

Batch Proofs of Partial Knowledge*

Ryan Henry and Ian Goldberg

Cheriton School of Computer Science
University of Waterloo
Waterloo, ON, Canada N2L 3G1
{rhenry, iang}@cs.uwaterloo.ca

Abstract. This paper examines “batch zero-knowledge” protocols for communication- and computation-efficient proofs of propositions composed of many simple predicates. We focus specifically on batch protocols that use Cramer, Damgård, and Schoenmakers’ *proofs of partial knowledge* framework (Crypto 1994) to prove propositions that may be true even when some of their input predicates are false. Our main result is a novel system for batch zero-knowledge arguments of knowledge and equality of k -out-of- n discrete logarithms. Along the way, we propose the first general definition for batch zero-knowledge proofs and we revisit Peng and Bao’s batch zero-knowledge proofs of knowledge and equality of one-out-of- n discrete logarithms (Inscrypt 2008). Our analysis of the latter protocol uncovers a critical flaw in the security proof, and we present a practical lattice-based attack to exploit it.

Keywords: Batch proof and verification, zero-knowledge, cryptanalysis, lattice-based attacks, efficiency.

1 Introduction

An interactive zero-knowledge proof is a conversation between two mutually distrusting parties—a *prover* and a *verifier*—in which the prover tries to convince the verifier that some proposition is true. The prover holds evidence (e.g., an NP witness) but is unwilling to reveal it to the verifier; the verifier, conversely, is skeptical of the prover and needs convincing. What makes a zero-knowledge proof special, therefore, is how much extra information the verifier learns: in a zero-knowledge proof, the verifier learns *nothing* beyond the veracity of the prover’s claim. Zero-knowledge proofs have had profound implications for cryptography since Goldwasser, Micali, and Rackoff introduced them back in 1985 [20]; indeed, they are integral to many cryptographic protocols in the literature ranging from end-to-end verifiable voting schemes [11, 25], through to anonymous blacklisting and reputation systems [2, 3, 33], protocols for priced symmetric private information retrieval [24], threshold ring signatures [34], verifiable mix networks [21, 31], and cryptographic auctions [9], among others.

* An extended version of this paper is available [23].

Alas, zero-knowledge does not come free. Each application above gives rise to at least one proposition whose “fan-in” scales with a critical system parameter (e.g., the global user count [11, 25, 34] or the size of a database [2, 3, 24, 33]). Take for example a prover and a verifier that share generators g, h and a set of n pairs of group elements $\{(g_i, h_i) = (g^{x_i}, h^{y_i}) \mid i \in [1, n]\}$ from some suitably chosen group \mathbb{G} . The prover wishes to prove a proposition about this entire *batch* of predicates, such as “For each $i \in [1, n]$, I know $x_i \in \mathbb{Z}_{|\mathbb{G}|}$ such that $\log_g g_i = \log_h h_i = x_i$ ” or “For some $i \in [1, n]$, I know $x_i \in \mathbb{Z}_{|\mathbb{G}|}$ such that $\log_g g_i = \log_h h_i = x_i$ ”. The first (“AND”) proposition naturally arises, e.g., in universally verifiable shuffling protocols for mix networks [31, §2.3], while the second (“OR”) proposition arises, e.g., in coercion-resistant Internet voting when each voter must prove that she appears on the election roster [11, §3.5]. In both cases, the standard zero-knowledge protocols scale linearly with the number of inputs n : the prover and verifier each compute $\Theta(n)$ full-length exponentiations in \mathbb{G} , the verifier sends $\Theta(1)$ group elements to the prover, and the prover sends $\Theta(n)$ group elements to the verifier. In 1998, Bellare, Garay, and Rabin [4, 5] suggested *batch verification* techniques to reduce verification costs. Their “small exponents” batch verification test [5, §3.3] reduces the verifier’s computation cost to just $\Theta(1)$ full-length exponentiations and $\Theta(\lambda n)$ multiplications in \mathbb{G} . The quantity λ is a *soundness parameter*; perhaps $\lambda = 40$ or 60 in practice.

Inspired by Bellare et al.’s small-exponent batch test, Peng, Boyd, and Dawson [30, §4.1] proposed a four-round *batch proof* of (complete) knowledge for the above “AND” proposition. The prover and verifier each compute just $\Theta(1)$ full-length exponentiations and $\Theta(\lambda n)$ multiplications in \mathbb{G} , the prover sends $\Theta(1)$ group elements to the verifier, and the verifier sends $\Theta(1)$ group elements and $\Theta(\lambda n)$ additional bits to the prover. More recently, Peng and Bao [28, §5.1] proposed a four-round batch proof of *partial knowledge* for the above “OR” proposition, which blends small-exponent batch testing with a special case of Cramer, Damgård, and Schoenmakers’ proofs of partial knowledge [12]. The prover and verifier each compute just $\Theta(1)$ full-length exponentiations and $\Theta(\lambda n)$ multiplications in \mathbb{G} , and they each send and receive just $\Theta(1)$ group elements and $\Theta(\lambda n)$ additional bits. A handful of other papers [8, 15, 21, 24, 29, 31] propose similar “batch proofs” with similar “sublinear” costs.

Our contributions.

1. We propose a novel system for batch zero-knowledge arguments of knowledge and equality of k -out-of- n discrete logarithms for fixed $k \in [1, n]$. As special cases, we obtain batch “AND” proofs (n -out-of- n) and batch “OR” proofs (one-out-of- n). Our protocol has similar costs to the protocols of Peng et al. [30, §4.1] and of Peng and Bao [28, §5.1].
2. We present a practical, lattice-based attack on the soundness of Peng and Bao’s protocol for batch zero-knowledge proofs of knowledge and equality of one-out-of- n pairs of discrete logarithms. We provide a fix that uses *all-but- k mercurial commitments* [22], a variant of mercurial vector commitments [10] with a strengthened binding property.

3. We propose formal definitions for *batch zero-knowledge* proofs and proofs of knowledge. Prior treatments of batch proofs have been informal and ad hoc. Our new definitions address this with a *conciseness* property describing the asymptotic performance of a “batch” protocol relative to one formed by sequential composition of corresponding single-instance protocols.

Outline. We examine Peng and Bao’s [28, §5.1] batch “OR” protocol in §2 and describe a practical attack on its soundness (the full details of which are in the extended version of this paper [23, Appendix A]). In §3, we discuss using *all-but-k mercurial commitments* [22] to repair Peng and Bao’s protocol. In §4, we propose our formal definitions for batch zero-knowledge proofs and batch zero-knowledge proofs of knowledge. Our new batch protocol follows in §5. We list some potential applications in §6 and conclude in §7.

