Groth [43]** | RDF integer comm.* | Strong RSA | N | N | N |
| CKLR [30] | $\overline{\mathsf{Ped}}$*** | DLOG (optionally DSLE) | N | N | N |
| CKLR [30] | ElGamal variant [30] | DXDH, ORD | Y (class groups) | N | N |
| $\mathsf{Sharp}_{\mathsf{GS}}$, $\mathsf{Sharp}_{\mathsf{SO}}^{\mathsf{PO}}$ [29]† | Pedersen | DLOG, SEI | Y | Y | Y |
| $\mathsf{Sharp}_{\mathsf{HO}}$ [29]† | Pedersen | 1/2-fROOT | N (RSA), Y (class groups) | Y | Y |

- An extension of the Dåmgard-Fujisaki commitment [33] that [56] constructs.
- **[43] is not exactly a new scheme; its contribution is observing a trick that can be applied to make [56] more efficient. Integers of a certain form can be written as a sum of *three* squares, and one can quickly find this decomposition.
- *** A bounded integer commitment scheme based on Pedersen commitments.
- † Sharp is only a *relaxed* range proof and not sufficient for all applications. [29] has a thorough discussion; it is sufficient for anonymous credentials and can be used for some but not all proofs in anonymous transactions, with some modifications. $\mathsf{Sharp}_{\mathsf{HO}}$ refers to a scheme where $\mathsf{Sharp}_{\mathsf{GS}}$ or $\mathsf{Sharp}_{\mathsf{SO}}^{\mathsf{PO}}$ is modified using an additional commitment requiring an RSA group or class group in order to achieve improved soundness.

**Batch verification.** A related property is batch verification, where there exists a process for verifying many proofs together that is more efficient than verifying each proof individually. Batch verification is especially useful in blockchain applications, where a block proposer can aggregate the range proofs for its block and other validators can batch verify this proof more efficiently. Bulletproofs provides batch verification [17], using an observation that verifying many statements of the form $g^x = 1$ can be done by carefully combining them into a single equation requiring fewer exponentiations.

Aggregated range proofs often naturally enable batch verification, as some of the work is effectively done by the aggregator. However, neither aggregation nor batch verification in general implies the other.

**Compatibility with homomorphic commitments.** A commitment scheme Com is homomorphic if $\mathsf{Com}(m_0, r_0) + \mathsf{Com}(m_1, r_1) = \mathsf{Com}(m_0 + m_1, r_0 + r_1)$. It is convenient for applications such as confidential transactions for the underlying commitments to be homomorphic; in particular, homomorphism makes it easier to prove that the sum of transaction output amounts is at least the sum of input amounts.

Most ZKRPs use Pedersen commitments, which are homomorphic. Some exceptions are HashWires [24] and various lattice-based constructions such as KTX [51], which often achieve weaker homomorphism.

## 5 Square Decomposition Constructions

Recall that the square decomposition method involves writing the committed value as the sum of four squares and proving that this equality holds over the integers. Integer commitments, which were discussed in greater detail in Section 3, are a useful tool here. (Recall: An integer commitment scheme is a commitment scheme for which binding holds over the integers: it is computationally infeasible for an adversary to find messages $m_0, m_1$ and randomness $r_0, r_1$ such that $\mathsf{Com}(m_0, r_0) = \mathsf{Com}(m_1, r_1)$, where $m_0 \neq m_1$ *over* $\mathbb{Z}$.) Below we discuss different approaches in this class and also compare them in Table 1. Our comparison is done in terms of the properties discussed in Section 4 except efficiency which as explained above, will be treated separately in Section 9.

Approaches in this class combine integer commitment schemes with a way to prove in zero knowledge that, given a commitments $\mathsf{Com}_x$ and $\mathsf{Com}_y$, the committed values satisfy $x^2 = y$. This implies that $y$ is non-negative. One can generalize this argument to work not just for squares $y$, but for all non-negative integers.

Boudot [14] introduced the approach of proving that a committed value is positive by representing an arbitrary integer as a sum of squares (although not four squares). It uses Fujisaki-Okamoto commitments [40], which require a group of unknown order such as an RSA group. Damgård-Fujisaki commitments [33] are slightly more efficient integer commitments used in subsequent work [56] which refined Boudot's idea and used Lagrange's four square theorem [46, Theorem 369] (which states that every integer can be written as the sum of the squares of four integers). In order to do so, it also introduced an efficient algorithm for finding this four-square decomposition. [43] similarly followed this approach and improved its efficiency by observing that $x$'s of a certain form can be written as the sum of only three squares rather than four. [31] further improved the efficiency and showed that the RSA assumption (rather than the strong RSA assumption, as previously shown) is sufficient to show the security of Damgård-Fujisaki commitments.

The integer commitments used by all of [14, 56, 43] require a group modulus whose factorization is unknown, and therefore require trusted setup. A newer line of work [30, 29] develops new integer commitment schemes, some of which do not require a trusted setup. These schemes also yield much better efficiency, though Bulletproofs and subsequent binary-decomposition-based proofs are still more efficient in practice due to compatibility with available optimized libraries.

CKLR [30] build a *bounded integer commitment* by modifying Pedersen commitments; their scheme essentially enforces that the Pedersen commitment can only be opened to values within some bounded range. They then use this bounded integer commitment to construct their ZKRP following the square decomposition approach. However, their commitment scheme operates over *rationals* rather than integers; while honest openers round these rationals to integers, malicious openers may open to rationals instead which can be problematic for some applications and results in a relaxed notion of soundness. Sharp [29] improves upon CKLR in several ways. In addition to improving over the efficiency of CKLR, Sharp is compatible with standard Pedersen commitments. This is because Sharp effectively moves CKLR's modifications of Pedersen commitments to the *proof* rather than modifying the commitment itself. Two variants of Sharp ($\mathsf{Sharp}_{\mathsf{GS}}$, $\mathsf{Sharp}_{\mathsf{SO}}^{\mathsf{PO}}$), like CKLR, achieve a relaxed notion of soundness. However, they show how to boost soundness by adding an additional commitment using a hidden-order group such as an RSA group or class group; the resulting variants $\mathsf{Sharp}_{\mathsf{HO}}$ achieve standard soundness but require longer proofs. The RSA version also requires a trusted setup. Class groups are hidden-order groups that can be instantiated without a trusted setup, though they are less well-supported than RSA groups from an engineering standpoint. Finally, Sharp improves over CKLR by also offering batching capabilities.

## 6 Binary Decomposition Constructions

CCs [22] introduced the $n$-ary decomposition paradigm to zero-knowledge range proofs. CCs [22] operates over Pedersen commitments and constructs a zero-knowledge set membership protocol by having the verifier publish a signature of each element in the set. The prover then shows in zero knowledge that it knows a signature of its committed value $x$ under the verifier's secret key; by unforgeability this is only possible if the value is in this set. By choosing this set to be $\{0, \ldots, n-1\}$ for base $n$, the prover can commit to the

digits of $x$ and prove that they are valid digits under that base. CCs then uses properties of Pedersen commitments to show that the committed digits indeed represent $x$. The size of the proof is linear in $\log_n 2^k$, where $n$ is the base used and the range is of size $2^k$. By optimizing the choice of the base $n$, this results in a slightly sublinear (in $k$) proof size for a range $[0, 2^k - 1]$. This scheme requires a trusted setup for the signature generation, and it does not offer aggregation.

