CDS Composition of Multi-round Protocols

Masayuki Abe[1,5], Andrej Bogdanov[2], Miyako Ohkubo[3], Alon Rosen[4], Zehua Shang[5(✉)], and Mehdi Tibouchi[1,5]

[1] NTT Social Informatics Laboratories, Tokyo, Japan
abe.masayuki.914@gmail.com, mehdi.tibouchi@ntt.com
[2] School of EECS, University of Ottawa, Ottawa, Canada
abogdano@uottawa.ca
[3] Security Fundamentals Laboratory, CSRI, NICT, Tokyo, Japan
m.ohkubo@nict.go.jp
[4] Bocconi University and Reichman University, Milan, Italy
alon.rosen@unibocconi.it
[5] Kyoto University, Kyoto, Japan
shang.zehua.23m@st.kyoto-u.ac.jp

Abstract. We revisit the Cramer, Damgård, Schoenmakers (CDS) approach for composing sigma protocols, and adapt it to a setting in which the underlying protocols have multiple rounds of interaction. The goal of CDS composition is to prove compound NP-relations by combining multiple "atomic" proof systems. Its key feature is that it interacts with the atomic proofs in a generic fashion, enabling simpler and more efficient implementation.

Recent developments in multi-round protocols call for the adaptation of CDS composition beyond its original scope, which not only was restricted to three-move protocols but in fact fails in the multi-round case, as well as in the composition of so-called k-special sound proofs.

We propose a new method for multi-round composition in the plain model, in a soundness preserving way and with an "offline" zero-knowledge simulation property. The need for handling arbitrary monotone access structures in mNC^1, which is all Boolean function families represented by polynomial-size formulas over some fixed complete basis, leads us to identify a complexity theoretic problem of independent interest.

Prior to our work, multi-round composition was either restricted to the random oracle model, or worked only for argument systems, and moreover required heavy "online" zero-knowledge simulation.

Keywords: Zero-Knowledge · Composition · Multi-round · Sigma Protocol · Online Offline

1 Introduction

Driven by efficiency considerations, proof systems are sometimes optimized for specific languages, and then more complex statements are proved by composing

L. Reyzin and D. Stebila (Eds.): CRYPTO 2024, LNCS 14928, pp. 391–423, 2024.
https://doi.org/10.1007/978-3-031-68400-5_12

them. Towards this end, one may be interested in a transformation that takes proof systems for relations, $R_1, \ldots, R_n$ as black-box, and produces a proof system for a compound relation R_Γ for access structure Γ, where R_Γ is satisfied if and only if there exists $(i_1, \ldots, i_t) \in \Gamma$ for which $R_{i_1}, \ldots, R_{i_t}$ are satisfied.

For a class of public-coin proofs known as Σ-protocols [19], Cramer, Damgård, and Schoenmakers presented a framework for generic compositions concerning general monotone access structures represented by monotone formulas in [22]. This framework, often referred to as CDS composition, was later extended to polynomial-size monotone span programs (MSP) in [20]. It is unconditional and preserves the round complexity, soundness, and zero-knowledge of the underlying Σ-protocols.

To attain its desirable features, CDS composition crucially relies on the structure and security of the underlying Σ-protocols. However, the requisite requirements for composition are not always satisfied. For instance, it is observed in [1,25] that if the 2-special soundness of underlying protocols is relaxed to $k(> 2)$-special soundness, e.g., [14,32], CDS composition is no longer secure. Moreover, as noted in [27] and elaborated on later here, its natural extension to multi-round public-coin protocols for more than three moves fails to provide knowledge soundness.

Recent years have witnessed the emergence of several notable multi-round public-coin protocols, including ones based on a pre-processing variant of the MPC-in-the-Head paradigm [34] known as KKW [36], Bulletproofs [15], and many IOP-based protocols, e.g., [12,16,35]. Adapting CDS-like composition to such protocols would be desirable, especially if this composition is generic so that the (already optimized) underlying protocols remain unchanged.

The question of generically turning multi-round protocols into non-interactive arguments has been previously considered in the random oracle model [27], albeit the composition was restricted to a simple 1-out-of-n structure. Another approach followed Lindell's paradigm [38] in which the prover commits to all but the last messages and opens them at the end of the protocol execution [30]. Combined with the idea of "stacking" protocols [28,29], the "commit and open" approach enjoys compressed communication. However, as pointed out in [27], Lindell's paradigm is best suited for arguments rather than proofs, as prover messages tend to be large and so committing to every one of them in a statistically binding manner is communication inefficient. Another shortcoming is that, even in cases as simple as 1-out-of-n composition, the honest prover is forced to defer $n - 1$ simulator runs to the very last message of the protocol.

1.1 Our Contribution

We address the question of whether it is possible to generically compose multi-round protocols in the plain model. We aim to preserve soundness and to allow the prover to perform the bulk of the simulation work before the protocol begins, and in fact even before the composition structure Γ is known. As mentioned earlier, such an offline simulation property is not satisfied by earlier constructions.

Theorem 1. *(Informal) Let Π_i be a $(2\mu + 1)$-round, $(k_1, k_2, \ldots, k_\mu)$-special sound public-coin honest verifier zero-knowledge protocol for* NP*-relation R_i. There exists a composition method that, given Π_i for $i \in [n]$ and a monotone access structure, Γ over $[n]$ as input, outputs a $(2\mu+1)$-round public-coin honest verifier zero-knowledge protocol Π for compound relation R_Γ.*

The resulting Π is as sound as Π_i and special honest verifier computational zero-knowledge. Prover's algorithm is efficient if Γ is in mNC^1*. Prover's online computation complexity is as close as that for proving Π_i's for all $i \in I$ for $I \in \Gamma$ where the prover knows the witness.*

In other words, we extend CDS composition to also work for multi-round and k-special sound protocols without having to sacrifice valuable features such as offline simulation, soundness preservation, and support for a broad class of access structures. Given our result, one can compose efficient protocols such as [6,15] with previously unsupported structures, and the added benefit of offline simulation. The price we pay is that we settle for computational zero-knowledge and linear communication complexity in the number of underlying protocols.

As briefly mentioned in [27], one could extend delayed-input related ideas [18] to multi-round protocols in the same way as [30] extends [29], and then apply techniques from [20] to handle $\Gamma \in$ MSP. The difference between those previous works and ours boils down to whether to commit to 'a' or 'c', where a and c stand for all outgoing messages from the prover and challenges, respectively. The former suits arguments where $\mathsf{Com}(a)$ can be small for large a, but not for proofs or for applications demanding high throughput due to "online" zero-knowledge simulations. We compare the approaches at a high-level in Table 1. A more detailed comparison is given later in Table 2.

Table 1. High-level comparison. Previous works refer to [30] and an extension of [18] based on Lindell's paradigm in the standard model.

	Previous Works	Ours
soundness preservation	Costly for large $\lvert\mathsf{Com}(a)\rvert$	$\surd$ $\lvert\mathsf{Com}(c)\rvert$ is small
performance	all-online $\surd$ compressed argument (t-out-of-n) CRS/one more round	$\surd$ offline[a] ZK simulation linear communication $\surd$ plain & round preserving
functionality	$\surd$ delayed-input $\surd$ MSP	all-fixed input mNC^1

[a]Our offline computation is not completely independent of the statement. It is useful for scenarios where atomic relations are known in advance, and their composition is determined when the proof protocol starts.

1.2 Techniques

We begin by pointing out the challenges arising when attempting to extend CDS composition to multi-round and k-special sound protocols. We assume that readers are familiar with relevant concepts, and refer them to Sect. 2.5 otherwise.

Inconsistent Extraction. Consider as an example a five-round public-coin protocol for proving statement $A \wedge B$ that consists of two sequential executions of three-round protocols proving statements A and B individually. By rewinding at the first and second challenges, respective witnesses ω_A and ω_B are obtained. Then consider composing two such five-round protocols to prove the disjunctive statement $(A \wedge B) \vee (C \wedge D)$. CDS composition suggests running the protocols in parallel, sharing a challenge into two used in each atomic protocol. Suppose that the first challenge, c^1, is shared into c^1_A and c^1_C and used to prove A and C respectively. Similarly, the second challenge, c^2, is shared into c^1_B and c^1_D and used to prove B and D respectively. Suppose that, by rewinding at the first challenge, new challenge $\tilde{c}^1 (\neq c^1)$ is shared into $\tilde{c}^1_A (\neq c^1_A)$ and $\tilde{c}^1_C (= c^1_C)$, yielding ω_A. We thus expect that the prover knows witness (ω_A, ω_B) for clause $(A \wedge B)$. But, by further rewinding at the second challenge, new challenge $\tilde{c}^2 (\neq c^2)$ would be shared into $\tilde{c}^2_B (= c^2_B)$ and $\tilde{c}^2_D (\neq c^2_D)$, yielding ω_D unexpectedly.

Adaptive Choice of Simulated Challenges. For k-special sound protocols, one could simulate executions with up to $k - 1$ different challenges for the same first message. This feature could be exploited by a cheating prover in CDS composition. By adaptively selecting challenges used in simulated executions, the cheating prover can ensure that k rewindings do not yield k colliding transcripts at any protocol executions. In other words, the composed protocol is not special k-sound. In [1], a solution that eliminates prover control by using random oracles to generate challenges was presented. It is not known how to circumvent this problem in the plain model.