2 Peng and Bao’s Batch Proof Protocol

The following protocol is due to Peng and Bao [28, §5.1]; they call it “batch ZK proof and verification of 1-out-of- n equality of logarithms”. The protocol incorporates Bellare et al.’s small-exponent batch testing [5, §3.3] into both the proof and verification phases of an otherwise standard sigma protocol for proving knowledge and equality of one-out-of- n pairs of discrete logarithms. Both the prover P and verifier V know the same two generators g, h of an order- p group \mathbb{G} and a set of n pairs of group elements $\{(g_i, h_i) \mid i \in [1, n]\}$, but only P knows an index $j \in [1, n]$ and exponent $x_j \in \mathbb{Z}_p^*$ such that $\log_g g_j = \log_h h_j = x_j$. The goal of the protocol is for P to convince V that she knows such a (j, x_j) pair without revealing any additional information. For ease of notation below, we define $H = [1, n] \setminus \{j\}$ for the (j, x_j) pair that honest P is proving knowledge of. We also introduce a *soundness parameter* $\lambda \in \mathbb{N}$, which tunes the cost versus soundness trade off in small-exponent batch testing.

Protocol 1. (Peng & Bao’s Batch Proof of Partial Knowledge [28, §5.1]).

- V1:** Choose $t_i \in_{\mathbb{R}} [0, 2^\lambda - 1]$ for each $i \in [1, n]$. Send (t_1, \dots, t_n) to P .
- P2:** Receive (t_1, \dots, t_n) from V . Choose $r \in_{\mathbb{R}} \mathbb{Z}_p^*$ and $c_i \in_{\mathbb{R}} [0, 2^\lambda - 1]$ for each $i \in H$. Compute $a = g^r \prod_{i \in H} g_i^{c_i t_i}$ and $b = h^r \prod_{i \in H} h_i^{c_i t_i}$. Send (a, b) to V .
- V3:** Receive (a, b) from P . Choose $c \in_{\mathbb{R}} [0, 2^\lambda - 1]$ and send it to P .
- P4:** Receive c from V . Compute $c_j = c - \sum_{i \in H} c_i \pmod{2^\lambda}$ and $v = r - t_j c_j x_j \pmod{p}$. Send (c_1, \dots, c_n, v) to V .
- V5:** Receive (c_1, \dots, c_n, v) from P . Output “true” if and only if $a \stackrel{?}{=} g^v \prod_{i=1}^n g_i^{c_i t_i}$, $b \stackrel{?}{=} h^v \prod_{i=1}^n h_i^{c_i t_i}$, and $c \stackrel{?}{=} \sum_{i=1}^n c_i \pmod{2^\lambda}$, otherwise output “false”.

Some remarks about Protocol 1 are in order. Perhaps surprisingly, we observe that V speaks before P does. What V sends to P in Step V1 is a list of short exponents for small-exponent batch testing. Step P2 ostensibly forces P to commit to an index j (such that $H = [1, n] \setminus \{j\}$) and to $\{c_i \mid i \in H\}$; if so, then V choosing

$c \in_{\mathbb{R}} [0, 2^\lambda - 1]$ in Step V3 is equivalent to V choosing the missing $c_j \in_{\mathbb{R}} [0, 2^\lambda - 1]$ for P in Step P4, which is exactly what we want for good soundness. It is trivial to verify that the protocol is *complete*; i.e., that honest P always convinces honest V. Theorem 1 in Peng and Bao’s paper [28] states that “Soundness in [the above protocol] only fails with an overwhelmingly small probability [in the soundness parameter λ].” Their soundness proof works by computing an upper bound of $1/2^\lambda$ on the probability that the verification equations hold if $\log_g g_j \neq \log_h h_j$, given the following two implicit assumptions: 1) P committed to $H = [1, n] \setminus \{j\}$ and to $\{c_i \mid i \in H\}$ in Step P2, and 2) V chose c and (hence, c_j) uniformly at random from $[0, 2^\lambda - 1]$ in Step V3. *However, it is easy to see that the pair (a, b) of “commitments” that P computes and sends to V in Step P2 does not bind her to using $H = [1, n] \setminus \{j\}$; hence, the first implicit assumption in Peng and Bao’s soundness proof is not guaranteed to hold when P is dishonest.* Dishonest P can exploit this observation to pass the verification equations even when the claimed equality of logarithms is false.

We give a high-level description of the attack below; interested readers can find further details in the extended version of this paper [23, Appendix A].

Overview of the attack. Suppose that P knows several (x_j, y_j) pairs such that $(g_j, h_j) = (g^{x_j}, h^{y_j})$ but $x_j \not\equiv y_j \pmod{p}$ for any of the known pairs. Partition the interval $[1, n]$ into two sets H and S , where S is a subset of indices for which P knows the above pair of discrete logarithms and H is a superset of indices for which she does not. (Note that in some reasonable settings P may know *every* such pair.) In Step P2, P computes (a, b) using this new H so that when V sends c to P in Step V3, P has $|S| - 1$ extra degrees of freedom to compute her response in Step P4. In particular, to find the missing $\{c_j \mid j \in S\}$ she solves the following system of two linear equations in $k = |S|$ unknowns:

$$0 \equiv \sum_{j \in S} c_j t_j (x_j - y_j) \pmod{p}, \text{ and} \tag{1}$$

$$c' \equiv \sum_{j \in S} c_j \pmod{2^\lambda}, \tag{2}$$

where $c' = c - \sum_{i \in H} c_i \pmod{2^\lambda}$. Equation (1) implies $\sum_{j \in S} c_j t_j x_j \equiv \sum_{j \in S} c_j t_j y_j \pmod{p}$; hence, if P sets $v = r - \sum_{j \in S} c_j t_j x_j \pmod{p}$ in Step P4, then (c_1, \dots, c_n, v) will satisfy each verification equation in Step V5. Of course, if P just naively solves the above system of equations and obtains a solution $\{c_j \mid j \in S\}$ containing $c_{j'} \geq 2^\lambda$ for some $j' \in S$, then V may notice that P is cheating. Therefore, what P really wants to do is find a solution to the above system subject to the additional restriction that $0 \leq c_j < 2^\lambda$ for all $j \in S$.

A counting argument suggests that such “suitably small” solutions are plentiful whenever $k \cdot \lambda$ is “sufficiently large” compared to $\lg p$.¹ If \mathbf{X} is an instance

¹ Recall that $k = |S|$ is a lower bound on the number of exponent pairs that P knows and that λ is the soundness parameter. Larger values of λ are *supposed* to result in better soundness; however, what we find is just the opposite: larger values of λ only make suitably small solutions more numerous and easier for P to find.

of the above system induced by some real interaction between P and honest V, then we heuristically expect the distribution of solutions of \mathbf{X} to be uniform among all possible $\langle c_{j_1}, \dots, c_{j_k} \rangle \in (\mathbb{Z}_p)^k$; in particular, we expect the proportion of solutions that are suitably small to be about $(2^\lambda/p)^k$. Now, only p^{k-1} of the $\langle c_{j_1}, \dots, c_{j_k} \rangle \in (\mathbb{Z}_p)^k$ can satisfy Equation (1), and of these only about $p^{k-1}/2^\lambda$ can simultaneously satisfy Equation (2). This leads us to conclude that \mathbf{X} has around $(p^{k-1}/2^\lambda)(2^\lambda/p)^k = (2^\lambda)^{k-1}/p$ suitably small solutions. In the extended version of this paper [23, Appendix A], we discuss how P can find one of these solutions by solving a short vector search problem in a particular lattice of dimension $k+3$. When k is reasonably small, P can use a standard basis reduction algorithm, such as Lenstra-Lenstra-Lovász (LLL) [27], to find a suitably small solution quickly. For example, setting $\lambda = 40$ and letting $\lg p \approx 160$, P only needs to know about $k = 5$ exponent pairs to find a suitably small solution, on average.