Subsequent constructions (which we call "Bulletproofs-style" and detail in the next subsection) improve on the efficiency of CCs to avoid this near-linear dependence on $k$. They use inner product arguments or polynomial commitment schemes in clever ways to avoid showing individually that each bit is in $\{0, 1\}$; instead, they are able to roll all of these checks into a shorter proof.

There are also several newer lattice- and code-based constructions that use binary decomposition, such as [61, 4]. While these schemes are less efficient and have very large proofs, their main merit is that they are plausibly post-quantum secure. Additionally, they do offer transparent setup. Developing more practical lattice-based ZKRPs is an interesting research direction as we discuss in more details in Section 10.

When surveying binary decomposition constructions, we separate them into two categories: Bulletproofs-style constructions, which are very practical; and lattice-based constructions, which are primarily of theoretical interest. We provide an overview of all the bulletproof style constructions described below in Table 2.

## 6.1 Bulletproofs-Style Constructions

Bulletproofs [17], arguably considered the state-of-the-art range proof scheme, uses the binary decomposition technique.

Bulletproofs combines the binary decomposition technique with an inner product argument to enable the prover to send only $O(\log k)$ elements. Bulletproofs improves and uses their improvement of an inner product argument (IPA) of [13] where the prover sends only $O(\log k)$ group elements for an IPA over length-$k$ vectors. The key idea in Bulletproofs is that the prover can use this IPA to execute the binary decomposition approach more efficiently; we give intuition for this idea here.

We write $x = a_0 \cdot 2^0 + a_1 \cdot 2^1 + \ldots + a_{k-1} \cdot 2^{k-1}$ and let $\mathbf{a_L} = [a_0, a_1, \ldots, a_{k-1}]$. We let $\mathbf{2^k} := [2^0, 2^1, \ldots, 2^{k-1}]$. The prover shows that it knows a vector $\mathbf{a_R}$ such that:

$$(1)\ \mathbf{a_L} \circ \mathbf{a_R} = \mathbf{0^k}, \qquad (2)\ \mathbf{a_L} - \mathbf{a_R} = \mathbf{1^k}, \qquad (3)\ \mathbf{a_L} \circ \mathbf{2^k} = x$$

Conditions (1) and (2) show that each component of $\mathbf{a_L}$ is in $\{0, 1\}$, using the standard inner product strategy described in Section 3. Condition (3) shows that indeed $\mathbf{a_L}$ contains the binary decomposition of $x$.

These three checks can be combined into a single invocation of the IPA. The IPA used employs a technique that reduces each IPA of length-$n$ vectors to an equivalent IPA over length-$\frac{n}{2}$ vectors. Using this IPA results in a proofs size of $O(\log_2 k)$.

Subsequent works [28, 68, 69] slightly optimize Bulletproofs but keep the scheme and its properties (in particular, its transparent setup and aggregation properties) largely the same. Bulletproofs+ [28] slightly optimizes the Bulletproofs argument to reduce the number of group elements sent by the prover. Bulletproofs++ [36] further improves efficiency by reducing both prover and verifier time. All of these Bulletproofs derivatives maintain the same aggregation properties.

**Table 2** Properties of Bulletproofs-style proofs (all support aggregation and batched verification).

| Bulletproofs-Style Range Proofs (all DLOG-based) | | |
|---|---|---|
| Scheme | Commitment Scheme | Transparent Setup |
| Bulletproofs [17] | Pedersen | Y |
| Bulletproofs+ [28] | Pedersen | Y |
| Bulletproofs++ [36] | Pedersen | Y |
| Flashproofs [68] | Pedersen | Y |
| SwiftRange [69] | Pedersen | Y |
| DRZ [34] | Pedersen | N |
| ZZT+ [72] | Pedersen | N |
| Libert [52] | Pedersen | N |
| BFGW [11] + $KZG_{\mathsf{Ped}}$ | Pedersen | N |
| BFGW [11] + DARKs [19] | DARK [19] | Y with class groups; N with RSA |

Bulletproofs++ [36] extends the recursive-style argument of Bulletproofs to work for any base, yielding asymptotic and concrete efficiency improvements. They do so using a *lookup argument*, which shows that committed values lie in some predefined table. Bulletproofs++ applies this lookup argument to show digit validity in arbitrary bases, allowing them to improve the proof size from Bulletproofs' $O(\log_2 k)$ to $O(\log_2 k / \log_2 \log_2 k)$.

BFGW [11] takes a different approach to the binary decomposition idea, using a polynomial commitment scheme. We detail this approach in Section 3. This scheme assumes that the commitment to a value $x$ is formed as commitment to a polynomial $f$ such that $f(1) = x$. For some polynomial commitment schemes, such a commitment is nonstandard; conveniently, there is a version of KZG commitments for which this is a Pedersen commitment.

BFGW works with any hiding and binding polynomial commitment scheme, yielding different properties based on the scheme used. Notably, when instantiated with KZG commitments [49], BFGW has constant-sized proofs and is competitive efficiency-wise with Bulletproofs. Though KZG commitments require a trusted setup, this setup ceremony is perhaps one of the most commonly run, and some blockchains such as Ethereum have run a KZG ceremony.[1] In Section 9, we provide the first efficiency (prover and verifier time) benchmarks that we know of for BFGW + KZG. If the Pedersen variant of KZG commitments is used, BFGW + KZG is compatible with Pedersen commitments. BFGW can also be instantiated with DARKs [19], which do not require a trusted setup. Both BFGW + KZG and BFGW + DARKs are aggregatable.

## 6.2 Lattice- and code-based constructions

There are several lattice- and code- based zero knowledge range proof schemes. These schemes have the advantages that they are plausibly post-quantum secure and have a transparent setup. However, they are concretely much less efficient than the discrete logarithm-based schemes such as Bulletproofs. In particular, they have very long proofs. Thus, one worthwhile research direction is to improve the efficiency of these lattice-based protocols, such as [4, 37, 58]. One area for improvement is in the repetition required to achieve negligible soundness error. Most of these schemes build on protocols with constant soundness and must repeat the protocol $\Omega(\lambda)$ times to achieve $\lambda$ bits of security. When made non-interactive, this amplification results in large proofs.

[1] `https://blog.ethereum.org/2023/01/16/announcing-kzg-ceremony`

Lattice- and code-based schemes typically use the binary decomposition approach, where the prover already holds a commitment to the bits $b_0, \ldots, b_{k-1}$ of the value in question. The prover wants to show that $\sum_{i=0}^{k-1} 2^i \cdot b_i \leq \beta$ for some $\beta$. This condition can be written equivalently as a system of equations over the bits modulo 2. Such systems of equations can be proven in zero-knowledge using *Stern-like* protocols [66].