Our Approach. Observe that both of the above issues are due to the *adaptive* choice of simulated challenges by the malicious prover. We thus bind all challenges in each protocol run with a commitment at the beginning. They are then opened at the end of the protocol. Using a dual-mode commitment scheme, commitments for real protocol runs are simulated at the beginning and eventually equivocated to actually observed challenges. In prior works, e.g., [18], binding keys are used to commit to the first and all succeeding messages from the prover in the real protocol runs. Our approach uses them in a "dual" way to commit to challenges in each simulated protocol run.

Adversarial Access Structures for Key Allocation. In the above approach, the verifier has to be convinced that the binding and equivocal commitment keys are allocated appropriately. Then, what is the "appropriate" allocation of commitment keys for monotone Γ in general? There must not be too many equivocal commitments for the strategy to be meaningful even if the prover has exceeding

witnesses and can complete real protocol runs than necessary. However, claiming an upper limit is a non-monotone statement, which is challenging to prove efficiently directly. Even if we prove about binding keys, it is unclear whether the structure, say Γ', which the binding keys should follow allows efficient proofs in general. For instance, when Γ is in a compact CNF, can we also have compact, monotone Γ'?

We define this *adversarial structure* Γ' precisely and study its computational complexity in terms of the complexity of Γ. By definition, the monotone DNF size of Γ' is equal to the monotone DNF size of Γ. If, however, Γ is represented as a size-s CNF, then the (unrestricted) circuit complexity of Γ' cannot in general be bounded by any polynomial in s unless $\mathsf{NP} \not\subseteq \mathsf{P/poly}$ (Theorem 3). On the positive side, we show that if Γ is represented as a read-once formula, then the formula size of Γ' is polynomially related to the formula size of Γ (Theorem 4).

To prove Theorem 3, we construct an explicit sequence of monotone 2CNFs Γ_n on with poly(n) variables such that the satisfiability of a 3CNF ϕ reduces to checking whether the adversarial structure Γ'_n accepts an input derived from ϕ. In graph-theoretic language, adversarial structures of monotone 2CNFs are monotone closures of maximal independent sets, allowing us to leverage ideas from Karp's canonical reduction from 3SAT to Independent Set. Theorem 4 is obtained by inductive application of composition rules for conjunction and disjunction of adversarial structures over disjoint inputs.

Challenge Space Extension. Another problem stems from k-special soundness. It is unrelated to the above mentioned adaptive choice of simulated challenges. Although our composition requires secret sharing over the challenge space, the challenge space of the atomic protocol may not be large enough to map the shares. This problem was also observed in the original CDS, but in the case of 2-special sound protocols, it is obvious that the challenge space can be extended t times in bits in t parallel runs. On the other hand, if, for example, two k-special sound protocols are executed in parallel, it is not always the case that either protocol execution will have k distinct challenges in k rewindings. The situation becomes even more difficult when there are multiple challenges in multiple rounds.

We introduce the notion of statistical k-special soundness, which allows for small errors, and show by careful combinatorial analysis that t parallel executions of a k-special sound protocol form a statistical k-special sound protocol (Sect. 5). This allows for a combined challenge space sufficient to map the shares. Similar ideas employing weaker notions of special soundness have been considered in [9,23,24,41,42]. They all follow the generic "special soundness $\Rightarrow$ knowledge soundness" framework and show a weaker notion of special soundness is still knowledge sound. Such weaker notion applies to a more general structure than thresholds. Our starting point, however, diverges from theirs. Since we are more concerned with the issue of soundness preserving during compositions, we instead focus on introducing a notion which is "flexible", rather than consider how much relaxed notions can imply knowledge soundness.

Table 2. Comparison of generic compositions of public-coin protocols. Notations: x-(C)SS: (Computational) x-Special Soundness, RBRS: Round-by-Round Soundness, CS: Computational Soundness. HVZK: Honest Verifier ZK, EHVZK: Extended (Special) Honest Verifier ZK, SHVZK: Special Honest Verifier ZK, HVCZK: Honest Verifier Computational ZK, AOK: Argument of Knowledge, AIS: Adaptive Input Soundness, WI: Witness Indistinguishability, IDTC: Instance Dependent Trapdoor Commitment, Cham-Σ: Chameleon Σ-protocol, PolyCom: Polynomial Commitment, CRH: Collision-Resistant Hash, NIPBC: Non-Interactive Partially Binding Commitment (available in CRS model), DualCom: Dual-mode Commitment.

Scheme	Underlying Protocol			Composed Protocol			Composition	Extra Assumptions	Offline Simulation
	# of Rounds	Soundness	ZK	# of Rounds	Soundness	ZK			
CDS94 [22]	3	2-SS	SHVZK	3	2-SS	SHVZK	MSP	-	yes
Delayed-Input [17]	3	2-SS	SHVZK	3	2-SS	SHVZK	1-OR	CRS, Cham-Σ	no
Delayed Threshold [18]	3	2-SS	SHVZK	3	$(2n+t)$-AIS	WI	t-out-of-n	IDTC	no
Share-then-Hash [1]	3	k-SS	SHVZK	1	AOK	HVZK	MSP	NPROM	-
Acyclicity Program [2]	3	k-SS	SHVZK	1	AOK	SHVZK	ACP	ROM	-
Compressed Σ [6]	3	2-SS	SHVZK	$O(\log(t+n))$	$(n, 2, 3, \ldots, 3)$-CSS	SHVZK	t-out-of-n	PolyCom, DL	yes
DAG-Σ [43]	3	2-SS	SHVZK	3	AOK	SHVZK	CNF	Cham-Σ, CRH	yes
Stacking-Σ [28,29]	3	2-SS	EHVZK	3	AOK	EHVZK	t-out-of-n	NIPBC(CRS)	no
Threshold Stacking [10]	3	2-SS	EHVZK	3	AOK	EHVZK	t-out-of-n	NIPBC(CRS),CRH	no
FGQRW23 [27]	$2\mu+1$	RBRS	HVZK	1	CS	HVZK	1-out-of-n	NPROM, CRS	-
Speed-Stacking [30]	$2\mu+1$	CS	EHVZK	$2\mu+1$	CS	EHVZK	t-out-of-n	NIPBC(CRS)	no
Ours(Sec.3)	$2\mu+1$	$(k_1, \ldots, k_\mu)$-SS	HVZK	$2\mu+1$	$(k_1, \ldots, k_\mu)$-SS	SHVCZK	mNC[1]	DualCom	yes

1.3 Related Work

In Table 2, we present a qualitative comparison of compositions in the literature. We refer to the respective papers for the formal notions of security in the table. The upper table above the line is about the compositions dedicated to three-round protocols, and those below the line support more than three-round protocols. The second groups of columns show the properties of the underlying protocols that are sufficient for the respective compositions to work. Since k-SS is more general than 2-SS, compositions that admit k-SS protocols as input are more general, and hence more widely applicable. Similarly, we have HVZK $>$ SHVZK $>$ EHVZK in terms of generality; compositions accepting HVZK protocols are the most general. (Some compositions admit a trade-off between the generality of the zero-knowledge property and the composition type. In that case, the table takes more general composition types.) The third groups of columns show the properties of the compound protocol. Some compositions noted as 1 round are dedicated to producing non-interactive proofs. In the Soundness column, 2-SS and $(k_1, \ldots, k_\mu)$-SS implies a proof of knowledge (POK), and AOK denotes an argument of knowledge. Regarding the ZK column, SHVZK is a stronger property than HVZK, thus making it preferable as an achieved goal. The rightmost groups of columns show the type of composition, extra assumptions, and capability of offline zero-knowledge simulation. MSP and ACP are orthogonal and the most general types of compositions. Regarding the extra assumptions, we focus on assumptions that highlight important differences and ignore mild ones such as assuming some common domains for all underlying protocols.

Many of previous works, e.g., [1,2,6,11,17,22,29,32,33,43], are dedicated to three-round protocols. Each has its own unique properties in terms that may not be listed in Table 1, in exchange of limitations or extra assumptions. The composition in [17] is for a single OR only but allows delayed inputs where the instances to prove can be determined after the protocol starts. [18] gives compositions of delayed-input protocols for threshold structures and states adaptive soundness. As explained earlier, it can be extended to multi-round composition following the structure of [30]. [29] achieves logarithmic communication complexity in the number of variables for a 1-out-of-n structure in the CRS model. Its threshold variant is improved in communication in [10]. [6] also enjoys logarithmic communication complexity for relations commonly used in the discrete-logarithm setting. [2] follows the idea of sequential-OR composition in [3,26] and extends the composition function to acyclicity programming producing non-interactive proofs in the random oracle model. Multi-round compositions are addressed explicitly in [27] and [30] as we mentioned earlier.

In [37], an OR-composition of Ligero [5] proof system is presented. The prover commits to the challenges of Ligero with a statistically hiding commitment and later proves that the commitment is consistent with one of the transcripts. This idea can be seen as the bare bones of approaches using commitments, leaving the generality and the key generation issue unsettled. In particular, their composition increases the number of rounds for their protocol.

1.4 Organization

In Sect. 2, we introduce basic notations and building blocks. We also revisit standard definitions regarding public-coin proof protocols and CDS composition with its natural extension to the multi-round case, which is our starting point. We present our composition in Sect. 3 followed by a complexity analysis of adversarial structures in Sect. 4 and challenge space extension of k-special sound protocols in Sect. 5. We finally conclude with some closing thoughts in Sect. 6.