3 All-but- k Mercurial Commitments

Our attack on Protocol 1 is possible because P can wait until *after* she sees the challenge c in Step P4 to choose $k > 1$ of the c_i . If the “commitment” in Step P2 actually bound P to using $H = [1, n] \setminus \{j\}$ and $\{c_i \mid i \in H\}$, then Peng and Bao’s upper bound of $1/2^\lambda$ on the protocol’s soundness error would hold. For a direct fix, we therefore desire a special commitment that will (i) force P to commit to all but one component of $\langle c_1, \dots, c_n \rangle$ in Step P2 and (ii) let P specify an arbitrary value for the missing component—*without betraying its position*—when she opens the commitment in Step P4. This, informally, is the binding and hiding guarantees that all-but- k mercurial commitments [22] provide when $k = 1$. More generally, an all-but- k mercurial commitment allows P to commit to an *arbitrary subset* of $n - k$ components from a length- n vector so that she is bound to these $n - k$ components but is still free to choose the k unspecified components prior to opening. V does not learn which components P chose before committing and which she chose after committing; V does, however, learn the total number ‘ k ’ of non-committed components in the opening.

We refer the reader to Henry and Goldberg’s paper [22] for a more comprehensive exposition of all-but- k mercurial commitments, including formal statements of the security properties. For our own purposes, we use an abridged notation that abstracts away certain technical details.

Informal definition. An *all-but- k mercurial commitment scheme* is a 4-tuple of probabilistic polynomial-time (PPT) algorithms (ABK-Init, ABK-Commit, ABK-Open, ABK-Verify) that work as follows:

- ABK-Init outputs a common reference string PK for use in the other protocols.
- ABK-CommitPK outputs commitments to vectors in which some subset of components is as-of-yet unspecified.

- ABK-OpenPK opens such commitments to fully specified vectors, explicitly revealing the number of components k not bound by the commitment.
- ABK-VerifyPK verifies the output of ABK-OpenPK, including the validity of k .

Repairing Peng and Bao’s protocol. Given secure all-but- k mercurial commitments, it is straightforward to protect Protocol 1 from attacks like the one in §2. In Step P2, P commits to $\langle c_i \mid i \in H \rangle$ for $H = [1, n] \setminus \{j\}$. After V sends c to P in Step V3, P computes the missing c_j as usual, then opens the above commitment to $\langle c_1, \dots, c_n \rangle$ as part of Step P4, proving as she does so that she chose only one of the c_j after committing in Step P2. Constructing a simulator and extractor for this augmented protocol is simple (and we give explicit simulator and extractor constructions for the generalized version in the extended version of this paper [23, Appendix B]); furthermore, the augmented protocol is still intuitively a “batch” protocol provided the all-but- k scheme satisfies certain efficiency requirements. In §5, we let the parameter k vary and thereby generalize the repaired Peng-Bao protocol to a system for batch zero-knowledge arguments of knowledge and equality of k -out-of- n discrete logarithms for any $k \in [1, n]$. Our protocol (including the special case just outlined) appears to be the first such batch protocol for $k \neq n$.

4 Defining Batch Zero-Knowledge Proofs

Several papers (many of which we listed in the introduction [8, 15, 21, 24, 29, 31]) propose protocols that implement what their respective authors refer to as “batch zero-knowledge proofs (of knowledge)”. Regrettably, the community has no agreed upon definition of what constitutes a “batch” zero-knowledge proof. Prior works, consequently, justify the terminology using ad hoc arguments that contrast the communication cost (counted in terms of group elements transfers) and computation cost (counted in terms of full-length exponentiations) of their protocols with those of the most “obvious” protocols to implement proofs of the same propositions. (Peng et al. did suggest one definition for batch zero-knowledge proofs [30, Definition 1]; however, their definition fails to address asymptotic communication and computation costs, which we believe to be the key property differentiating the abovementioned “batch” proofs from their “non-batch” counterparts.) We therefore offer our own, very general definition for batch zero-knowledge proofs (of knowledge). We model our new definition after the standard zero-knowledge definitions (specifically, [18, Definition 3] and [6, Definition 3.1]), but add a new parameterized *conciseness* criterion that places asymptotic restrictions on how the communication and the computation costs of the interaction scale with respect to the size of the proposition under consideration. In particular, our conciseness criterion characterizes the asymptotic relationship between the number of predicates under consideration, the *soundness (or knowledge) error* of the proof, and the communication and computation cost of the resulting interaction.

Formal model. We model our prover P and verifier V as a pair of interactive functions and consider the interaction $(P_x(y), V_x(z))$ that occurs when both functions take $x = \langle x_1, \dots, x_n \rangle$ as common input, P takes $y = \langle y_1, \dots, y_n \rangle$ as private input, and V takes string z as private auxiliary input. In general, some (possibly trivial) subset of (x_i, y_i) pairs satisfy a given witness relation R and z encodes arbitrary prior knowledge of V, such as a set of transcripts from earlier interactions with P. (The *transcript* of an interaction $(P_x(\cdot), V_x(z))$, which is denoted by $\text{tr}_{P,V}(x, z)$, is the string-valued random variable that records V’s inputs and all correspondence with P up to the end of an interaction.)

We let φ_R be the function that maps pairs of n -tuples (x, y) as above to n -bit strings in which the i^{th} bit is 1 if and only if $(x_i, y_i) \in R$. Intuitively, we are interested in interactions $(P_x(y), V_x(z))$ that implement zero-knowledge proofs of the proposition $p(x, y)$ induced by R and a given language L in the following sense: $p(x, y)$ is true if and only if $\varphi_R(x, y) \in L$. Note that the pair (L, R) uniquely determines the proposition p , and vice versa. For ease of notation below, we define the language of n -tuples induced by (L, R) as $L_R = \{x \mid \exists y \text{ for which } \varphi_R(x, y) \in L\}$. We parameterize the lengths of the x_i and y_i in a given interaction by τ ; in particular, we assume throughout that the x_i are all τ bits long and the y_i are all $\text{poly}(\tau)$ bits long, where $\text{poly}(\cdot)$ is some fixed polynomial. Let $T = \{(\{0, 1\}^\tau)^n \mid \tau, n \in \mathbb{N}_{\geq 1}\}$ denote set of n -tuples of fixed-length strings.

Example. For propositions $p(x, y)$ asserting knowledge and equality of discrete logarithms, as in the actual protocols we consider in this paper, $R = \{(x_i, y_i) = (g_i, h_i), y_i) \in \mathbb{G}^2 \times \mathbb{Z}_p \mid \log_g g_i = \log_h h_i = y_i\}$. If $p(x, y)$ is the “AND” proposition, then L is the language of strings comprised entirely of 1s; if $p(x, y)$ is the “OR” proposition, then L is the language of strings with nonzero Hamming weight. For our own k -out-of- n proofs, L is the language of strings with Hamming weight at least k . Note that in general P is proving *partial knowledge* of witnesses for R , with the strings in L reflecting which subsets of witnesses P might actually know.