In this section, we present several ideas involved in lattice-based schemes. We first present a lattice-based commitment scheme, KTX [51], that is used in some of these ZKRPs. In doing so, we emphasize several challenges common to many lattice-based schemes. We then give a high-level description of Stern-like protocols, a standard technique for lattice-based zero-knowledge proofs. We also include a table with newer lattice-based schemes that offer constructions tailored to range proofs. We do not include all generic lattice-based zero-knowledge proof constructions.

**KTX commitment scheme ([51]).** The KTX commitment scheme is based on the hardness of the Short Integer Solution (SIS) problem. Let $\lambda$ be the security parameter, $L$ be the number of bits to be committed to, and $q$ be a prime modulus of size $O(\lambda\sqrt{L})$. Let $m = 2\lambda\lceil \log q \rceil$. The scheme uses public parameters $(\mathbf{A}, \mathbf{B})$ chosen uniformly from $\mathbb{Z}_q^{\lambda \times L} \times \mathbb{Z}_q^{\lambda \times m}$. The commitment to a bit vector $\mathbf{x} \in \{0,1\}^L$ is the vector

$$\mathbf{c} = \mathbf{A} \cdot \mathbf{x} + \mathbf{B} \cdot \mathbf{r} \pmod{q}$$

where $\mathbf{r}$ is sampled uniformly from $\{0,1\}^m$. This scheme is statistically hiding and computationally binding assuming that the public parameters are sampled uniformly.

Note that KTX commitments are only approximately homomorphic. While it holds that:

$$\mathbf{A} \cdot \mathbf{x_1} + \mathbf{B} \cdot \mathbf{r_1} + \mathbf{A} \cdot \mathbf{x_2} + \mathbf{B} \cdot \mathbf{r_2} = \mathbf{A}(\mathbf{x_1} + \mathbf{x_2}) + \mathbf{B}(\mathbf{r_1} + \mathbf{r_2}) \pmod{q},$$

note that $(\mathbf{x_1} + \mathbf{x_2})$ and $(\mathbf{r_1} + \mathbf{r_2})$ may not be 0/1 vectors. Therefore, $A(x_1 + x_2) + B(r_1 + r_2)$ is not necessarily a valid commitment to a message in the message space. Many commitment schemes used by schemes in this section have similar limited homomorphism.

Note also that KTX commitments do not require a trusted setup to generate the public parameters $\mathbf{A}$, $\mathbf{B}$, and $q$, as these matrices are uniformly random and $q$ can be publicly known. Many lattice-based commitment schemes similarly use random matrices as the public parameters. All of the range proofs in this section offer transparent setup.

**Stern-like protocols.** Stern's original protocol [66] proves in zero knowledge that a committed bit vector has a certain Hamming weight; that is, it is a zero-knowledge argument of knowledge for the following relation:

$$\{((\mathbf{H}, \mathbf{y}, w), \mathbf{s}) \in \mathbb{Z}_2^{n \times m} \times \mathbb{Z}_2^n \times \mathbb{Z} \times \mathbb{Z}_2^m : (\mathrm{wt}(\mathbf{s}) = w) \wedge (\mathbf{H} \cdot \mathbf{s} = \mathbf{y})\}$$

The key idea behind Stern's protocol is that the prover permutes the bits of $\mathbf{s}$ to obtain $\mathbf{s}'$ which it reveals to the verifier. It also convinces the verifier that $\mathbf{s}'$ is indeed a permutation of $\mathbf{s}$ under some $\pi$. $\mathbf{s}'$ has the same Hamming weight as $\mathbf{s}$, and the distribution of $\mathbf{s}'$ is identical for any $\mathbf{s}$ satisfying the relation – therefore, $\mathbf{s}'$ reveals no information about $\mathbf{s}$. At a high level, the prover samples a random blinding factor $\mathbf{r}$ and constructs three commitments, which it sends to the verifier, as follows:

$$c_1 = \mathsf{Com}(\pi, \mathbf{H} \cdot \mathbf{r}), \quad c_2 = \mathsf{Com}(\pi(\mathbf{r})), \quad c_3 = \mathsf{Com}(\pi(\mathbf{r} \oplus \mathbf{s}))$$

Here, $\pi(\mathbf{v})$ denotes the vector obtained by permuting the components of $\mathbf{v}$ under $\pi$. We now run one of three randomized checks: the verifier sends the prover $b \in \{0,1,2\}$. In each of these tests, the prover opens a different combination of the commitments and sends some

additional information, e.g., $\pi(\mathbf{s})$ for $b = 2$. The cheating prover cannot pass all of these tests simultaneously and therefore fails with probability at least 1/3. Note that running all of these tests at once would reveal information about $\mathbf{s}$.

This permute-then-reveal strategy can be used for other relations with similar properties. [61] provides an abstraction of such relations, in terms of some set VALID, which in Stern's original protocol was $\mathsf{VALID} = \{\mathbf{s} : \mathrm{wt}(\mathbf{s}) = \mathsf{w}\}$:

$$R = \{((\mathbf{H}, \mathbf{y}), \mathbf{s}) \in \mathbb{Z}_2^{n \times m} \times \mathbb{Z}_2^n \times \mathsf{VALID} : \mathbf{H} \cdot \mathbf{x} = \mathbf{y}\}$$

**Correctness under permutation:** For all $((\mathbf{H}, \mathbf{y}), \mathbf{s}) \in \mathbb{Z}_2^{n \times m} \times \mathbb{Z}_2^n \times \mathbb{Z}_2^m$ and all permutations $\pi$ over $[m]$,

$$\mathbf{s} \in \mathsf{VALID} \iff \pi(\mathbf{s}) \in \mathsf{VALID}$$

**Hiding under permutation:** For all $\mathbf{s} \in \mathsf{VALID}$, the distribution of $\pi(\mathbf{s})$ where $\pi$ is a random permutation over $[m]$ is uniform over the set VALID

Even given a relation that does not fit the above requirements, one can sometimes construct an associated relation (e.g., using a common technique called *extension*) that does fall into this paradigm and allows one to construct the desired argument.

Other relations that can be proven under Stern's paradigm include proving knowledge of one secret bit that may appear in multiple equations [54], or proving the knowledge of the product of two secret bits [53]. Stern-like techniques underlie many older lattice- and code-based zero-knowledge protocols. However, recall that due to the randomized tests, Stern's original protocol has soundness error 2/3. In general, Stern-like protocols have constant soundness error and thus require roughly $\lambda$ repetitions for $\lambda$ bits of security. Thus, once made non-interactive via Fiat-Shamir, these protocols result in long proofs.

Only recently have techniques emerged for avoiding Stern-like protocols in constructing lattice-based ZKRPs, whose state-of-the-art is thus not reflected in the previous ZKRP survey [60]. These new techniques resulted in a surge of lattice-based constructions with greatly improved efficiency, with proofs on the order of 10,000 KB rather than 100,000 KB. However, this efficiency still lags behind many non-lattice-based constructions with $\tilde{5}$00-byte proofs, as seen in Table 4. Improving lattice-based schemes remains a fruitful research direction.