2 Preliminaries

2.1 Notation

For a finite set S, we write $a \leftarrow\!\$\, S$ to denote that a is uniformly sampled from S. We denote the security parameter by $\lambda \in \mathbb{N}$. Given two functions $f, g : \mathbb{N} \to [0,1]$, we write $f \approx g$ if the difference $|f(\lambda) - g(\lambda)|$ is asymptotically smaller than the inverse of any polynomial. A function f is said to be negligible if $f \approx 0$, whereas it is said to be overwhelming when $f \approx 1$. For integers m, n, such that $m \leq n$, we denote by $[m,n]$ the range $\{m, m+1, \ldots, n\}$. We denote by $[n]$ the range $[1,n]$. A sequence of indexed values, $(x_1, \ldots, x_n)$ is denoted by $\{x_i\}_{i\in[n]}$.

Similarly, we denote $(x_{i_1}, \ldots, x_{i_k})$ by $x_{i\in I}$ for $I = \{i_1, \ldots, i_k\}$. We assume that index set I is trivial from $x_{i\in I}$. When $\mathcal{A}$ is a probabilistic algorithm, we denote by $\mathcal{A}(x;r)$ an execution of $\mathcal{A}$ on input x and random coins r taken from

an appropriate domain defined for $\mathcal{A}$. If the random coins are not important, we simply write $\mathcal{A}(x)$. We generally assume stateful adversaries, e.g., in the game execution $a \leftarrow \mathcal{A}(x); e \leftarrow\!\!\$\, \{0,1\}^\lambda; z \leftarrow \mathcal{A}(e)$ the adversary $\mathcal{A}$ in the second call knows the state of $\mathcal{A}$ after the first call (in particular, it knows x and a).

2.2 Public-Coin Proof System

We follow the definitions of [7]. Let $R : \mathcal{X} \times \mathcal{W} \to \{0,1\}$ be a binary relation defined over a set of instances $\mathcal{X}$ and a set of witnesses $\mathcal{W}$. It is also denoted as $R(instance, witness) = \{description\}$ that evaluates to 1 if and only if *description* is fulfilled for *instance* with *witness*. Language L_R associated by R is $L_R := \{x \in \mathcal{X} \mid \exists w \in \mathcal{W} : R(x,w) = 1\}$. By L_{RW}, we denote set $L_{RW} := \{(x,w) \mid R(x,w) = 1\}$. An interactive proof system Π for relation R is a pair of interactive algorithms, prover P and verifier V, where a witness w is given to P as private input and instance x is given to both P and V as common input, and V outputs $b \in \{0,1\}$ at the end of the execution. By $\langle P(w), V\rangle(x) \to b$ we denote an execution of the algorithms. A transcript of a protocol execution consists of x, b, and all content of input and output communication tapes of V. It is (perfectly) complete if, for any $(x,w) \in L_{RW}$, $\langle P(w), V\rangle(x) \to 1$.

Definition 1 (Knowledge Soundness). *An interactive proof system for relation R is knowledge sound with knowledge error ϵ if there exists an expected polynomial-time algorithm $\mathcal{E}$ (extractor) that, for every algorithm P^*, every $x \in \{0,1\}^\lambda$, and $aux \in \{0,1\}^*$, there exists a negligible function ϵ that* $\Pr[w \leftarrow E^{P^*(aux)}(x) : R(x,w) = 1] \geq \Pr[\langle P^*(aux), V\rangle(x) = 1] - \epsilon(\lambda)$.

An interactive proof system is called a proof of knowledge (PoK) if it is knowledge sound against unbound P^* as above. It is called an argument of knowledge (AoK) if it is knowledge sound against polynomial-time P^*.

Definition 2 (Honest-verifier Zero-Knowledge). *An interactive proof system is honest-verifier zero-knowledge if there exists a polynomial-time algorithm $\mathcal{S}$ (simulator) that, for any $(x,w) \in L_{RW}$, distribution of outputs from $\mathcal{S}(x)$ and that of transcripts observed in $\langle P(w), V\rangle(x)$ are indistinguishable.*

Definition 3 (Public-Coin Proof Protocol). *A $(2\mu+1)$-round public-coin proof protocol for relation R is a set of polynomial-time algorithms $\mathcal{A}$, $\{\mathcal{Z}_i\}_{i\in[\mu]}$, $\mathcal{V}$ and efficiently and uniformly sampleable space $\{\mathcal{C}_i\}_{i\in[\mu]}$ that constitutes an interactive proof system (P,V) that:*

Step 1: P runs $\mathcal{A}(x,w;r) \to a$ and sends a to V.
Step $2i$: V uniformly choose c_i from $\mathcal{C}_i$ and send it to P.
Step $2i+1$: P runs $\mathcal{Z}_i(x,w,\{c_j\}_{j\in[i]};r) \to z_i$ and sends z_i to V
Output: V runs $\mathcal{V}(x,a,\{c_i\}_{i\in[\mu]},\{z_i\}_{i\in[\mu]}) \to b$ and outputs b.

A tuple $(x,a,\{c_i\}_{i\in[\mu]},\{z_i\}_{i\in[\mu]})$ is called an accepting transcript if $\mathcal{V}(x,a,\{c_i\}_{i\in[\mu]},\{z_i\}_{i\in[\mu]}) = 1$.

Definition 4 (Tree of Transcripts [7]). *Let $(k_1, \ldots, k_\mu)$-tree of transcripts for a $(2\mu+1)$-round public-coin proof protocol is a set of $K = \prod_{i=1}^{\mu} k_i$ transcripts arranged in the following tree structure. Nodes in this tree correspond to responses from the prover and the edges to challenges from the verifier. Every node at depth i has k_i children and k_i pairwise distinct challenges. Every transcript can be represented as one path from the root node to a leaf node.*

Definition 5 ($(k_1, \ldots, k_\mu)$-Special Soundness [7]). *A $(2\mu+1)$-round public-coin proof protocol is $(k_1, \ldots, k_\mu)$-special sound if there exists a polynomial-time algorithm that, given a distinct $(k_1, \ldots, k_\mu)$-tree of accepting transcripts, outputs w that satisfies $R(x, w) = 1$.*

In [7], it is shown that the above implies knowledge soundness.

Definition 6 (Special Honest-Verifier Zero-Knowledge). *A $(2\mu+1)$-round public-coin proof protocol is special honest-verifier zero-knowledge if there exists a polynomial-time algorithm $\mathcal{S}$ that, for any $(x, w) \in L_{RW}$ and $c_i \in \mathcal{C}_i$ for $i \in [\mu]$, distribution of $(a, \{c_i\}_{i\in[\mu]}, \{z_i\}_{i\in[\mu]})$ generated as $(a, \{z_i\}_{i\in[\mu]}) \leftarrow \mathcal{S}(x, \{c_i\}_{i\in[\mu]})$ and that of $(a', \{c_i\}_{i\in[\mu]}, \{z'_i\}_{i\in[\mu]})$ generated as $a' \leftarrow \mathcal{A}(x, w; r)$ and $z'_i \leftarrow \mathcal{Z}_i(x, w, \{c_j\}_{j\in[i]}; r)$ are indistinguishable.*

A Σ-protocol [19] is a three-round public coin proof protocol that is 2-special sound and special honest verifier zero-knowledge.

2.3 Monotone Access Structure

First we recall the definition of the monotone access structure from [22].

Definition 7 (Monotone Access Structure [22]). *An access structure $\Gamma \subset 2^M$ defined over a set M is called a monotone access structure if for all $A \in \Gamma$ and for all $B \supset A$ it holds that $B \in \Gamma$. Sets in Γ are called authorized sets, and sets not in Γ are called unauthorized sets.*

Definition 8 (Dual Structure [22]). *Let Γ be an access structure defined over a set M. If $A \subseteq M$, then $\bar{A}$ denotes the complement of A in M. Now Γ^*, the dual access structure is defined as follows:*

$$A \in \Gamma^* \Leftrightarrow \bar{A} \notin \Gamma.$$

The dual Γ^* of a monotone access structure is also monotone, and satisfies $(\Gamma^*)^* = \Gamma$.

2.4 Secret Sharing Scheme

A semi-smooth perfect secret sharing scheme [22], SSS_Γ, over domain S and access structure Γ over a set M consists of four polynomial-time algorithms, Share, Rec, CheckShares and Complete that:

- $\mathsf{Share}_{\Gamma}(s) \to \{s_i\}_{i\in M}$ is a probabilistic algorithm that takes secret $s \in S$ and outputs shares $\{s_i\}_{i\in M}$.
- $\mathsf{Rec}_{\Gamma}(\{s_i\}_{i\in A}) \to s$ is a reconstruction algorithm that takes a qualified set of shares $\{s_i\}_{i\in A}$, $A \in \Gamma$, and recovers secret $s \in S$.
- $\mathsf{CheckShares}_{\Gamma}(s, \{s_i\}_{i\in M})$ is a share verification algorithm that takes secret s and all shares $\{s_i\}_{i\in M}$, returns 1 or 0.
- $\mathsf{Complete}_{\Gamma}(s, \{s_i\}_{i\in \bar{A}})$ takes shares of a non-qualified set of shares $\{s_i\}_{i\in \bar{A}}$ for $A \in \Gamma$ and secret s, and outputs $\{s_i\}_{i\in A}$ that $\{s_i\}_{i\in M}$ constitute a complete set of shares of s.