We now recall the standard notions of a *simulator* for verifier V, which we use to formalize what it means for $(P_x(y), V_x(z))$ to be “zero-knowledge”, and of a *knowledge extractor* for P, which we use to formalize what it means for $(P_x(y), V_x(z))$ to be a “proof of knowledge”.

Definition 1. A probabilistic function S_{V^*} is a *simulator* for verifier-language pair (V^*, L_R) if the probability ensembles $\{\text{tr}_{P,V^*}(x, z)\}_{x \in L_R}$ and $\{S_{V^*}(x, z)\}_{x \in L_R}$ are (computationally, statistically, or perfectly [17, Definitions 3,4]) indistinguishable, where $S_{V^*}(x, z)$ is the string-valued random variable describing the output of S_{V^*} on input (x, z) .

An *oracle machine* for P^* is a function E^{P^*} that is endowed with *rewinding black box oracle access* to P^* . In other words, E^{P^*} is able to 1) submit arbitrary challenges to P^* and get truthful responses in a single time step, and 2) “rewind” P^* to a previous state to get several responses for the same input and random coin flips but different challenges. (Note that E^{P^*} is generally *not* privy to P^* ’s inputs or internal state.)

Definition 2. Let $\kappa: T \rightarrow [0, 1]$ and let $q(x)$ denote the probability that V outputs “true” in $(P_x^*(y), V_x(z))$. An oracle machine E^{P^*} is a *knowledge extractor* (with knowledge error $\kappa(\cdot)$) for the prover-language pair (P^*, L_R) if there exists a positive polynomial $g(\cdot)$ such that, for all n -tuples $x \in L_R$, if $q(x) \geq \kappa(x)$ then, with probability at least $\frac{q(x) - \kappa(x)}{g(|x|)}$, $E^{P^*}(x)$ outputs an n -tuple y' for which $\varphi_R(x, y') \in L$.

Given the above definitions, we formally define what it means for a pair of interactive functions to implement a system of *batch zero-knowledge proofs* or a system of *batch zero-knowledge proofs of knowledge* for the language-relation pair (L, R) . What sets our Definition 3 apart from the standard zero-knowledge definitions is that we include an explicit *conciseness* condition, which characterizes the cost of proving $\varphi_R((x_1, \dots, x_n), (y_1, \dots, y_n)) \in L$ in terms of the cost of proving $(x_1, y_1) \in R$.

For example, consider an interactive protocol A between P and V in which, on common input $x_0 \in \{0, 1\}^\tau$, P convinces V that there exists (or, perhaps, that it “knows”) some y_0 such that $(x_0, y_0) \in R$. Let $a_0(\tau)$ and $a_1(\tau)$ respectively denote the computation cost (for both P and V) and the bidirectional communication cost of A . Now, consider a second interactive protocol B between P and V in which, on common input $x \in (\{0, 1\}^\tau)^n$, P convinces V that there exists (or it “knows”) some y such that $\varphi_R(x, y) \in L$. Let $b_0(\tau, n)$ and $b_1(\tau, n)$ respectively denote the computation cost (for both P and V) and the bidirectional communication cost of B . For a fixed pair of constants $\alpha, \beta \in [0, 1]$, we say that B is (α, β) -concise if there exists a constant $\delta > 0$ such that, for all $\epsilon > 0$, we have

$$b_0(\tau, n) \in O\left(n^\alpha a_0(\tau) + n^{\beta+\epsilon} \tilde{a}_0(\tau)\right) \text{ for some } \tilde{a}_0(\tau) \in o\left(a_0(\tau)^{1-\delta}\right),$$

and

$$b_1(\tau, n) \in O\left(n^\alpha a_1(\tau) + n^{\beta+\epsilon} \tilde{a}_1(\tau)\right) \text{ for some } \tilde{a}_1(\tau) \in o\left(a_1(\tau)^{1-\delta}\right).$$

(That is, the computation cost and communication cost of B grow no faster than n^α times the corresponding cost of A plus at most about n^β times some function that grows at least polynomially slower than the corresponding cost of A as τ grows large. The ϵ and δ factors are present so that we may ignore the contribution of polylogarithmic terms.) Recalling that $\alpha, \beta \in [0, 1]$, we call B a *batch proof (of knowledge)* for (L, R) if it is (α, β) -concise for any $\alpha < 1$, or, very roughly, if the cost of the protocol grows slower than n times the cost of A as we let both n and τ tend to infinity.

Example.

1. Consider Peng et al.’s protocol for proofs of complete knowledge [30, §4.1], our repaired version of Peng and Bao’s protocol in §3 for proofs of partial knowledge, and the forthcoming protocol in §5 for proofs of partial knowledge. In each, we find that $a_0(\tau) \in \Omega(\tau^2 \lg^2 \tau)$ and $a_1(\tau) \in \Omega(\tau)$ while, for all $\epsilon > 0$ and for any soundness parameter $\lambda \in \mathbb{N}$, we find that $b_0(\tau, n) \in O(a_0(\tau) + n^{1+\epsilon} \lambda \tau \lg^2 \tau)$ and $b_1(\tau, n) \in O(a_1(\tau) + n \lambda)$.² For any fixed $\delta < 1/2$, we have that $\tilde{a}_0(\tau) =$

² We assume here that multiplication in \mathbb{G} requires $O(\tau \lg \tau \lg \lg \tau) \in O(\tau \lg^2 \tau)$ bit operations using the Fast Fourier Transform (FFT) [14].

$\lambda \tau \lg^2 \tau \in o(a_0(\tau)^{1-\delta})$ and $\tilde{a}_1(\tau) = \lambda \in o(a_1(\tau)^{1-\delta})$, and therefore each of these protocols is $(0, 1)$ -concise; moreover, since $\alpha = 0 < 1$, each protocol satisfies our conciseness criterion for a batch proof of knowledge.

2. Brands, Demuyck, and De Decker describe a protocol [8, §3.4] to prove that a given commitment commits to a different value than every other commitment on a list. As in the previous example, we find that $a_0(\tau) \in \Omega(\tau^2 \lg^2 \tau)$ and $a_1(\tau) \in \Omega(\tau)$ but, in this case, we have $b_0(\tau, n) \in O(n^{1/2} \tau^2 \lg^2 \tau)$ and $b_1(\tau, n) \in O(n^{1/2} \tau)$. Thus, letting $\delta = 1/2$ and letting $\tilde{a}_0(\tau)$ and $\tilde{a}_1(\tau)$ be arbitrary constant functions, we see that Brands et al.’s protocol is $(\frac{1}{2}, 0)$ -concise; moreover, since $\alpha = \frac{1}{2} < 1$, it satisfies our conciseness criterion for a batch proof.