[37] proposes techniques for avoiding the repetition that Stern-like protocols require for soundness. Their *one-shot* protocol saves a factor of $\lambda$ computation time over repeated Stern-like protocols, though the proofs are still quite long as shown in Table 4. One-shot approaches are a fruitful direction for developing a more practical (in terms of both communication and computation) lattice-based ZKRP.

ALS [4] uses an inner product argument in the $n$-ary decomposition approach, which results in significantly shorter proofs compared to other lattice-based constructions; see Table 4. Its proofs are roughly an order of magnitude larger than those of the most efficient non-lattice schemes, such as Bulletproofs. Another barrier to practical efficiency is that the proofs of ALS cannot be aggregated.

## 7 Hash chain constructions

Payword [65] was the first to use hash chains to construct a range proof for electronic payments, and HashWires [24] more recently revisited this idea with great efficiency improvements. In this approach, the core idea is that a commitment $C_x$ to a value $x$ is the output of a hash function evaluated $x$ times on a random value. That is, $C_x = H^x(r)$ for a random $r$. The proof that $x$ is at least some threshold $t$ is a value $\pi = H^{x-t}(r)$ such that applying the

**Table 3** Properties of lattice- and code-based range proofs. CKLR supports proof aggregation and batch verification, while it is unclear if the other schemes natively do so.

| Lattice- and Code-Based Range Proofs | | | |
|---|---|---|---|
| Scheme | Commitment Scheme | Assumptions | Transp. Setup |
| LLNW [55] | KTX [51] | SIVP | Y |
| ESLL [37] | UMC, HMC [10, 5, 38] | Module-SIS, Module-LWE | Y |
| YAZ+ [71] | KTX [51] | LWE, SIS | Y |
| ALS [4] | BDLOP [5] | Module-SIS, Module-LWE | Y |
| CKLR [30]† | BDLOP [5], as modified by [71] | LWE, SIS | Y |
| LNS [58]* | BDLOP [5] | Module-SIS, Module-LWE | Y |
| LNP [57] | ABDLOP [1, 5] | Module-SIS, Module-LWE | Y |
| Code-based [61] | [61] | 2-RNSD | Y |

- †CKLR [30] uses the square decomposition approach, but one of their constructions is lattice-based.
- *In addition to their standard range proof, LNS [58] also constructs an *approximate* range proof, showing that $z \in [0, n \cdot 2^k - 1]$ for some small $n$. While relaxed, this kind of approximate range proof is sufficient for showing smallness of vectors, which is an application they target. Its efficiency does not depend on $k$.

hash function $t$ more times to $\pi$ yields $C_x$; that is, $H^t(\pi) = C_x$. Since the hash function is hard to invert, if $x - t$ is negative it should be hard for the prover to find an accepting $\pi$. Importantly, though, $C_x$ must be well-formed to ensure soundness. Thus, the setting where hash chain constructions can be used is slightly more restricted.

HashWires [24] defines a relaxation of zero-knowledge range proofs called *credential-based range proofs* (CBRPs). This notion is weaker than general ZKRPs in that the commitment is assumed to be well-formed. Soundness is shown only under this assumption, which is motivated by a setting where a trusted authority distributes commitments to parties that later prove that their committed values exceed some threshold. For example, the trusted authority may be a government, and the commitments might be used for credentials including citizens' ages. If a commitment is signed by this trusted authority, a verifier can be confident that the commitment is properly formed. Technically, this implies that when defining soundness for CBRPs, the adversary cannot produce the commitment (as defined in the statement of Def. 3, but instead is honestly generated (the full definition of CBRPs can be found in [24]).

As described, the time to generate $\pi$ and $C_x$ is linear in $x$, and the verifier time is linear in $t$. This is very expensive if we wish to prove that $x$ is in some large range $[0, 2^k - 1]$; ideally, these costs should grow at most linearly with $k$. HashWires achieves this by observing that $x$ can be written in some base $u$, and the proof can be broken into several sub-chains to greatly improve this efficiency (they called this a *minimum dominating partition*). This base can be chosen to trade off between proof size and prover/verifier efficiency. In our later discussion of efficiency, we include benchmarks for a variety of bases. We will see in Section 9 that HashWires is extremely concretely efficient, in terms of both verifier time and prover time. Its proof sizes are also competitive with other constructions.

## 8 Choosing the construction family for your application

As there are dozens of ZKRP constructions, choosing the appropriate scheme for a particular application can be challenging. In Figure 1, we give a flowchart for narrowing down the class of range proofs depending on constraints. The next section gives an efficiency comparison to help choose a scheme within this class.

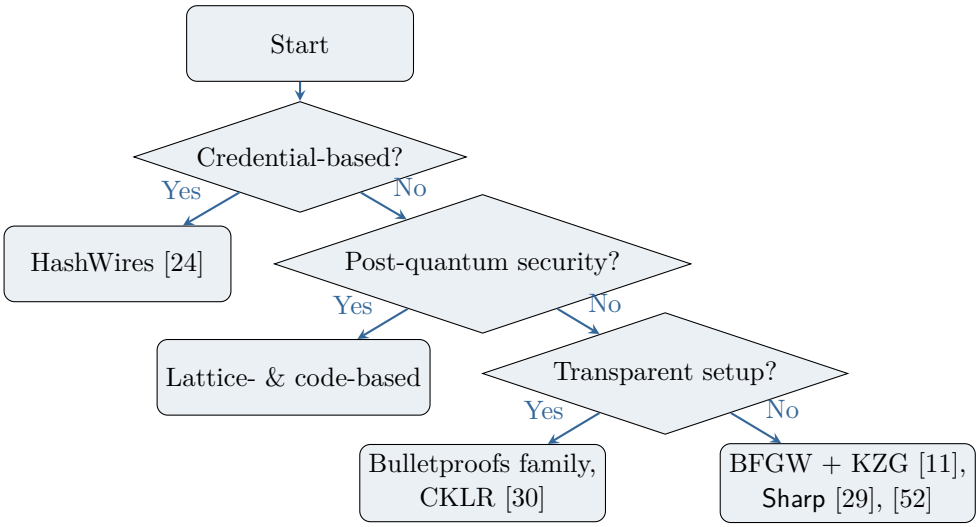


**Figure 1** Flowchart for choosing a range proof based on desired properties.

HashWires [24] are concretely quite efficient and use only hash functions; thus, they are plausibly post-quantum secure and do not require a trusted setup. However, they're in a more stringent trust model (they are *credential-based range proofs* as defined in Section 7), where there is a trusted issuer distributing commitments; that is, soundness holds only if the commitment is well-formed. If the desired use case does have this type of trusted issuer, HashWires is likely the most efficient scheme.

Among the remaining constructions, only the lattice-and code-based constructions are plausibly post-quantum secure, and thus if this is a requirement this class is the only option. These schemes have relatively large proof sizes (on the order of 10KB). Hash-based generic zero-knowledge proof systems may be considered as well.