It provides the following properties:

- Correctness: For all $s \in S$, $\{s_i\}_{i\in M} \leftarrow \mathsf{Share}_{\Gamma}(s)$, and $A \in \Gamma$, it holds that $s \leftarrow \mathsf{Rec}(\{s_i\}_{i\in A})$.
- Perfect hiding: For any $A \in \Gamma$, $s \in S$, and $\{s_i\}_{i\in M} \leftarrow \mathsf{Share}_{\Gamma}(s)$, distribution of $\{s_i\}_{i\in \bar{A}}$, denoted by $S_{\bar{A}}$, is independent of s.
- Consistency testing: $\mathsf{CheckShares}_{\Gamma}(s, \{s_i\}_{i\in M})$ returns 1 if and only if, for all $A \in \Gamma$, $\mathsf{Rec}_{\Gamma}(\{s_i\}_{i\in A}) = s$.
- Share completion: For any $A \in \Gamma$, $s \in S$, $\{s_i\}_{i\in \bar{A}} \leftarrow S_{\bar{A}}$, and $\{s_i\}_{i\in A} \leftarrow \mathsf{Complete}_{\Gamma}(s, \{s_i\}_{i\in \bar{A}})$, it holds that $1 \leftarrow \mathsf{CheckShares}_{\Gamma}(s, \{s_i\}_{i\in M})$.

Efficient SSS_{Γ} exists for Γ being a threshold structure [39], monotone circuit [13], and monotone span program [21]. If, for every $A \in \Gamma$, $S_{\bar{A}}$ equals uniform distribution over $S^{|\bar{A}|}$, then it is called a smooth perfect secret sharing scheme. Shamir's secret sharing scheme for threshold structures is an example.

2.5 CDS Composition

We describe CDS composition in a general form for $2\mu + 1$-round protocols. It matches the original one at $\mu = 1$. We warn that it is for introducing consistent notations for succeeding sections, and indeed not sound for $\mu \geq 2$ and $k(> 3)$-special sound as mentioned earlier.

Let Π_{R_i} be a $(2\mu + 1)$-round public-coin proof protocol for relation R_i, and (x_i, w_i) be a pair of instance and witness satisfying $R_i(x_i, w_i) = 1$. Let Γ be a monotone access structure over $[n]$, and $R_{\Gamma}((x_i, \ldots, x_n), (w_i, \ldots, w_n))$ be a compound relation that returns 1 if and only if there exists $A \in \Gamma$ that $R_i(x_i, w_i) = 1$ for all $i \in A$. We denote the honest verifier zero-knowledge simulator of Π_{R_i} by $\mathcal{S}_i$. Given Π_{R_i} for $i \in [n]$ and access structure Γ, CDS composition constructs Prover and Verifier as follows where steps $2l$ and $2l + 1$ are repeated for $l \in [\mu]$:

$\Pi^{\mathsf{CDS}}_{\mu,\Gamma}(\mathsf{Prover}(\{x_i\}_{i\in[n]}, \{w_i\}_{i\in A}), \mathsf{Verifier}(\{x_i\}_{i\in[n]}))$:

1. For each $i \in \bar{A}$, Prover calls the honest verifier zero-knowledge simulator $(a_i, c_i^{l\in[\mu]}, z_i^{l\in[\mu]}) \leftarrow \mathcal{S}_i(x_i)$. For those $i \in A$, Prover commits to a_i by running $a_i \leftarrow \mathcal{A}_i(x_i; r_i)$. Prover sends $\{a_i\}_{i\in[n]}$ to Verifier.

$2l$. Verifier samples $c^l \xleftarrow{\$} \mathcal{C}^l$, and sends it to Prover.

$2l+1$. Prover completes shares by $\{c_i^l\}_{i\in A} \leftarrow \mathsf{Complete}_{\Gamma^*}(c^l, \{c_i^l\}_{i\in\bar{A}})$, and computes $z_i^l \leftarrow \mathcal{Z}_i(x_i, w_i, c_i^{j\in[l]}; r_i)$ for all $i \in A$. It then sends $\{\{c_i^l\}_{i\in[n]}, \{z_i^l\}_{i\in[n]}\}$ to Verifier.

Final. Verifier runs $\mathcal{V}_i(x_i, a_i, c_i^{l\in[\mu]}, z_i^{l\in[\mu]})$ for $i \in [n]$ and $\mathsf{CheckShares}_{\Gamma^*}(c^l, \{c_i^l\}_{i\in[n]})$ for $l \in [\mu]$, and outputs 1 if all outputs are 1, outputs 0, otherwise.

It is shown in [22] that, if every Π_{R_i} is a 3-round public-coin proof protocol that is 2-special sound and honest verifier zero-knowledge, and Γ admits a smooth perfect secret sharing scheme, then the above protocol $\Pi_{1,\Gamma}^{\mathsf{CDS}}$ is a Σ-protocol for relation R_Γ. It admits offline simulation since zero-knowledge simulator $\mathcal{S}_i$ is invoked only in the first step of the prover algorithm. If every Π_{R_i} is special honest verifier zero-knowledge, the above can be augmented to accept access structure Γ that admits a semi-smooth secret sharing scheme. It is done by modifying the prover's first step in a way that it first shares a random secret to obtain challenge c_i for $i \in \bar{A}$ and runs the special honest verifier zero-knowledge simulator on input x_i and c_i.

2.6 Dual-Mode Commitment

Definition 9 (Dual-Mode Commitment Scheme). *A dual-mode commitment scheme,* $\mathsf{dmCom} := (\mathsf{CGen}_{\mathsf{BIND}}, \mathsf{CGen}_{\mathsf{HIDE}}, \mathsf{Com}, \mathsf{Eqv})$*, is a tuple of polynomial-time algorithms that:*

- $\mathsf{CGen}_{\mathsf{BIND}}(1^\lambda) \to \mathsf{ck}$*: The binding key generation algorithm that generates a binding commitment key* ck*.*
- $\mathsf{CGen}_{\mathsf{HIDE}}(1^\lambda) \to (\mathsf{ck}, \mathsf{td})$*: The hiding key generation algorithm that outputs a hiding commitment key* ck *and a trapdoor* td*.*
- $\mathsf{Com}_{\mathsf{ck}}(m; r) \to \mathsf{com}$*: The commitment algorithm that takes a message* m *and commitment key* ck *as input, and outputs a commitment* com*.*
- $\mathsf{Eqv}_{\mathsf{td}}(\mathsf{com}, m', r', m) \to r$*: The equivocation algorithm that takes trapdoor* td*, and target* $(\mathsf{com}, m', r', m)$ *of equivocation as input, outputs a randomness* r *matching* m*.*

We require the dual-mode commitment scheme to provide the following properties. It is noted that we do not require binding property in the hiding mode.

Definition 10 (Security of Dual-mode Commitment).

- *Mode Indistinguishability: Distribution of binding keys and that of hiding keys are computationally indistinguishable.*
- *Equivocality: For any* $\lambda \in \mathbb{N}$*,* $(\mathsf{ck}, \mathsf{td}) \leftarrow \mathsf{CGen}_{\mathsf{HIDE}}(1^\lambda)$*,* $(m', m) \in \{0,1\}^*$*,* $\mathsf{com} \leftarrow \mathsf{Com}_{\mathsf{ck}}(m'; r')$*,* $r \leftarrow \mathsf{Eqv}_{\mathsf{td}}(\mathsf{com}, \mathsf{ck}, m', r', m)$*, it holds that* $\mathsf{Com}_{\mathsf{ck}}(m; r) = \mathsf{com}$*.*
- *Perfect Binding in Binding Mode: For any* $\lambda \in \mathbb{N}$*,* $\mathsf{ck} \leftarrow \mathsf{CGen}_{\mathsf{BIND}}(1^\lambda)$*, and* $\mathsf{com} \in \{0,1\}^*$*, there exists at most one* $(m, r) \in \{0,1\}^*$ *that satisfies* $\mathsf{com} = \mathsf{Com}_{\mathsf{ck}}(m; r)$*.*

- *Computational Hiding in Binding Mode: For any* $\lambda \in \mathbb{N}$, $\mathsf{ck} \leftarrow \mathsf{CGen}_{\mathsf{BIND}}(1^\lambda)$, *and* $(m_0, m_1) \in \{0,1\}^{\mathsf{poly}(\lambda)}$ *of equal length, distributions* $D_{\mathsf{ck}}(m_b) := \{(\mathsf{ck}, \mathsf{com}) \mid \mathsf{com} \leftarrow \mathsf{Com}_{\mathsf{ck}}(m_b; r)\}$ *for* $b \in \{0,1\}$ *are computationally indistinguishable.*
- *Perfect Hiding in Hiding Mode: For any* $\lambda \in \mathbb{N}$, $(\mathsf{ck}, \mathsf{td}) \leftarrow \mathsf{CGen}_{\mathsf{HIDE}}(1^\lambda)$, *and* $(m_0, m_1) \in \{0,1\}^{\mathsf{poly}(\lambda)}$ *of equal length, above distributions* $D_{\mathsf{ck}}(m_0)$ *and* $D_{\mathsf{ck}}(m_1)$ *are identical.*

As an additional requirement, we assume that there exists a Σ-protocol, Π^{bind}, for proving the binding-key relation:

$$R^{\mathsf{bind}}(\mathsf{ck}, t) = \{\mathsf{ck} = \mathsf{CGen}_{\mathsf{BIND}}(1^\lambda; t)\}.$$

In Appendix A, we present an instantiation of a dual-mode commitment scheme with an efficient Σ-protocol for the binding relation based on the decision Diffie-Hellman assumption.