Definition 3. (System of batch zero-knowledge proofs (of knowledge)).

Let $\Lambda: \mathbb{N} \rightarrow \mathbb{N}$ be a nondecreasing *soundness function* and let $\alpha, \beta \in [0, 1]$ be constants such that $\alpha < 1$. An interactive protocol $(P_x(y), V_x(z))$ is a *system of (α, β, Λ) -batch zero-knowledge proofs* for the language-relation pair (L, R) if there exists a negligible function $\varepsilon_0: \mathbb{N} \rightarrow \mathbb{R}$ for which $(P_x(y), V_x(z))$ satisfies each of the following four conditions.

- 1. Complete:** For any $n \in \mathbb{N}$ and pair (x, y) such that $\varphi_R(x, y) \in L$, if y is input to honest P and x is input to P and honest V, then V outputs “true”.
- 2. (Unconditionally) sound:** For every (possibly malicious) prover P^* , $\tau \in \mathbb{N}$, $n \in \mathbb{N}$, and $x \in (\{0, 1\}^\tau)^n \setminus L_R$, if P^* and honest V receive x as common input then, with probability at least $1 - \varepsilon_0(\Lambda(\tau))$, V outputs “false”.
- 3. (General) zero-knowledge:** For every (possibly malicious) PPT verifier V^* , there exists a PPT simulator S_{V^*} for (V^*, L_R) .
- 4. (α, β) -concise:** If $a_0(\tau)$ and $a_1(\tau)$ respectively denote the computation and communication cost of $(P_x(y), V_x(z))$ when n is fixed as 1, then there exists some constant $\delta > 0$ and functions $\tilde{a}_0(\tau) \in o(a_0(\tau)^{1-\delta})$ and $\tilde{a}_1(\tau) \in o(a_1(\tau)^{1-\delta})$ such that, for every $\epsilon > 0$, we have that
 - a.** for every (possibly malicious) PPT verifier V^* , $\tau \in \mathbb{N}$, and pair (x, y) such that $\varphi_R(x, y) \in L$, if y is input to honest P and x is input to P and V^* , then P runs in $O(n^\alpha a_0(\tau) + n^{\beta+\epsilon} \tilde{a}_0(\tau))$ time and sends $O(n^\alpha a_1(\tau) + n^{\beta+\epsilon} \tilde{a}_1(\tau))$ bits to V; and
 - b.** for every (possibly malicious) prover P^* , $\tau \in \mathbb{N}$, and n -tuple x , if x is input to P^* and honest V, then V runs in $O(n^\alpha a_0(\tau) + n^{\beta+\epsilon} \tilde{a}_0(\tau))$ time and sends $O(n^\alpha a_1(\tau) + n^{\beta+\epsilon} \tilde{a}_1(\tau))$ bits to P^* .

If $(P_x(y), V_x(z))$ additionally satisfies the following condition, then it is a system of (α, β, Λ) -batch zero-knowledge proofs *of knowledge* for (L, R) .

- 5. (Unconditionally) knowledge extractable:** There exists an oracle machine E and function $\kappa: T \rightarrow [0, 1]$ such that, for every (possibly malicious) prover P^* , E^{P^*} is an expected PPT knowledge extractor for (P^*, L_R) with knowledge error $\kappa(\cdot) \leq \varepsilon_0(\Lambda(\tau))$.

We also consider the following two (standard) relaxations of the above definition. First, if $(P_x(y), V_x(z))$ satisfies Conditions 1, 2, 4 (and 5) as stated above, but it only satisfies the weaker Condition 3b as stated below (instead of Condition 3),

then it is a system of *honest-verifier* (α, β, Λ) -batch zero-knowledge proofs (of knowledge) for (L, R) .

3b. (Honest-verifier) zero-knowledge: There exists a PPT simulator S_V for (V, L_R) , where V is the *honest verifier*.

If there exists a negligible function $\varepsilon_1: \mathbb{N} \rightarrow \mathbb{R}$ for which $(P_x(y), V_x(z))$ satisfies Conditions 1, 3[b], and 4 as stated above, but only satisfies the weaker Conditions 2b (and 5b) as stated below, then it is a computationally convincing system of [honest-verifier] (α, β, Λ) -batch zero-knowledge *arguments* (of knowledge) for (L, R) .

2b. (Computationally) sound: For every (possibly malicious) PPT prover P^* , there exists a constant τ_0 such that, for every $\tau > \tau_0$ and $n \in \mathbb{N}$, if P^* and honest V receive $x \in (\{0, 1\}^\tau)^n \setminus L_R$ as common input then, with probability at least $1 - \varepsilon_0(\Lambda(\tau)) - \varepsilon_1(\tau)$, V outputs “false”.

5b. (Computationally) knowledge extractable: There exists an oracle machine E^{P^*} and function $\kappa: T \rightarrow [0, 1]$ such that, for every (possibly malicious) PPT prover P^* , there exists a constant τ_0 such that, for every $\tau > \tau_0$ and $n \in \mathbb{N}$, E^{P^*} is an expected PPT knowledge extractor for (P^*, L_R) with knowledge error $\kappa(\cdot) \leq \varepsilon_0(\Lambda(\tau)) + \varepsilon_1(\tau)$.

5 Batch Proof of Knowledge and Equality of k -out-of- n Pairs of Discrete Logarithms

Our new protocol draws inspiration from the repaired version of Peng and Bao’s protocol outlined in §3, but it improves on that protocol by letting k vary in the all-but- k mercurial commitments, which allows us to prove a more general class of propositions. More precisely, the new protocol generalizes from a system for proofs of knowledge and equality of one-out-of- n pairs of discrete logarithms to a system for arguments of knowledge and equality of k -out-of- n pairs of discrete logarithms for any $k \in [1, n]$. We defer a formal security analysis of the new protocol to the extended version of this paper [23, Appendix B] wherein we prove that, for any fixed k and soundness parameter $\lambda \in \mathbb{N}$, it is a system for honest-verifier $(0, 1, \min\{\tau, \lambda\})$ -batch zero-knowledge proofs of knowledge (in the sense of Definition 3). The latter analysis uses efficiency characteristics of the underlying construction for all-but- k mercurial commitments [22]. Using Henry and Goldberg’s all-but- k mercurial commitment scheme [22], which is computationally binding under the n -Strong Diffie-Hellman assumption [7, §3], yields a system of honest-verifier batch zero-knowledge arguments. We note that in this particular instantiation, the prover is assumed to be computationally bounded not in the bit-length τ but in the security parameter for the all-but- k mercurial commitments. Standard tricks from the literature [19] can relax the honest-verifier assumption at only a small cost to efficiency. Swapping in unconditionally binding all-but- k mercurial commitments would yield a system of proofs rather than arguments.