If trusted setup is allowed, there are several schemes with very short proofs and efficient verifier and prover. BFGW + KZG [11], Sharp [29], and Libert's DLOG-based scheme [52] all have constant-sized proofs.

If trusted setup is undesired, the Bulletproofs family is recommended. Although many lattice- and code-based constructions do not require a trusted setup, all Bulletproofs-style constructions have much shorter proofs. Even if a trusted setup is allowed, Bulletproofs-style constructions may still be worth considering depending on how much one values short proofs. Though their proof sizes are not constant, they seem to be the most commonly used in practice. We list CKLR [30] as well because it has comparable efficiency to Bulletproofs on paper and also does not require trusted setup. However, it has several drawbacks: it does not allow batching, it is less efficient in practice due to its incompatibility with optimized libraries for common elliptic curves, and it offers a more relaxed notion of security. For certain applications where these drawbacks are less important, CKLR may be worth considering.

## 9 Efficiency Comparison

This section includes an efficiency comparison of various ZKRPs. In Table 4, we compile both concrete and asymptotic proof sizes for schemes of particular interest. The concrete proof sizes have been extracted directly from the schemes' respective papers, as the proof sizes are largely the same across machine configurations. Groth16 has the shortest range proofs for a 64-bit range at 192 bytes whereas HashWires has the shortest range proofs at 177 bytes for a 32-bit range.

In Table 5, we record prover and verifier times for various schemes. We add many of our own benchmarks to ensure that configurations are normalized. In particular, we add a benchmark for Groth16 [44] that was absent in prior work. The configurations for benchmarks that we pull from other papers are noted below.

**Other benchmarks.** The Sharp paper's [29] benchmark was run on a MacBook Pro with a 2.3 GHz Intel core i7 processor and uses the library libsecp256k1 [70]. The HashWires paper [24] includes a benchmark for Bulletproofs which is significantly faster than ours. They used an AVX2 backend was used which significantly speeds up curve arithmetic. We include this benchmark in addition to ours, to reflect the speedup possible with their configuration.

**Our benchmarks.** We add our own benchmarks for Hashwires (base 16 and base 256), Bulletproofs, BFGW + KZG, and Groth16. In all cases, we record the median running time over 100 runs. We plan to open source all of our benchmarks for reproducibility.

For Groth16, we implement range proofs with two versions of the commitment scheme: the well-established Pedersen commitments and the new zk-friendly Poseidon commitments. We've used Circom [8] for writing circuits and rapidsnark [47] for generating and verifying the Groth16 proofs.

The implementations for Hashwires, Bulletproofs and BFGW + KZG are in Rust. All the benchmarks were run on a AMD EPYC 7443P 24-Core with 512GB of RAM (a c3.large.x86 machine hosted by latitude.sh). We explicitly chose a non-Mac machine because rapidsnark leverages Intel Assembly to speed up Groth16 proof generation.

Hashwires has the fastest proof generation and verification times. Both BFGW + KZG and Groth16 have constant-sized proofs but they are less computationally efficient than others. Groth16 has the longest proof generation times. This is expected because we are instantiating range proofs within a general-purpose zk proof system.

It is worth noting that in practice the availability of a reliable library may outweigh mild efficiency gains. Bulletproofs is the most widely used range proof in practice and is likely a good choice. Groth16, though not tailored to range proofs, is one of the most popular general-purpose zero-knowledge proof systems and offers several well supported libraries; we use Circom [8] and rapidsnark [47]. From our benchmarks, one can see the efficiency gains offered by tailored range proof solutions over generic solutions, which can be seen especially in the long prover times required for Groth16 relative to the other range proofs.

## 10 Research Gaps

▶ **Research Gap 1.** *Practical transparent constant-sized range proofs.*

No zero-knowledge range proofs are practical, transparent, and have constant-sized proofs. Bulletproofs and its close relatives have transparent setup but have proofs of size $O(\log k)$ for a $k$-bit range. BFGW + KZG has constant-sized proofs but requires a trusted setup; BFGW + DARKs has a transparent setup but requires $O(\log k)$-sized proofs. CKLR has a transparent setup and has constant-sized proofs but achieves only a relaxed notion of soundness. Furthermore, its proofs are not as practically efficient as the above schemes because they use less common curves that optimized libraries do not support.

▶ **Research Gap 2.** *Shorter (even amortized) lattice- or code-based ZKRPs.*

**Table 4** Proof sizes in bytes for 64- and 32-bit ranges. The benchmark for each of these schemes is from that scheme's original paper, except where otherwise noted.

| Scheme | Proof size (bytes) | | Proof size (asymptotic) |
|---|---|---|---|
| | 32-bit range | 64-bit range | $k$-bit range |
| Bulletproofs | 610 | 675 | $O(\log k)$ |
| BFGW + KZG † | 576 | 576 | $O(1)$ |
| $\mathsf{Sharp}_{\mathsf{GS}}$ | 318 | 360 | $O(1)$ |
| $\mathsf{Sharp}^{\mathsf{Po}}_{\mathsf{SO}}$ | 335 | 389 | $O(1)$ |
| $\mathsf{Sharp}_{\mathsf{RSA}}$ | 751 | 793 | $O(1)$ |
| HashWires (Base 16)† | 231 | 263 | $O(\log k)$ |
| HashWires (Base 256)† | 167 | 199 | $O(\log k)$ |
| Groth16 [44]§ | 192 | 192 | $O(1)$ |
| Lattice-based ALS [4]** | 5,900 | - | $O(k)$ |
| Lattice-based ESLL [37] | 58,000 | 93,000 | $\Omega(k)$* |
| Lattice-based LNS [58]** | 11,800 | - | $o(k)^{\ddagger}$ |

- † Our own benchmark.
- § Benchmark from HashWires [24], over the BLS12-381 curve.
- *See [37] for the exact expression, which includes several other parameters not described here. It is $\Omega(k)$ and is large relative to the other schemes.
- **The proof sizes for 64-bit ranges were not included in [4, 58]. Note that [4] has linear growth, so extrapolating from its 5,900-bit proof for 32-bit ranges, its proof for 64-bit ranges would be large.
- ‡ See [58] for the exact expression, which is complicated; it is sublinear in $k$.

**Table 5** Verifier and prover times. †Our own benchmark.

| Scheme | Verifier Time (ms) | | Prover Time (ms) | |
|---|---|---|---|---|
| | 32-bit range | 64-bit range | 32-bit range | 64-bit range |
| Bulletproofs† | 1.37 | 2.51 | 6.32 | 11.96 |
| $\mathsf{Sharp}^{\mathsf{Po}}_{\mathsf{SO}}$ | 0.74 | 0.75 | 0.97 | 1.17 |
| Bulletproofs AVX2 (HashWires benchmark) | - | 0.938 | - | 6.516 |
| HashWires base 16† | 0.002 | 0.002 | 0.003 | 0.061 |
| HashWires base 256† | 0.009 | 0.01 | 0.083 | 0.194 |
| BFGW + KZG† | 5.653 | 5.682 | 9.572 | 12.569 |
| Groth16-Poseidon† | 4 | 4 | 34.23 | 34.46 |
| Groth16-Pedersen† | 4 | 4 | 31.18 | 33.57 |

The proofs of lattice-based and code-based ZKRPs are concretely quite long, as shown in Table 4. For blockchain applications where one must pay for the space used on-chain, this length is problematic, especially as these constructions do not support aggregation. In order to be competitive with constructions using other techniques shown in Table 4, the proof size must be under 1 KB.