3 Our Composition

We first introduce building blocks for our construction in Sect. 3.1 including some new notions referred in the construction and further studied in succeeding sections. The main construction is presented in Sect. 3.2 followed by analysis of security and performance in Sect. 3.3 and 3.4. We also discuss some extensions in Sect. 3.5.

3.1 Building Blocks

We define some structures associated to Γ and present its useful properties.

Definition 11 (Adversarial Structure). *For a monotone access structure* Γ *over a set* M*, minimal authorized structure* Γ_{min} *, minimal adversarial structure* $\bar{\Gamma}_{min}$ *, and monotone adversarial structure* Γ' *are defined as follows:*

$$\begin{aligned} \Gamma_{min} &= \{A \mid A \in \Gamma \wedge \forall a \in A, A \setminus \{a\} \notin \Gamma\}, \\ \bar{\Gamma}_{min} &= \{A \mid \bar{A} \in \Gamma_{min}\}, \\ \Gamma' &= \{B \mid \exists A \in \bar{\Gamma}_{min} \ \ s.t.\ B \supseteq A\}. \end{aligned}$$

As an example, if $M = [n]$ and Γ is a (t, n)-threshold structure, then Γ^* is a $(n - t + 1, n)$-threshold structure, $\Gamma_{\min} = \binom{n}{t}$, $\bar{\Gamma}_{\min} = \binom{n}{n-t}$, and Γ' is a $(n - t, n)$-threshold structure.

Our main construction involves the following building blocks:

- $\Pi_i = (\mathcal{A}_i, \mathcal{C}_i^{l \in [\mu]}, \mathcal{Z}_i^{l \in [\mu]}, \mathcal{V}_i, \mathcal{S}_i)$: a $(2\mu + 1)$-round public-coin proof protocol for relation R_i.
- $\mathsf{dmCom} := (\mathsf{CGen}_{\mathsf{BIND}}, \mathsf{CGen}_{\mathsf{HIDE}}, \mathsf{Com}, \mathsf{Eqv})$: a dual-mode commitment scheme equipped with Σ-protocol Π^{bind} for relation R^{bind}.

– $\mathsf{SSS}_{\Gamma^*} = (\mathsf{Share}, \mathsf{Rec}, \mathsf{CheckShares}, \mathsf{Complete})$: a smooth secret sharing scheme for Γ.

From Π^{bind}, we construct a Σ-protocol for proving the allocation of binding keys satisfying adversarial structure Γ'. Let $R^{\mathsf{bind}}_{\Gamma'}$ be a relation composed of R^{bind} for n keys $\mathsf{ck}_1, \ldots, \mathsf{ck}_n$ according to access structure Γ'. That is, $R^{\mathsf{bind}}_{\Gamma'}(\mathsf{ck}_i, t_i)$ is satisfied if and only if there exists $A \in \Gamma'$ that all $\{\mathsf{ck}_i\}_{i\in A}$ are binding keys. Let $\Pi^{\mathsf{bind}}_{\Gamma'} = (\mathcal{A}_{\Pi^{\mathsf{bind}}_{\Gamma'}}, \mathcal{C}_{\Pi^{\mathsf{bind}}_{\Gamma'}}, \mathcal{Z}_{\Pi^{\mathsf{bind}}_{\Gamma'}}, \mathcal{V}_{\Pi^{\mathsf{bind}}_{\Gamma'}})$ be a Σ-protocol for relation

$$R^{\mathsf{bind}}_{\Gamma'} = \{(\{\mathsf{ck}_i\}_{i\in[n]}; \{t_i\}_{i\in\bar{A}}) : \mathsf{ck}_i = \mathsf{CGen}_{\mathsf{BIND}}(1^\lambda; t_i)\ \forall i \in \bar{A}\}.$$

It can be constructed by composing Π^{bind} with respect to Γ' according to CDS composition.

Finally, we introduce statistical k-special soundness that slightly relaxes k-special soundness by allowing small statistical errors with adaptation to the multi-round setting. This careful treatment is important for our composition to cover a broader range of realistic atomic protocols. Section 5 presents a detailed study of this notion.

Definition 12 (Statistical $(k_1, \ldots, k_\mu)$-Special Soundness). *A $(2\mu + 1)$-round public-coin proof protocol is statistical $(k_1, \ldots, k_\mu)$-special sound if there exists a polynomial-time algorithm that, given a distinct $(k_1, \ldots, k_\mu)$-tree of accepting transcripts chosen from all possible trees of accepting transcripts, outputs w that satisfies $R(x, w) = 1$ except for some negligible probability κ. We denote this κ as the special soundness error.*

Remark 1. In this definition we consider the case that the special soundness error κ distributes over the choices of challenges. To make this precise, we simplify trees of accepting transcripts as follows: we ignore differences of nodes among different trees, i.e., we regard all trees with the same edges but different nodes as the same tree. We denote a tree is "bad" if the extractor fails to extract a witness from one of the original trees associated to it. The special soundness error $\kappa = \frac{|\text{bad trees}|}{|\text{all accepting trees}|}$. There exists an extractor that taken as input a tree of accepting transcripts succeeds in extracting a witness with overwhelming probability $1 - \kappa$ over the choices of challenges.

3.2 Main Construction

Our soundness-preserving composition is presented in Fig. 1. The private input to the prover algorithm is a minimal set of witnesses $w_{i\in A}$ that $A \in \Gamma_{\min}$. For monotone Γ, a prover holding non-minimal authorized set $A' \in \Gamma$ can find a minimal authorized subset $A \subseteq A'$ efficiently just by eliminating elements one by one until it reaches A for which no more element can be removed.

Protocol $\Pi_{\Gamma}^{\mathsf{sound}}$

Statement: $\{x_i\}_{i\in[n]}$

Witness: $\{w_i\}_{i\in A\in\Gamma_{\min}}$

Prover in Round 1 :

- **for all** $i \in \bar{A}$:
  - Generate $\mathsf{ck}_i \leftarrow \mathsf{CGen}_{\mathsf{BIND}}(1^\lambda; t_i)$.
  - Simulate $\pi_i = (a_i^1, c_i^1, z_i^1, \ldots, c_i^\mu, z_i^\mu) \leftarrow \mathcal{S}_i(x_i)$.
  - Compute $\mathsf{com}_i = \mathsf{Com}_{\mathsf{ck}_i}(\{c_i^l\}^{l\in[\mu]}; r_{\mathsf{com}_i})$.
- **for all** $j \in A$:
  - Generate $(\mathsf{ck}_j, \mathsf{td}_j) \leftarrow \mathsf{CGen}_{\mathsf{HIDE}}(1^\lambda; t_j)$.
  - Compute $a_j \leftarrow \mathcal{A}_j(x_j, w_j; r_j)$.
  - Compute $\mathsf{com}_j = \mathsf{Com}_{\mathsf{ck}_j}(\{0\}^{l\in[\mu]}; r'_{\mathsf{com}_j})$.
- Compute $a_{\mathsf{bind}} \leftarrow \mathcal{A}_{\Pi^{\mathsf{bind}}_{\Gamma'}}((\{\mathsf{ck}_i\}_{i\in[n]}, \{t_i\}_{i\in\bar{A}}); r_{\mathsf{bind}})$.
- Send $(\{a_i, \mathsf{com}_i, \mathsf{ck}_i\}_{i\in[n]}, a_{\mathsf{bind}})$ to the verifier.

For $l \in [\mu]$:

Verifier in Round $2l$:

- Sample $c^l \leftarrow\!\!\$\ \mathcal{C}^l$ and sends c^l to the prover.
- **if** $l = 1$, sample $c_{\mathsf{bind}} \leftarrow\!\!\$\ \mathcal{C}_{\Pi^{\mathsf{bind}}_{\Gamma'}}$ and sends c_{bind} to the prover.

Prover in Round $2l+1$:

- Compute $\{c_i^l\}_{i\in[n]} \leftarrow \mathsf{Complete}_{\Gamma^*}(c^l, \{c_i^l\}_{i\in\bar{A}})$.
- **for all** $j \in A$:
  - Compute $z_j^l \leftarrow \mathcal{Z}_j^l(x_j, w_j, a_j^1, \{c_j^m, z_j^m\}^{m\in[l-1]}, c_j^l; r_j)$.
  - **if** $l = \mu$, compute $r_{\mathsf{com}_j} = \mathsf{Eqv}_{\mathsf{td}_j}(\mathsf{ck}_j, \{0\}^{l\in[\mu]}, \{c_j^l\}^{l\in[\mu]}, r'_{\mathsf{com}_j})$.
- **if** $l = 1$, compute $z_{\mathsf{bind}} \leftarrow \mathcal{Z}_{\Pi^{\mathsf{bind}}_{\Gamma'}}(\{\mathsf{ck}\}_{i\in[n]}, a_{\mathsf{bind}}, c_{\mathsf{bind}}, r_{\mathsf{bind}})$.
- **if** $l = 1$, send $\{c_i^1, z_i^1, z_{\mathsf{bind}}\}$ to the verifier.
- **if** $l = \mu$, send $\{c_i^\mu, z_i^\mu, r_{\mathsf{com}_i}\}_{i\in[n]}$ to the verifier.
- **else** Send $\{c_i^l, z_i^l\}_{i\in[n]}$ to the verifier.