Table 1. This table compares the communication cost (in bits) and the computation cost (in τ -bit multiplications) for four different protocols that each implement honest-verifier zero-knowledge proofs of knowledge and equality of k -out-of- n pairs of discrete logarithms for some k in a group with τ -bit order. The “**Concise**” column indicates the conciseness of the protocol (in the sense of Definition 3); the “**Batch?**” column indicates if the protocol satisfies our definition of a batch proof; the “**k-out-of-n**” column lists values of k that the protocol supports; the “**Sound?**” column indicates if the protocol achieves overwhelming soundness in the soundness parameter λ . Note that $\lambda = \tau$ in the protocol by Cramer et al.; for the other protocols, typically $\lambda \ll \tau$ and λ is fixed as the smallest value yielding a palatable soundness error.

| Protocol | Communication | Computation | Concise | Batch? | k-out-of-n | Sound? |
|--------------------|----------------------------|----------------------------------|---------|--------|----------------|--------|
| Cramer et al. [12] | $\Theta(\tau n)$ | $\Theta(\tau n)$ | (1, 0) | ✗ | $k \in [1, n]$ | ✓ |
| Peng-Bao [28] | $\Theta(\tau + \lambda n)$ | $\Theta(\tau + \lambda n)$ | (0, 1) | ✓ | $k = 1$ | ✗ |
| Peng et al. [30] | $\Theta(\tau + \lambda n)$ | $\Theta(\tau + \lambda n)$ | (0, 1) | ✓ | $k = n$ | ✓ |
| This work | $\Theta(\tau + \lambda n)$ | $\Theta(\tau + \lambda n \lg n)$ | (0, 1) | ✓ | $k \in [1, n]$ | ✓ |

Table 1 compares the cost of our protocol and those arising from a naive application of Cramer et al.’s framework [12], Peng and Bao’s protocol [28], and Peng et al.’s protocol [30]. The latter three protocols are all systems for *proofs* of knowledge; ours is a system for *arguments* of knowledge. Observe that Peng et al.’s protocol is both sound and a batch protocol, *but it only handles the simple $k = n$ case*, and that Peng and Bao’s protocol is a batch protocol and handles the interesting $k = 1$ case, *but it is not sound*. Cramer et al.’s framework is sound and handles every $k \in [1, n]$, *but it is not a batch protocol*.

5.1 The Protocol

Suppose that **ABK** = (ABK-Init, ABK-Commit, ABK-Open, ABK-Verify) is a secure all-but- k mercurial commitment scheme. Fix a soundness parameter $\lambda \in \mathbb{N}$ and use ABK- to generate a common reference string PK. Protocol 2 implements a system for batch zero-knowledge proofs or arguments of knowledge and equality of k -out-of- n pairs of discrete logarithms for any pair of nonnegative integers (k, n) with $k \leq n$ and $n \leq n_0$. In the protocol, we use $\mathbf{V}_q^{n \times k}$ to denote the column-wise $n \times k$ *rectangular Vandermonde matrix* with entries reduced modulo q :

$$\mathbf{V}_q^{n \times k} = \begin{bmatrix} 1 & 1 & 1 & \dots & 1 \\ 1 & 2 & 2^2 & \dots & 2^k \\ 1 & 3 & 3^2 & \dots & 3^k \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & n & n^2 & \dots & n^k \end{bmatrix} \pmod q.$$

Note that because q is prime with $n < q$ and $k \leq n$, every subset of k rows of $\mathbf{V}_q^{n \times k}$ has full rank and thus forms a non-singular (i.e., invertible) square matrix modulo q . If desired, one could replace $\mathbf{V}_q^{n \times k}$ with any other matrix that has this property when all arithmetic is modulo q .

The setting for the new protocol is similar to before. Both P and V know the same two generators g, h of an order- p group \mathbb{G} , the above-generated all-but- k reference string PK, and a set of $n \in [1, n_0]$ pairs of group elements $\{(g_i, h_i) \mid i \in [1, n]\}$, but only P knows a size- k subset $S \subseteq [1, n]$ of indices and corresponding set $x_S = \{x_j \in \mathbb{Z}_p \mid j \in S\}$ of exponents such that $\log_g g_j = \log_h h_j = x_j$ for all $j \in S$. The goal of the protocol is for P to convince V that she knows such a (S, x_S) pair without revealing any additional information. For ease of notation below, we let $H = [1, n] \setminus S$ for the (S, x_S) pair that honest P is proving knowledge of.

Intuitively, our k -out-of- n proof replaces the all-but-one mercurial commitment from the repaired Peng-Bao proof with an all-but- k mercurial commitment. P commits to $\{c_i \mid i \in H\}$ in Step P2, thus assuring V that she can choose at most $k = |S|$ of the c_i after V sends the challenge in Step V3. Rather than challenge P to produce c_i that sum to c modulo q , V challenges P to produce c_i that obey a system of k non-degenerate linear constraints induced by $\mathbf{V}_q^{n \times k}$ and the k free components in $\langle c_1, \dots, c_n \rangle$. V verifies the constraints in Step V5 by checking if $\langle c_1, \dots, c_n \rangle \cdot \mathbf{V}_q^{n \times k} \stackrel{?}{\equiv} \langle c, 0, \dots, 0 \rangle \pmod{q}$; the all-but- k commitment ensures that P chose all but k of the c_i before she received c . This assures V that c uniquely determined a size- k subset of the c_i , although he learns no information about which subset. From here, P essentially uses Peng et al.’s batch “AND” proof for the size- k subset she is proving knowledge of, and “simulates” the proof for the remaining $n - k$ predicates, as in a standard proof of partial knowledge.

Protocol 2. (Generalized batch proof of partial knowledge).

- V1:** Choose $t_i \in_{\mathbb{R}} [0, 2^\lambda - 1]$ for each $i \in [1, n]$. Send (t_1, \dots, t_n) to P.
- P2:** Receive (t_1, \dots, t_n) from V. Choose $r \in_{\mathbb{R}} \mathbb{Z}_p^*$ and $c_i \in_{\mathbb{R}} [0, q - 1]$ for $i \in H$. Compute $a = g^r \prod_{i \in H} g_i^{c_i t_i}$, $b = h^r \prod_{i \in H} h_i^{c_i t_i}$, and $\mathcal{C} \leftarrow \text{ABK-PR}(\langle c_1, \dots, c_n \rangle)$. Send (a, b, \mathcal{C}) to V.
- V3:** Receive (a, b, \mathcal{C}) from P. Choose $c \in_{\mathbb{R}} [0, q - 1]$ and send it to P.
- P4:** Receive c from V. Solve for $\mathbf{c} = \langle c_1, \dots, c_n \rangle \in \mathbb{Z}_q^n$ such that $\mathbf{c} \cdot \mathbf{V}_q^{n \times k} \equiv \langle c, 0, \dots, 0 \rangle \pmod{q}$, then compute $v = r - \sum_{j \in S} c_j t_j x_j \pmod{p}$ and $\omega \leftarrow \text{ABK-PR}(\mathcal{C}, k, \mathbf{c})$. Send (\mathbf{c}, v, ω) to V.
- V5:** Receive (\mathbf{c}, v, ω) from P. Output “true” if and only if $a \stackrel{?}{=} g^v \prod_{i=1}^n g_i^{c_i t_i}$, $b \stackrel{?}{=} h^v \prod_{i=1}^n h_i^{c_i t_i}$, $\mathbf{c} \cdot \mathbf{V}_q^{n \times k} \stackrel{?}{\equiv} \langle c, 0, \dots, 0 \rangle \pmod{q}$, and $\text{ABK-PR}(\mathcal{C}, \mathbf{c}, k, \omega) \stackrel{?}{=} \text{“true”}$; otherwise, output “false”.