▶ **Research Gap 3.** *Lattice- or code-based ZKRPs with multi-prover aggregation.*

Lattice-based ZKRPs with short proofs are desirable for confidential transactions, as blockchains transition to post-quantum security. In such settings, this size issue may be mitigated by multi-prover aggregation. Each block would then contain only an aggregate range proof for all included transactions. However, this aggregation must be multi-prover as these transactions may be made by many different parties, each holding commitments to private values. Lattice- and code-based ZKRPs with multi-prover aggregation have not yet been constructed, leading us to the this related research gap.

▶ **Research Gap 4.** *Un-replayable credential-based range proofs.*

For credential applications, one might want an interactive range proof that cannot be replayed. Suppose that Alice has a commitment of her age signed by a trusted credential issuer. Alice should be able to visit the DMV and prove in zero knowledge that her committed age is above 16. An observer Bob should not be able to copy Alice's commitment and re-use the transcript of the protocol to prove (possibly falsely) that his age is above 16. If this range proof is non-interactive, Bob can simply copy the proof and re-use it. This re-use might be avoided if the protocol is public-coin interactive, and the DMV issues a random challenge that requires knowledge of the committed value to respond to.

Can we make hash-chain-based range proofs that are un-replayable in this way? As credentials are a primary motivation for HashWires, un-replayability would be a nice property to add.

▶ **Research Gap 5.** *Integer commitments with full soundness with transparent setup.*

CKLR [30] and Sharp [29] construct integer commitments with a relaxed notion of soundness. In order to be used for confidential transactions, they must be augmented with additional proof elements from an RSA group or class group. The RSA version requires a trusted setup, and the class group solution is not compatible with existing optimized libraries. Rather than patching soundness issues by adding these extra elements, it would be preferred to construct practically efficient integer commitments with full soundness and transparent setup.

▶ **Research Gap 6.** *Efficient post-quantum ZKRPs compatible with LWE-based ciphertexts.*

Zero-knowledge range proofs can be used to build verifiable LWE-based encryption schemes as discussed in our full version[27].However, existing verifiable LWE-based encryption schemes constructed using ZKRPs [35, 52] use discrete logarithm-based ZKRPs. Thus, while they obtain privacy against quantum adversaries due to the LWE-based encryption used, they lack soundness in verification due to the DLOG-based range proofs. If there were efficient post-quantum range proofs compatible with LWE-based ciphertexts, one could obtain verifiable encryption with soundness against quantum adversaries as well. While a lattice-based zkSNARK (e.g., [2]) may work in theory, it may not be efficient (yielding long ciphertexts and heavy computation). An efficient lattice-based ZKRP that is compatible with lattice-based encryption would be more satisfactory.

## References


1 Miklós Ajtai. Generating hard instances of lattice problems. In *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing*, pages 99–108, 1996.

2 Martin R Albrecht, Valerio Cini, Russell WF Lai, Giulio Malavolta, and Sri AravindaKrishnan Thyagarajan. Lattice-based snarks: Publicly verifiable, preprocessing, and recursively composable. In *Annual International Cryptology Conference*, pages 102–132. Springer, 2022.

3 Sebastian Angel and Michael Walfish. Verifiable auctions for online ad exchanges. In *Proceedings of the ACM SIGCOMM 2013 conference on SIGCOMM*, pages 195–206, 2013.

4 Thomas Attema, Vadim Lyubashevsky, and Gregor Seiler. Practical product proofs for lattice commitments. In *Annual International Cryptology Conference*, pages 470–499. Springer, 2020.

5 Carsten Baum, Ivan Damgård, Vadim Lyubashevsky, Sabine Oechsner, and Chris Peikert. More efficient commitments from structured lattice assumptions. In *International Conference on Security and Cryptography for Networks*, pages 368–385. Springer, 2018.

6 Mira Belenkiy, Jan Camenisch, Melissa Chase, Markulf Kohlweiss, Anna Lysyanskaya, and Hovav Shacham. Randomizable proofs and delegatable anonymous credentials. In Shai Halevi, editor, *Advances in Cryptology – CRYPTO 2009, 29th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 16-20, 2009. Proceedings*, volume 5677 of *Lecture Notes in Computer Science*, pages 108–125. Springer, 2009. doi:10.1007/978-3-642-03356-8_7.

7 James Bell, Adrià Gascón, Tancrède Lepoint, Baiyu Li, Sarah Meiklejohn, Mariana Raykova, and Cathie Yun. {ACORN}: Input validation for secure aggregation. In *32nd USENIX Security Symposium (USENIX Security 23)*, pages 4805–4822, 2023.

8 Marta Bellés-Muñoz, Miguel Isabel, Jose Luis Muñoz-Tapia, Albert Rubio, and Jordi Baylina. Circom: A circuit description language for building zero-knowledge applications. *IEEE Transactions on Dependable and Secure Computing*, 2022.

9 Eli Ben-Sasson, Iddo Bentov, Yinon Horesh, and Michael Riabzev. Scalable, transparent, and post-quantum secure computational integrity. *Cryptology ePrint Archive*, 2018.

10 Fabrice Benhamouda, Stephan Krenn, Vadim Lyubashevsky, and Krzysztof Pietrzak. Efficient zero-knowledge proofs for commitments from learning with errors over rings. In *European symposium on research in computer security*, pages 305–325. Springer, 2015.

11 D Boneh, B Fisch, A Gabizon, and Z Williamson. A simple range proof from polynomial commitments, 2020. URL: `https://hackmd.io/@dabo/B1U4kx8XI`.

12 Dan Boneh and Matthew Franklin. Efficient generation of shared rsa keys. In *Advances in Cryptology – CRYPTO'97: 17th Annual International Cryptology Conference Santa Barbara, California, USA August 17–21, 1997 Proceedings 17*, pages 425–439. Springer, 1997.

13 Jonathan Bootle, Andrea Cerulli, Pyrros Chaidos, Jens Groth, and Christophe Petit. Efficient zero-knowledge arguments for arithmetic circuits in the discrete log setting. In *Advances in Cryptology–EUROCRYPT 2016: 35th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Vienna, Austria, May 8-12, 2016, Proceedings, Part II 35*, pages 327–357. Springer, 2016.

14 Fabrice Boudot. Efficient proofs that a committed number lies in an interval. In *International Conference on the Theory and Applications of Cryptographic Techniques*, pages 431–444. Springer, 2000.

15 Ernest F Brickell, David Chaum, Ivan B Damgård, and Jeroen van de Graaf. Gradual and verifiable release of a secret. In *Advances in Cryptology – CRYPTO'87: Proceedings 7*, pages 156–166. Springer, 1988.