Verification: Accept iff

- $\mathcal{V}_{\Pi^{\mathsf{bind}}_{\Gamma'}}(\{\mathsf{ck}_i\}_{i\in[n]}, a_{\mathsf{bind}}, c_{\mathsf{bind}}, z_{\mathsf{bind}}) = 1$.
- **for all** $i \in [n]$:
  - $\mathcal{V}_i(x_i, \pi_i = (a_i^1, c_i^1, z_i^1, \ldots, c_i^\mu, z_i^\mu)) = 1$.
  - $\mathsf{com}_i = \mathsf{Com}_{\mathsf{ck}_i}(\{c_i^l\}^{l\in[\mu]}; r_{\mathsf{com}_i})$.
- **for all** $l \in [\mu]$:
  - $\mathsf{CheckShares}_{\Gamma^*}(c^l, \{c_i^l\}_{i\in[n]}) = 1$.

Fig. 1. Our Multi-Round CDS Composition.

3.3 Security

We claim the following security properties of $\Pi^{\mathsf{sound}}_{\Gamma}$.

Theorem 2 ($\Pi^{\mathsf{sound}}_{\Gamma}$). *If* dmCom *is secure as defined in Definition 10,* SSS_{Γ^*} *is a smooth secret sharing scheme for access structure* Γ^**,* Π^{bind} *is complete, 2-special sound, and special honest verifier zero-knowledge, and every* Π_i *for relation* R_i *is complete, perfect (resp. statistical)* $(k_1, \ldots, k_\mu)$*-special sound, and honest verifier zero-knowledge, then, protocol* $\Pi^{\mathsf{sound}}_{\Gamma}$ *is a* $(2\mu+1)$*-round public-coin proof protocol for relation* $R_\Gamma((x_i, \ldots, x_n), (w_i, \ldots, w_n)) := \{\exists A \in \Gamma, \forall i \in A,\ R_i(x_i, w_i) = 1\}$*. It is perfectly complete, perfect (resp. statistical)* $(k_1, \ldots, k_\mu)$*-special sound, and special honest verifier computational zero-knowledge.*

Proof. **Completeness.** Completeness directly follows from the correctness of dmCom and SSS_{Γ^*}, the completeness and honest verifier zero-knowledge of Π_i, and completeness of $\Pi^{\mathsf{bind}}_{\Gamma'}$. We note that $\mathsf{Complete}_{\Gamma^*}$ works for set $\bar{A}$ of shares $\{c_i\}_{i\in\bar{A}}$ since $\bar{A} \notin \Gamma^*$ when $A \in \Gamma_{\min}$.

Special Soundness. We first prove for the case in which the underlying protocols are perfect $(k_1, \ldots, k_\mu)$-special sound. We prove that given any accepting $(k_1, \ldots, k_\mu)$-tree of transcripts as input, there exists an efficient extractor $\mathcal{E}$ that outputs a valid set of witnesses $\mathbf{w} = \{w_i\}_{i\in A}$ where $A \in \Gamma_{\min}$. The following facts are implied by the accepting tree of transcripts:

Fact 1. For all $i \in [n]$, the verifier $\mathcal{V}_i$ accepts on any distinct statement and atomic transcript (x_i, Π_i);
Fact 2. For all $i \in [n]$, the commitment com_i is generated with corresponding commitment key ck_i;
Fact 3. The shares $\{c_i^{l,k}\}_{i\in[n],k\in[k_l]}$ of each round $l \in [\mu]$ are consistent with the corresponding challenges $\{c_k^{l,k}\}_{k\in[k_l]}$.

Before we show the extractor, we first introduce the following lemma which is useful in our proof.

Lemma 1. *For any* $A \in \Gamma_{min}$ *and* $a \in A$*, it holds that* $\bar{A} \cup \{a\} \in \Gamma^*$*.*

Proof. Let $B = \bar{A} \cup \{a\}$. It suffices to show $\bar{B} \notin \Gamma$. Observe that $\bar{B} = A \setminus \{a\}$. Since $A \in \Gamma_{\min}$, it holds that $\bar{B} = A \setminus \{a\} \notin \Gamma_{\min} \subseteq \Gamma$. □

The extractor $\mathcal{E}$ works as follows: It first runs the extractor of $\Pi^{\mathsf{bind}}_{\Gamma'}$ on statements $\{\mathsf{ck}_i\}_{i\in[n]}$ and transcripts taken from the $(k_1, \ldots, k_\mu)$-tree of transcripts. This extraction will succeed since the commitment keys $\{\mathsf{ck}_i\}_{i\in[n]}$ and the first message of $\Pi^{\mathsf{bind}}_{\Gamma'}$ are all fixed in the initial message (i.e. the root) in the tree. Moreover, since any branches in the tree are accepting and $\Pi^{\mathsf{bind}}_{\Gamma'}$ is perfect 2-special sound, any 2 branches at $l = 1$ are sufficient to extract witnesses for $\Pi^{\mathsf{bind}}_{\Gamma'}$. The extracted witnesses identify a set $A' \in \Gamma'$ where the prover commits to challenges in the binding mode.

Suppose that $A' \supset \bar{A}$ contains more elements than $\bar{A}$. The intuition is as follows: first, any superset of $\bar{A}$ is a qualified set in Γ^*, due to Lemma 1 and Γ^*

is a monotone access structure. Second, if A' is a qualified set, the perfect secret sharing scheme on Γ^* and the perfect binding commitment will guarantee that the composed challenge is the same one since the shares in A' are the same. Same composed challenge contradicts to the definition of tree of transcripts. From Lemma 1 we can conclude that A' is an authorized set in Γ^* and thus shares $\{c_i\}_{i\in A'}$ uniquely determine a certain challenge c. From the tree of transcripts in round $l \in [\mu]$, the algorithm $\mathsf{CheckShares}(c^l, \{c_i^l\}_{i\in[n]})$ can only be satisfied if all the challenges $\{c_k^{l,k}\}_{k\in[k_l]}$ are equal, which contradicts to the fact that those challenges are distinct according to the definition of $(k_1, \ldots, k_\mu)$-tree of accepting transcripts.

Now $A' = \bar{A}$, which implies that all the shares in $\bar{A}$ are fixed in advance. Here A' is the set that the commitments work in the perfect binding mode, which means that the committed (shared) challenges must be the same, since the commitments are given in the root node of the tree of transcripts (they are included in the first message). The extractor then runs extractors $\mathcal{E}_{\Pi_i}$ for all $i \in A$ expecting to obtain witnesses $\{w_i\}_{i\in A}$. We then need to prove that in each round $l \in [\mu]$, distinct challenges yield all-different shared challenges $\{c_i^l\}_{i\in A}$. For the sake of brevity we omit round index l here and assume that there are k distinct challenges $\{c^j\}_{j\in[k]}$ in this round. We prove that for all $m, n \in [k]$ such that challenges $c^m \neq c^n$ and for all $i \in A$, the shared challenges $c_i^m \neq c_i^n$. Suppose that there exist some indices $m, n \in [k]$ such that $c_i^m = c_i^n$. From Lemma 1 we obtain that $\{c_j^m\}_{j\in\bar{A}} \cup \{c_i^m\}$ (resp. n) is an authorized set in Γ^*. Then, the algorithm $\mathsf{CheckShares}$ would only be satisfied on index m, n if $c^m = c^n$, which contradicts to $c^m \neq c^n$. Finally, we can split a $(k_1, \ldots, k_\mu)$-tree of accepting transcripts for any underlying protocol $\Pi_i(x_i, w_i)$ for all $i \in A$. The extractor $\mathcal{E}_{\Pi_i}$ is able to extract witnesses w_i successfully due to the $(k_1, \ldots, k_\mu)$-special soundness of Π_i.

We stress that the above proof is almost orthogonal to the soundness of underlying protocols, except that in the last step the extractor $\mathcal{E}$ invokes special soundness extractors $\mathcal{E}_{\Pi_i}$ for all $i \in A$. In the perfect special soundness case, $\mathcal{E}_{\Pi_i}$ should always extract witnesses $\{w_i\}_{i\in A}$. When it comes to statistical special soundness, the extractors are only able to extract witnesses except with some special soundness error κ_i. Due to the smoothness of the perfect secret sharing scheme, the perfect binding property of the dual-mode commitment and perfect 2-special soundness of $\Pi_{\Gamma'}^{\mathsf{bind}}$, the challenges $\{c_i^l\}_{i\in A}$ are uniformly distributed and independent of A. The overall success probability to extract witnesses $\{w_i\}_{i\in A}$ is $\prod_{i\in A}(1-\kappa_i) \geq 1-\sum_{i\in A}\kappa_i$, which proves that the composition is still statistical $(k_1, \ldots, k_\mu)$-special sound with error $\sum_{i\in A}\kappa_i$.

Special Honest Verifier Computational Zero-Knowledge. A special honest verifier zero-knowledge simulator can be constructed straightforwardly based on the facts that dmCom is mode indistinguishable, Π_i for all $i \in [n]$ are honest verifier zero-knowledge, and $\Pi_{\Gamma'}^{\mathsf{bind}}$ preserves special honest verifier zero-knowledge of Π^{bind} due to the property of CDS composition. □

3.4 Performance

The total amount of communication in $\Pi^{\mathsf{sound}}_{\Gamma}$ is $|\Pi^{\mathsf{sound}}_{\Gamma}| = |\Pi^{\mathsf{CDS}}_{\mu,\Gamma}| + |\Pi^{\mathsf{bind}}_{\Gamma'}| + n(|\mathsf{ck}| + |\mathsf{com}| + |r_{\mathsf{com}}|)$ where $|\Pi^{\mathsf{CDS}}_{\mu,\Gamma}| = \sigma/\rho \cdot |\Pi_i| + \mu \cdot \sigma$ for total share size σ of SSS_{Γ^*} and challenge size ρ is the communication complexity of the generalized CDS in Sect. 2.5. $|\Pi^{\mathsf{bind}}_{\Gamma'}|$ is the communication complexity of $\Pi^{\mathsf{bind}}_{\Gamma'}$, which is expanded to $\sigma'/\rho' \cdot |\Pi^{\mathsf{bind}}| + \sigma'$ where σ' is the total share size of $\mathsf{SSS}_{\Gamma'}$ and ρ' is the challenge size of Π^{bind}.