As before, some remarks about this protocol are in order. Protocol 2 follows the same basic recipe as Protocol 1, with V starting the conversation in Step V1 by sending to P a list of short exponents for small-exponent batching. In fact, one easily sees by inspection that the repaired version of Protocol 1 is just the special case of Protocol 2 with k fixed to one. (The only difference being that the former protocol uses $q = 2^\lambda - 1$ since it does not require a prime q to guarantee linearly independent constraints.) Completeness holds trivially by inspection and constructing a simulator for honest V is equally straightforward. In the extended version of this paper [23, Appendix B], we prove that using Henry and

Goldberg’s all-but- k mercurial commitment construction [22] in our protocol yields a system for honest-verifier $(0, 1, \min\{\lg p, \lambda\})$ -batch zero-knowledge arguments of knowledge of a size- k subset $S \subseteq [1, n]$ of indices and corresponding set $x_S = \{x_j \in \mathbb{Z}_p \mid j \in S\}$ of exponents such that $\log_g g_j = \log_h h_j = x_j$ for all $j \in S$.

6 Applications

In the introduction, we listed the following example applications in which the need to prove propositions about large batches of predicates naturally arise: cryptographic voting [11, 25], anonymous blacklisting and reputation systems [2, 3], priced symmetric private information retrieval [24], threshold ring signatures [34], verifiable mix networks [21, 31], and cryptographic auctions [9]. We now briefly discuss how our new protocol can directly speed up and extend two such constructions from the literature.

Symmetric private information retrieval. Henry, Olumofin, and Goldberg [24] describe a symmetric variant of Goldberg’s information-theoretic private information retrieval protocol [16] that achieves data privacy by having each client commit to her query using polynomial commitments [26] and then exhibit a zero-knowledge proof that the committed query is “well formed”. The final step in their proof—which dominates the computation cost of their enhancements and contributes considerable communication overhead to the protocol—is a proof of equality of one-out-of- r pairs of discrete logarithms, where r is the number of records in the database. The authors suggest small-exponent batch testing to speed up the verification at the database servers; however, playing the role of prover in that interaction still accounts for a significant fraction of a client’s per-query computational expenditure. Simply swapping in our protocol leads to significant reductions in both the computation overhead and the communication overhead of their protocol.

Cryptographic voting systems. The *JCJ protocol* of Juels, Catalano, and Jakobsson [25] underlies a number of protocols for coercion-resistant, receipt-free verifiable Internet voting [1, 11, 35]; indeed, Spycher et al. [32] opine that “[JCJ is] the only known protocol for remote e-voting that offers individual verifiability and receipt-freeness simultaneously under somewhat acceptable trust assumptions”. The bottleneck operation in JCJ is its *vote authorization* phase, which eliminates fake votes and duplicate votes prior to tallying. The computational cost of both steps grows quadratically in the number of votes cast: JCJ detects fake votes by having voters attach zero-knowledge proofs (made non-interactive via the *Fiat-Shamir heuristic* [13]) that they are on the registered voters roster, and it detects duplicate votes by employing a pairwise *plaintext equivalence test* on each vote. Several papers suggest strategies that can detect duplicate votes in linear time [1, 32, 35]; however, eliminating fake votes in linear time appears to necessitate a weakening the protocol’s security guarantees. In particular, some

existing schemes have voters prove membership within some smaller *anonymity set* rather than the entire roster [11, 32]. Our batch proof of partial knowledge may help to reduce the cost of this step and thereby allow for significantly larger anonymity sets; for smaller elections, it might even make the quadratic algorithm practical.

A second cryptographic operation that frequently arises in end-to-end verifiable voting systems is “proofs of re-encryption” of El Gamal ciphertexts: that is, given two sets of pairs $\{(g^{y_i}, m_i h^{y_i}) \mid i \in [1, n]\}$ and $\{(g^{y'_i}, m'_i h^{y'_i}) \mid i \in [1, n]\}$ of El Gamal ciphertexts encrypted under public key $h = g^x$, prove that $m_i = m'_i$ for all $i \in S$ (where S is a subset of indices suitably defined by the application). Such proofs work by considering the quotients $(m_i h^{y_i}) / (m'_i h^{y'_i}) = (m_i / m'_i) h^{y_i - y'_i}$ and $g^{y_i} / g^{y'_i} = g^{y_i - y'_i}$ and noting that $\log_h (m_i / m'_i) h^{y_i - y'_i} = \log_g g^{y_i - y'_i}$ if and only if $m_i = m'_i$. Thus, batch proofs of knowledge and equality for k -out-of- n pairs of discrete logarithms imply batch proofs of re-encryption of k -out-of- n El Gamal ciphertexts.

7 Conclusion

We have examined “batch zero-knowledge” protocols for communication- and computation-efficient proofs of propositions composed of many simple predicates. Our primary contribution is a novel system for batch zero-knowledge arguments of knowledge and equality of k -out-of- n discrete logarithms for fixed $k \in [1, n]$. We also suggested the first general definitions for *batch zero-knowledge proofs and arguments (of knowledge)*. Our new definitions introduce a *conciseness* property that describes the asymptotic performance of a protocol relative to one formed by sequential composition of single-instance protocols. Our new argument system came about when we analyzed and uncovered a critical flaw in the security proof for Peng and Bao’s [28] batch proofs of knowledge and equality of one-out-of- n discrete logarithms. A malicious prover can exploit the flaw to cause unsuspecting verifiers to accept proofs when the claimed equality of logarithms is false. Fortunately, we showed that the flaw is not fatal: we sketched a fix based on *all-but- k* mercurial commitments with $k = 1$ and then generalized to our main result by varying k in the repaired protocol. In addition, we illustrated the usefulness of our new protocol by sketching some example applications where its adoption could result in noteworthy speedups.

Acknowledgements. We thank Jalaj Uphadyay and Colleen Swanson for helpful discussions and the anonymous reviewers for their comments. The first author is supported by a GO-Bell Graduate Scholarship and by the Natural Sciences and Engineering Research Council of Canada (NSERC) through a Vanier Canada Graduate Scholarship. The second author thanks NSERC for further funding this research through a Discovery Grant and a Discovery Accelerator Supplement.