16 Benedikt Bünz, Shashank Agrawal, Mahdi Zamani, and Dan Boneh. Zether: Towards privacy in a smart contract world. In *International Conference on Financial Cryptography and Data Security*, pages 423–443. Springer, 2020.

17 Benedikt Bünz, Jonathan Bootle, Dan Boneh, Andrew Poelstra, Pieter Wuille, and Greg Maxwell. Bulletproofs: Short proofs for confidential transactions and more. In *2018 IEEE symposium on security and privacy (SP)*, pages 315–334. IEEE, 2018.

18 Benedikt Bünz, Ben Fisch, and Alan Szepieniec. Transparent snarks from dark compilers. In *Advances in Cryptology–EUROCRYPT 2020: 39th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zagreb, Croatia, May 10–14, 2020, Proceedings, Part I 39*, pages 677–706. Springer, 2020.

19 Benedikt Bünz, Ben Fisch, and Alan Szepieniec. Transparent snarks from dark compilers. In *Advances in Cryptology–EUROCRYPT 2020: 39th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zagreb, Croatia, May 10–14, 2020, Proceedings, Part I 39*, pages 677–706. Springer, 2020.

20 Philippe Camacho. Secure protocols for provable security, 2014. URL: `https://www.slideshare.net/philippecamacho/protocols-for-provable-solvency-38501620`.

21 Jan Camenisch, BRICS, and Markus Michels. Separability and efficiency for generic group signature schemes. In *Annual International Cryptology Conference*, pages 413–430. Springer, 1999.

22 Jan Camenisch, Rafik Chaabouni, and Abhi Shelat. Efficient protocols for set membership and range proofs. In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 234–252. Springer, 2008.

23 Jan Camenisch and Markus Michels. Proving in zero-knowledge that a number is the product of two safe primes. In *Advances in Cryptology – EUROCRYPT'99: International Conference on the Theory and Application of Cryptographic Techniques Prague, Czech Republic, May 2–6, 1999 Proceedings 18*, pages 107–122. Springer, 1999.

**24** Konstantinos Chalkias, Shir Cohen, Kevin Lewi, Fredric Moezinia, and Yolan Romailler. Hashwires: Hyperefficient credential-based range proofs. *Proceedings on Privacy Enhancing Technologies*, 4:76–95, 2021.

**25** Agnes Hui Chan, Yair Frankel, and Yiannis Tsiounis. Easy come – easy go divisible cash. In Kaisa Nyberg, editor, *Advances in Cryptology – EUROCRYPT '98, International Conference on the Theory and Application of Cryptographic Techniques, Espoo, Finland, May 31 – June 4, 1998, Proceeding*, volume 1403 of *Lecture Notes in Computer Science*, pages 561–575. Springer, 1998. doi:10.1007/BFb0054154.

**26** Panagiotis Chatzigiannis and Foteini Baldimtsi. Miniledger: Compact-sized anonymous and auditable distributed payments. In Elisa Bertino, Haya Shulman, and Michael Waidner, editors, *Computer Security – ESORICS 2021 – 26th European Symposium on Research in Computer Security, Darmstadt, Germany, October 4-8, 2021, Proceedings, Part I*, volume 12972 of *Lecture Notes in Computer Science*, pages 407–429. Springer, 2021. doi:10.1007/978-3-030-88418-5_20.

**27** Miranda Christ, Foteini Baldimtsi, Konstantinos Kryptos Chalkias, Deepak Maram, Arnab Roy, and Joy Wang. SoK: Zero-knowledge range proofs. Cryptology ePrint Archive, Paper 2024/430, 2024. URL: https://eprint.iacr.org/2024/430.

**28** Heewon Chung, Kyoohyung Han, Chanyang Ju, Myungsun Kim, and Jae Hong Seo. Bulletproofs+: Shorter proofs for a privacy-enhanced distributed ledger. *IEEE Access*, 10:42081–42096, 2022.

**29** Geoffroy Couteau, Dahmun Goudarzi, Michael Klooß, and Michael Reichle. Sharp: Short relaxed range proofs. In *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security*, pages 609–622, 2022.

**30** Geoffroy Couteau, Michael Klooß, Huang Lin, and Michael Reichle. Efficient range proofs with transparent setup from bounded integer commitments. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 247–277. Springer, 2021.

**31** Geoffroy Couteau, Thomas Peters, and David Pointcheval. Removing the strong rsa assumption from arguments over the integers. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 321–350. Springer, 2017.

**32** Gaby G Dagher, Benedikt Bünz, Joseph Bonneau, Jeremy Clark, and Dan Boneh. Provisions: Privacy-preserving proofs of solvency for bitcoin exchanges. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*, pages 720–731, 2015.

**33** Ivan Damgård and Eiichiro Fujisaki. An integer commitment scheme based on groups with hidden order. *Cryptology ePrint Archive*, 2001.

**34** Vanesa Daza, Carla Ràfols, and Alexandros Zacharakis. Updateable inner product argument with logarithmic verifier and applications. In *Public-Key Cryptography–PKC 2020: 23rd IACR International Conference on Practice and Theory of Public-Key Cryptography, Edinburgh, UK, May 4–7, 2020, Proceedings, Part I 23*, pages 527–557. Springer, 2020.

**35** Rafaël Del Pino, Vadim Lyubashevsky, and Gregor Seiler. Short discrete log proofs for fhe and ring-lwe ciphertexts. In *IACR International Workshop on Public Key Cryptography*, pages 344–373. Springer, 2019.

**36** Liam Eagen. Bulletproofs++. *Cryptology ePrint Archive*, 2022.

**37** Muhammed F Esgin, Ron Steinfeld, Joseph K Liu, and Dongxi Liu. Lattice-based zero-knowledge proofs: new techniques for shorter and faster constructions and applications. In *Annual International Cryptology Conference*, pages 115–146. Springer, 2019.

**38** Muhammed F Esgin, Ron Steinfeld, Amin Sakzad, Joseph K Liu, and Dongxi Liu. Short lattice-based one-out-of-many proofs and applications to ring signatures. In *Applied Cryptography and Network Security: 17th International Conference, ACNS 2019, Bogota, Colombia, June 5–7, 2019, Proceedings 17*, pages 67–88. Springer, 2019.

**39** Yair Frankel, Philip D MacKenzie, and Moti Yung. Robust efficient distributed rsa-key generation. In *Proceedings of the thirtieth annual ACM symposium on Theory of computing*, pages 663–672, 1998.

**40** Eiichiro Fujisaki and Tatsuaki Okamoto. Statistical zero knowledge protocols to prove modular polynomial relations. In *Advances in Cryptology – CRYPTO'97: 17th Annual International Cryptology Conference Santa Barbara, California, USA August 17–21, 1997 Proceedings 17*, pages 16–30. Springer, 1997.

**41** Ariel Gabizon, Zachary J Williamson, and Oana Ciobotaru. Plonk: Permutations over lagrange-bases for oecumenical noninteractive arguments of knowledge. *Cryptology ePrint Archive*, 2019.