Regarding offline/online computation, observe the prover algorithm in round 1. Computation for $i \in \bar{A}$ can be done for all $i \in [n]$ in advance, and those for $i \in A$ can be done for all i knowing the witness. Thus, these computations can be done offline without Γ. Computing $\mathcal{A}_{\Pi^{\mathsf{bind}}_{\Gamma'}}$ depends on Γ'. Thinking of its actual computation more carefully, it simulates on $i \in A$ and runs the prover algorithm on $i \in \bar{A}$. Since both binding and hiding keys are prepared for each $i \in [n]$, they are precomputable for all $i \in [n]$ without Γ'. Accordingly, almost all the prover algorithm in round 1 can be done offline for n fixed statements.

The online computation includes running prover algorithm $\mathcal{Z}^l_j$ for $j \in A$. It also computes functions $\mathsf{Complete}_{\Gamma^*}$ and $\mathsf{Eqv}_{\mathsf{td}_j}$ for secret sharing that would require small number of field arithmetic. It remains to consider $\mathcal{Z}_{\Pi^{\mathsf{bind}}_{\Gamma'}}$. Though it runs the real prover algorithm for $i \in \bar{A}$, the last-round computation of a Σ-protocol is usually much less expensive than the first-message computation. It is indeed the case in our instantiation in Appendix A where the last-round message is computed by two arithmetic operations in $\mathbb{Z}_p$, which is often ignored compared to exponentiation in computing the first-round messages. Accordingly, the Γ'-dependent change of the computational cost for $\mathcal{Z}_{\Pi^{\mathsf{bind}}_{\Gamma'}}$ would be small compared to the gain obtained by precomputation.

3.5 Extensions

In the construction in Fig. 1, Π_i is honest verifier zero-knowledge and SSS_{Γ^*} is smooth. We can extend Γ^* to admit semi-smooth SSS_{Γ^*}, by requiring Π_i to be *special* honest verifier zero-knowledge. This is analogous to the general CDS case, as mentioned in Sect. 2.5.

If zero-knowledge is to be preserved, one can do so by replacing the perfectly binding dual-mode commitment scheme with a computationally binding one. To setup the keys based on a computational assumption, the protocol would require more interaction or a stronger assumption like a common reference strings or random oracles. We observe that the non-interactive partially binding vector commitment scheme in [27,29] is not suitable for our construction due to its inability to provide a witness of a binding key, but it could be modified for this purpose.

Finally, we stress that our framework can be easily modified to support computational special soundness as well with no harm to the main body of the framework: the atomic protocol can be regarded as a proof of knowledge of knowing a witness or a solution to computational hard problems. Our extractor

will invoke the (computational special soundness) extractors as black-boxes to extract witnesses of atomic protocols. As a result, the composed protocol is also computational special sound.

4 Complexity of Adversarial Structure

In this section we sometimes view access structures over $[n]$ as Boolean functions from $\{0,1\}^n$ to $\{0,1\}$ by identifying subsets of $[n]$ with their indicator vectors in $\{0,1\}^n$.

4.1 A Lower Bound on the Circuit Size of Γ'

Theorem 3. *Assuming* $\mathsf{NP} \not\subseteq \mathsf{P/poly}$*, for all constants* c *and* s_0 *there exists a monotone 2CNF* Γ *of size* $s \geq s_0$ *for which the circuit complexity of* Γ' *is greater than* s^c*.*

As any reasonable monotone model of computation, including monotone circuits over some fixed basis, monotone span programs, and monotone branching programs can be simulated by (possibly nonmonotone) circuits with polynomial blowup in size, Theorem 3 rules out the existence of a polynomial-size Γ' in any of these models when Γ is a 2CNF.

A monotone 2CNF Γ has a faithful graph-theoretic representation by associating variables with vertices and clauses with edges. In this representation, Γ' has the following interpretation:

Claim. Assuming Γ is a monotone 2CNF over n variables, Γ' consists of those subsets of $[n]$ that contain some maximal independent set of the graph representing Γ.

Proof. Let G be the graph representation of Γ. The satisfying assignments of Γ are the vertex covers of G. The minterms of Γ are the minimal vertex covers. By definition, the minterms of Γ' are the complements of the minimal vertex covers, namely the maximal independent sets. □

To prove Theorem 3 we construct a fixed family of graphs $\{G_n\}$, where G_n has $64\binom{n}{3}$ vertices, for which the following problem is NP-hard:

> Maximal Independent Superset (MISS): On inputs 1^n and $Y \subseteq [64\binom{n}{3}]$, does Y contain a maximal independent set of G_n?

Proof (Theorem 3). We prove the contrapositive. Assume the adversarial structures $\{\Gamma'_n\}$ for the 2CNF family represented by $\{G_n\}$ had size polynomial in $|G_n|$ and so polynomial in n. Then $\{\Gamma'_n\}$ is a polynomial-size circuit family for MISS. As MISS is NP-hard, every NP-problem has a polynomial-size circuit family, hence $\mathsf{NP} \subseteq \mathsf{P/poly}$. □

The NP-hardness of MISS is proved by reducing from 3SAT. The proof is inspired by Karp's canonical proof for NP-hardness of independent set.

Proof (NP-hardness of MISS). There are two types of vertices in G_n. Type 1 vertices are labeled by the $56\binom{n}{3}$ pairs (c, a), where c ranges over all $8\binom{n}{3}$ possible 3CNF clauses in n variables and a ranges over the 7 possible satisfying assignments of c. For example, $(x_2 \vee \overline{x_3} \vee x_5, 001)$ is a vertex in G_n because $x_2 = 0, x_3 = 0, x_5 = 1$ is a satisfying assignment for this clause. Type 2 vertices are labeled by $(c, \star)$ where c again ranges over possible 3CNF clauses. There are two types of edges in G_n:

1. Inconsistency edges between type 1 vertices $(c, a), (c', a')$ when assignments a and a' are inconsistent. For example, there is an edge between $(x_2 \vee \overline{x_3} \vee x_5, 001)$ and $(x_3 \vee x_5 \vee x_6, 001)$ because the corresponding assignments to x_5 are inconsistent.
2. Clique edges between all 8 vertices (c, a) including when $a = \star$ for every c.

In particular the clique edges ensure that no independent set of size greater than $8\binom{n}{2}$ exists in G_n.

The portion of the G_n induced by the 16 vertices of the form $(x_2 \vee \overline{x_3} \vee x_5, \cdot)$ and $(x_3 \vee x_5 \vee x_6, \cdot)$ is illustrated in Fig. 2.

Given a 3CNF ϕ with n variables the reduction from 3SAT to MISS outputs the set

$$Y = \{(c, a) \text{ of type } 1\colon c \text{ is a clause of } \phi\} \cup \{(c, \star) \text{ of type } 2\colon c \text{ is not a clause of } \phi\}.$$

We now argue that the reduction is correct: ϕ is satisfiable if and only if Y contains some maximal independent set in G_n.

Assume ϕ is satisfiable and let x be the satisfying assignment. Let I be the set consisting of all type 1 vertices (c, a) where a is the partial assignment of x to the variables in c together with all type 2 vertices $(c, \star)$ where c is not a clause of ϕ. By construction I is a subset of Y, and moreover I contains exactly one vertex of the form $(c, \cdot)$ for all $8\binom{n}{3}$ possible n-variate 3CNF clauses c. Therefore, I is maximal.

If ϕ is not satisfiable then every independent set $I \subseteq Y$ must miss all type 1 vertices (c, a) for at least one clause c of ϕ. Otherwise, the partial assignments a would be mutually consistent, so they could be stitched together into a satisfying assignment for ϕ. Then $I \cup \{(c, \star)\}$ is also an independent set in G_n so I is not maximal. □

4.2 Efficient Adversarial Structures for Read-Once Formulas

A formula is a circuit with fan-out one. The formula is read-once if all leaves are labelled by different variables. The size of the formula is the number of leaves.

Theorem 4. *If $\Gamma\colon \{0,1\}^n \to \{0,1\}$ is computed by a size-s read-once AND/OR formula with $s_\vee$ fan-in two OR gates, then Γ' can be computed by an AND/OR formula of size at most $s + n \cdot s_\vee$.*

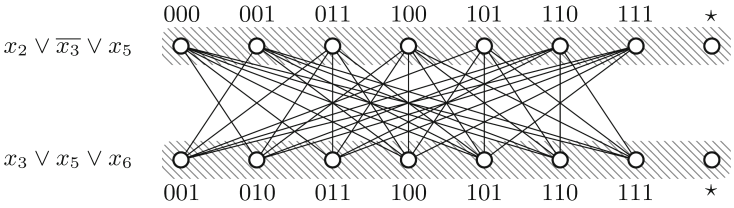

Fig. 2. The portion of G_n induced by the 16 vertices of the form $(x_2 \vee \overline{x_3} \vee x_5, \cdot)$ and $(x_3 \vee x_5 \vee x_6, \cdot)$. The hatched rectangles represent cliques.