References

1. Araujo, R., Foulle, S., Traoré, J.: A practical and secure coercion-resistant scheme for remote elections. In: *Frontiers of Electronic Voting*, Schloss Dagstuhl, Germany. Dagstuhl Seminar Proceedings, vol. 7311 (July 2007)
2. Au, M.H., Tsang, P.P., Kapadia, A.: PEREA: Practical TTP-free revocation of repeatedly misbehaving anonymous users. *ACM Transactions on Information and System Security* 14(4), Article No. 29 (2011)
3. Au, M.H., Kapadia, A., Susilo, W.: BLACR: TTP-free blacklistable anonymous credentials with reputation. In: *Proceedings of NDSS 2012*, San Diego, CA, USA (February 2012)
4. Bellare, M., Garay, J.A., Rabin, T.: Batch verification with applications to cryptography and checking. In: Lucchesi, C.L., Moura, A.V. (eds.) *LATIN 1998*. LNCS, vol. 1380, pp. 170–191. Springer, Heidelberg (1998)
5. Bellare, M., Garay, J.A., Rabin, T.: Fast batch verification for modular exponentiation and digital signatures. In: Nyberg, K. (ed.) *EUROCRYPT 1998*. LNCS, vol. 1403, pp. 236–250. Springer, Heidelberg (1998)
6. Bellare, M., Goldreich, O.: On defining proofs of knowledge. In: Brickell, E.F. (ed.) *CRYPTO 1992*. LNCS, vol. 740, pp. 390–420. Springer, Heidelberg (1993)
7. Boneh, D., Boyen, X.: Short signatures without random oracles and the SDH assumption in bilinear groups. *Journal of Cryptology* 21(2), 149–177 (2008)
8. Brands, S., Demuynck, L., De Decker, B.: A practical system for globally revoking the unlinkable pseudonyms of unknown users. In: Pieprzyk, J., Ghodosi, H., Dawson, E. (eds.) *ACISP 2007*. LNCS, vol. 4586, pp. 400–415. Springer, Heidelberg (2007)
9. Cachin, C.: Efficient private bidding and auctions with an oblivious third party. In: *Proceedings of CCS 1999*, Singapore, pp. 120–127 (November 1999)
10. Catalano, D., Fiore, D., Messina, M.: Zero-knowledge sets with short proofs. In: Smart, N.P. (ed.) *EUROCRYPT 2008*. LNCS, vol. 4965, pp. 433–450. Springer, Heidelberg (2008)
11. Clark, J., Hengartner, U.: Selections: Internet voting with over-the-shoulder coercion-resistance. In: Danezis, G. (ed.) *FC 2011*. LNCS, vol. 7035, pp. 47–61. Springer, Heidelberg (2012)
12. Cramer, R., Damgård, I., Schoenmakers, B.: Proof of partial knowledge and simplified design of witness hiding protocols. In: Desmedt, Y.G. (ed.) *CRYPTO 1994*. LNCS, vol. 839, pp. 174–187. Springer, Heidelberg (1994)
13. 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)
14. Fürer, M.: Faster integer multiplication. *SIAM Journal on Computing* 39(3), 979–1005 (2009)
15. Gennaro, R., Leigh, D., Sundaram, R., Yerazunis, W.S.: Batching schnorr identification scheme with applications to privacy-preserving authorization and low-bandwidth communication devices. In: Lee, P.J. (ed.) *ASIACRYPT 2004*. LNCS, vol. 3329, pp. 276–292. Springer, Heidelberg (2004)
16. Goldberg, I.: Improving the robustness of private information retrieval. In: *Proceedings of IEEE S&P 2007*, Oakland, CA, USA, pp. 131–148 (May 2007)
17. Goldreich, O.: A note on computational indistinguishability. *Information Processing Letters* 34(6), 277–281 (1990)

18. Goldreich, O., Micali, S., Wigderson, A.: Proofs that yield nothing but their validity or languages in NP have zero-knowledge proof systems. *Journal of the ACM* 38(3), 691–729 (1991)
19. Goldreich, O., Sahai, A., Vadhan, S.P.: Honest-verifier statistical zero-knowledge equals general statistical zero-knowledge. In: *Proceedings of STOC 1998*, Dallas, TX, USA, pp. 399–408 (1998)
20. Goldwasser, S., Micali, S., Rackoff, C.: The knowledge complexity of interactive proof systems (extended abstract). In: *Proceedings of STOC 1985*, Providence, RI, USA, pp. 291–304 (May 1985)
21. Groth, J.: A verifiable secret shuffle of homomorphic encryptions. *Journal of Cryptology* 23(4), 546–579 (2010)
22. Henry, R., Goldberg, I.: All-but-k mercurial commitments and their applications. Technical Report CACR 2012-26, University of Waterloo, Waterloo, ON, Canada (November 2012)
23. Henry, R., Goldberg, I.: Batch proofs of partial knowledge. Technical Report CACR 2013-08, University of Waterloo, Waterloo, ON, Canada (February 2013)
24. Henry, R., Olumofin, F.G., Goldberg, I.: Practical PIR for electronic commerce. In: *Proceedings of CCS 2011*, Chicago, IL, USA, pp. 677–690 (October 2011)
25. Juels, A., Catalano, D., Jakobsson, M.: Coercion-resistant electronic elections. In: *Proceedings of WPES 2005*, Alexandria, VA, USA, pp. 61–70 (November 2005)
26. 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)
27. Lenstra, A.K., Lenstra, H.W., Lovász, L.: Factoring polynomials with rational coefficients. *Mathematische Annalen* 261(4), 515–534 (1982)
28. Peng, K., Bao, F.: Batch ZK proof and verification of OR logic. In: Yung, M., Liu, P., Lin, D. (eds.) *Inscrypt 2008*. LNCS, vol. 5487, pp. 141–156. Springer, Heidelberg (2009)
29. Peng, K., Bao, F.: Batch range proof for practical small ranges. In: Bernstein, D.J., Lange, T. (eds.) *AFRICACRYPT 2010*. LNCS, vol. 6055, pp. 114–130. Springer, Heidelberg (2010)
30. Peng, K., Boyd, C., Dawson, E.: Batch zero-knowledge proof and verification and its applications. *ACM Transactions on Information and System Security* 10(2), Article No. 6 (2007)
31. Sako, K., Kilian, J.: Receipt-free mix-type voting scheme. In: Guillou, L.C., Quisquater, J.-J. (eds.) *EUROCRYPT 1995*. LNCS, vol. 921, pp. 393–403. Springer, Heidelberg (1995)
32. Spycher, O., Koenig, R.E., Haenni, R., Schläpfer, M.: A new approach towards coercion-resistant remote e-voting in linear time. In: Danezis, G. (ed.) *FC 2011*. LNCS, vol. 7035, pp. 182–189. Springer, Heidelberg (2012)
33. Tsang, P.P., Au, M.H., Kapadia, A., Smith, S.W.: BLAC: Revoking repeatedly misbehaving anonymous users without relying on TTPs. *ACM Transactions on Information and Systems Security* 13(4), Article No. 39 (2010)
34. Tsang, P.P., Wei, V.K., Chan, T.K., Au, M.H., Liu, J.K., Wong, D.S.: Separable linkable threshold ring signatures. In: Canteaut, A., Viswanathan, K. (eds.) *INDOCRYPT 2004*. LNCS, vol. 3348, pp. 384–398. Springer, Heidelberg (2004)
35. Weber, S.G., Araujo, R., Buchmann, J.: On coercion-resistant electronic elections with linear work. In: *Proceedings of ARES 2007*, Vienna, Austria, pp. 908–916 (April 2007)