**42** Craig Gentry, Shai Halevi, and Vadim Lyubashevsky. Practical non-interactive publicly verifiable secret sharing with thousands of parties. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 458–487. Springer, 2022.

**43** Jens Groth. Non-interactive zero-knowledge arguments for voting. In *Applied Cryptography and Network Security: Third International Conference, ACNS 2005, New York, NY, USA, June 7-10, 2005. Proceedings 3*, pages 467–482. Springer, 2005.

**44** Jens Groth. On the size of pairing-based non-interactive arguments. In *Advances in Cryptology–EUROCRYPT 2016: 35th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Vienna, Austria, May 8-12, 2016, Proceedings, Part II 35*, pages 305–326. Springer, 2016.

**45** Jens Groth. Non-interactive distributed key generation and key resharing. *Cryptology ePrint Archive*, 2021.

**46** G.H. Hardy, E.M. Wright, D.R. Heath-Brown, and J. Silverman. *An Introduction to the Theory of Numbers*. Oxford mathematics. OUP Oxford, 2008.

**47** iden3. rapidsnark. `https://github.com/iden3/rapidsnark`, 2023.

**48** Yan Ji and Konstantinos Chalkias. Generalized proof of liabilities. In *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security*, pages 3465–3486, 2021.

**49** Aniket Kate, Gregory M Zaverucha, and Ian Goldberg. Constant-size commitments to polynomials and their applications. In *Advances in Cryptology-ASIACRYPT 2010: 16th International Conference on the Theory and Application of Cryptology and Information Security, Singapore, December 5-9, 2010. Proceedings 16*, pages 177–194. Springer, 2010.

**50** Jonathan Katz and Yehuda Lindell. *Introduction to modern cryptography: principles and protocols*. Chapman and hall/CRC, 2007.

**51** Akinori Kawachi, Keisuke Tanaka, and Keita Xagawa. Concurrently secure identification schemes based on the worst-case hardness of lattice problems. In *Advances in Cryptology-ASIACRYPT 2008: 14th International Conference on the Theory and Application of Cryptology and Information Security, Melbourne, Australia, December 7-11, 2008. Proceedings 14*, pages 372–389. Springer, 2008.

**52** Benoit Libert. Vector commitments with short proofs of smallness. *Cryptology ePrint Archive*, 2023.

**53** Benoît Libert, San Ling, Fabrice Mouhartem, Khoa Nguyen, and Huaxiong Wang. Zero-knowledge arguments for matrix-vector relations and lattice-based group encryption. In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 101–131. Springer, 2016.

**54** Benoît Libert, San Ling, Khoa Nguyen, and Huaxiong Wang. Zero-knowledge arguments for lattice-based accumulators: logarithmic-size ring signatures and group signatures without trapdoors. In *Advances in Cryptology–EUROCRYPT 2016: 35th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Vienna, Austria, May 8-12, 2016, Proceedings, Part II 35*, pages 1–31. Springer, 2016.

**55** Benoît Libert, San Ling, Khoa Nguyen, and Huaxiong Wang. Lattice-based zero-knowledge arguments for integer relations. In *Annual International Cryptology Conference*, pages 700–732. Springer, 2018.

**56** Helger Lipmaa. On diophantine complexity and statistical zero-knowledge arguments. In *Advances in Cryptology-ASIACRYPT 2003: 9th International Conference on the Theory and Application of Cryptology and Information Security, Taipei, Taiwan, November 30–December 4, 2003. Proceedings 9*, pages 398–415. Springer, 2003.

**57** Vadim Lyubashevsky, Ngoc Khanh Nguyen, and Maxime Plançon. Lattice-based zero-knowledge proofs and applications: shorter, simpler, and more general. In *Annual International Cryptology Conference*, pages 71–101. Springer, 2022.

**58** Vadim Lyubashevsky, Ngoc Khanh Nguyen, and Gregor Seiler. Practical lattice-based zero-knowledge proofs for integer relations. In *Proceedings of the 2020 ACM SIGSAC conference on computer and communications security*, pages 1051–1070, 2020.

**59** Greg Maxwell. Confidential transactions, 2016. URL: `https://people.xiph.org/~greg/confidential_values.txt`.

**60** Eduardo Morais, Tommy Koens, Cees van Wijk, and Aleksei Koren. A survey on zero knowledge range proofs and applications. *CoRR*, abs/1907.06381, 2019. `arXiv:1907.06381`.

**61** Khoa Nguyen, Hanh Tang, Huaxiong Wang, and Neng Zeng. New code-based privacy-preserving cryptographic constructions. In *Advances in Cryptology–ASIACRYPT 2019: 25th International Conference on the Theory and Application of Cryptology and Information Security, Kobe, Japan, December 8–12, 2019, Proceedings, Part II 25*, pages 25–55. Springer, 2019.

**62** Valeria Nikolaenko, Sam Ragsdale, Joseph Bonneau, and Dan Boneh. Powers-of-tau to the people: Decentralizing setup ceremonies. In *International Conference on Applied Cryptography and Network Security*, pages 105–134. Springer, 2024.

**63** Torben Pryds Pedersen. Non-interactive and information-theoretic secure verifiable secret sharing. In *Annual international cryptology conference*, pages 129–140. Springer, 1991.

**64** Andrew Poelstra. Mimblewimble, 2016.

**65** Ronald L Rivest and Adi Shamir. Payword and micromint: Two simple micropayment schemes. In *International workshop on security protocols*, pages 69–87. Springer, 1996.

**66** Jacques Stern. A new paradigm for public key identification. *IEEE Transactions on Information Theory*, 42(6):1757–1768, 1996.

**67** Alin Tomescu. Range proofs from polynomial commitments, re-explained, March 2020. URL: `https://decentralizedthoughts.github.io/2020-03-03-range-proofs-from-polynomial-commitments-reexplained/`.

**68** Nan Wang and Sid Chi-Kin Chau. Flashproofs: Efficient zero-knowledge arguments of range and polynomial evaluation with transparent setup. In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 219–248. Springer, 2022.

**69** Nan Wang, Sid Chi-Kin Chau, and Dongxi Liu. Swiftrange: A short and efficient zero-knowledge range argument for confidential transactions and more. *Cryptology ePrint Archive*, 2023.

**70** Pieter Wuille. libsecp256k1, 2018. URL: `https://github.com/bitcoin/secp256k1`.

**71** Rupeng Yang, Man Ho Au, Zhenfei Zhang, Qiuliang Xu, Zuoxia Yu, and William Whyte. Efficient lattice-based zero-knowledge arguments with standard soundness: construction and applications. In *Advances in Cryptology–CRYPTO 2019: 39th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 18–22, 2019, Proceedings, Part I 39*, pages 147–175. Springer, 2019.

**72** Zibo Zhou, Zongyang Zhang, Hongyu Tao, Tianyu Li, and Boyu Zhao. Efficient inner product arguments and their applications in range proofs. *IET Information Security*, 17(3):485–504, 2023.