The class mNC^1 [31] consists of all Boolean function families that are represented by polynomial-size formulas over some fixed complete basis of monotone functions. The basis may be fixed to $\{AND, OR\}$ without loss of generality. Equivalently mNC^1 can be defined as the class of bounded fan-in monotone circuits of depth $O(\log n)$.

In the context of secure computation, mNC^1 contains access structures of interest beyond a single threshold. This allows greater flexibility with fine-grained access control in credential systems. Popular vote aggregation rules such as thresholds of thresholds (Corollary 2) and ranked choice (as in Australian elections) can be efficiently implemented in mNC^1. The resulting composed zero-knowledge proof assures correctness of the ballot count.

As a consequence of Theorem 4 we obtain the closure of read-once projections to mNC^1 under the transformation $\Gamma \to \Gamma'$. We say that Boolean function $f(y_1, \ldots, y_m)$ *projects to* $g(x_1, \ldots, x_n)$ if there exists a map $\iota\colon [m] \to [n]$ for which $g(x_{\iota(1)}, \ldots, x_{\iota(m)}) = f(x_1, \ldots, x_n)$ for all x.

Corollary 1. *For every $\Gamma \in \mathsf{mNC}^1$ there exists a read-once family Γ_{RO} that projects to Γ and for which Γ'_{RO} is in mNC^1. Moreover, Γ'_{RO} can be computed from Γ in polynomial time.*

As a special case, we obtain the following bound for bounded-depth formulas with threshold gates:

Corollary 2. *If Γ is a depth-d formula over unweighted threshold gates of arity at most k then Γ'_{RO} has AND/OR formula size at most $k^{O(d)}$.*

Proof. As majority over k bits has AND/OR formula size $k^{5.3}$ [4,40], any threshold of arity at most k can be computed by a size-$(2k)^{5.3}$ AND/OR formula. Replacing each threshold gate with this formula gives a size-$(2k)^{5.3d}$ AND/OR formula for Γ. As the number of inputs to Γ is at most k^d by Theorem 4 the AND/OR formula size of Γ'_{RO} is at most $(2k)^{6.3d}$. □

A possible direction for improving the constant 6.3 in the exponent is to generalize the composition rules (1) below to arbitrary thresholds.

Proof. (Theorem 4). The formula for Γ' is inductively constructed from the formula for Γ using the following composition rules (assuming Γ and Δ are not constant):

$$\begin{aligned}(\Gamma(x) \wedge \Delta(y))' &= \Gamma'(x) \wedge \Delta'(y) \\ (\Gamma(x) \vee \Delta(y))' &= (\Gamma'(x) \wedge (\wedge_i y_i)) \vee (\Delta'(y) \wedge (\wedge_i x_i)) \\ id' &= 1.\end{aligned} \tag{1}$$

In the AND and OR rules, the sets of x and y-variables $X = \{x_1, \ldots, x_n\}$ and $Y = \{y_1, \ldots, y_m\}$ are assumed to form a partition (they are disjoint and cover all variables). In the last rule, $id\colon \{0,1\} \to \{0,1\}$ stands for the identity function $id(x) = x$.

The claimed complexity bound follows by strong induction over size. It remains to verify the correctness of the composition rules. We will show that the formulas on the left and right hand side of all three rules have the same minterms and are therefore logically equivalent. Let $\mathcal{M}$ and $\mathcal{N}$ stand for the minterms of Γ and Δ. We identify the ground set with $X \cup Y$.

For the AND rule, the minterms of $\Gamma \wedge \Delta$ are $A \cup B$ where $A \in \mathcal{M}$ and $B \in \mathcal{N}$: All such sets are accepted, and omitting any element from either causes Γ or Δ to reject. As A and B are disjoint, the minterms of $(\Gamma \wedge \Delta)'$ are all sets of the form $(X \cup Y) \setminus (A \cup B) = (X \setminus A) \cup (Y \setminus B)$, which are the unions of all minterms of Γ' and all minterms of Δ'. These are the minterms of $\Gamma' \wedge \Delta'$.

For the OR rule, the minterms of $\Gamma \vee \Delta$ are $\mathcal{M} \cup \mathcal{N}$: All of the sets are accepted, and owing to disjointness taking out an element from a set in say $\mathcal{M}$ will cause both Γ and (by default) Δ to reject. The minterms of $(\Gamma \vee \Delta)'$ are therefore all sets of the form $(X \cup Y) \setminus A = (X \setminus A) \cup Y$ for $A \in \mathcal{M}$ as well as $(X \cup Y) \setminus B = X \cup (Y \setminus B)$ for $B \in \mathcal{N}$. These are the minterms of $\Gamma'(x) \wedge (\wedge_i y_i)$ and $\Delta'(y) \wedge (\wedge_i x_i)$, respectively. We argue that their union equals the minterms of the OR of these two formulas: Clearly the OR accepts all terms in the union. Moreover, removing any element, say x_i causes both these formulas to reject.

For the identity rule, id has a unique minterm $\{x\}$ and its complement is the empty set, which is the unique minterm of the constant function 1. □

5 Extending Challenge Space for k-Special Sound Protocols

In the original protocol of [22], if the challenge space of the Σ-protocol is bigger than the domain of the secret sharing scheme, that is, the bit lengths of challenges in the Σ-protocols are larger than the lengths of shared challenges, then there exist natural maps from shared secrets to shared challenges. This is also true when the two spaces are the same. The problem occurs when the challenge space is smaller than the domain of secret sharing. Some shared secrets can not be distributed to the challenge space naturally.

[22] proposed the following theorem, i.e., parallel repetition to solve this issue for Σ-protocols.

Theorem 5 (Challenge-length Amplification [17,22]). *Let Π be a Σ-protocol for relation R and challenge length l. Running Π t-times in parallel for the same instance x corresponds to running Σ-protocol for R with challenge length $t \cdot l$.*

Parallel repetition works perfectly for Σ-protocols, which is 2-special sound. Unfortunately, such method does not apply to k-special sound protocols with secret sharing. First recall the definition of k-special soundness: the extractor takes as input a statement x and k accepting transcripts $(a, c_1, z_1), \ldots, (a, c_k, z_k)$ with common first message a and pairwise distinct challenges $c_1, \ldots, c_k$ outputs a witness w s.t. $R(x, w) = 1$. In a t-fold parallel repetition, each transcript $\Pi_i, i \in [k]$ consists of t-tuples $(a^1, c_i^1, z_i^1), \ldots, (a^t, c_i^t, z_i^t)$. The extractor expects that given k distinct all accepting transcripts all least one out of t tuples contains k pairwise distinct transcripts $(a^j, c_1^j, z_1^j), \ldots, (a^j, c_k^j, z_k^j)$. The expectation cannot be satisfied in general, since for any distinct transcript Π_i and Π_j there may be only one position out of t that is different. In the worst case, only after collecting $((k-1)^t + 1)$ distinct transcripts will the extractor extract a witness. In fact, a t-fold k-special sound protocol is actually a $((k-1)^t+1)$-special sound protocol.

We attempt to solve this issue in this section. Our intuition is to guarantee that when collecting k t-wise accepting transcripts, there exists at least 1-out-of-t repetition having k-distinct valid transcripts with high probability. We show that if the k-special soundness property allows for some negligible knowledge error, then a t-fold parallel repetition of a perfect k-special sound protocol is still k-special sound with negligible error.

In Sect. 5.1, we first start from the simpler case of k-special sound Σ-protocols. We prove that parallel repetition of a perfect k-special sound protocol is still statistical k-special sound. In Sect. 5.2, we further extend this theorem to the multi-round version.

5.1 Statistical k-Special Soundness

We prove that if the k-special soundness property allows for some negligible knowledge error, then a t-fold parallel repetition of a perfect k-special sound protocol is still statistical k-special sound with negligible error. The definition of k-special soundness is modified as follows: the extractor which takes as input a statement x and k accepting transcripts $(a, c_1, z_1), \ldots, (a, c_k, z_k)$ with common first message a and pairwise distinct challenges $c_1, \ldots, c_k$ outputs a witness w s.t. $R(x, w) = 1$ except for some negligible probability κ. We denote this κ as the special soundness error.

Definition 13 (Statistical k-Special Soundness). *A 3-round public-coin proof is statistical k-special sound if there exists a polynomial-time algorithm that, given any k distinct transcripts $\left(x, a, c_{i\in[k]}, z_{i\in[k]}\right)$ that satisfies $c_i \neq c_j$ for every $i \neq j$ and $\mathcal{V}(x, a, c_i, z_i) = 1$ for $i \in [k]$, outputs w such that $R(x, w) = 1$ except for some negligible probability κ (over the choice of challenges). We denote this κ as the special soundness error.*

Theorem 6. *The t-fold parallel repetition Π^t of a perfect k-special sound protocol Π with challenge space $\mathcal{C}$ is a statistical k-special sound protocol with special soundness error κ decaying exponentially with the growth of t.*

Proof. The extractor is given k t-fold accepting transcripts. Each t-fold transcript has at least one challenge different from other transcripts (in the same fold). The extractor expects that if t is large enough, then with high probability in one certain fold all the challenges are pairwise distinct. We can view the game as follows:

Suppose a t-column matrix C of which the elements $c \in \mathcal{C}$. Without loss of generality we assume $|\mathcal{C}| = q$. In matrix C every row is different from each other row (at least for one column). There are q^t such rows. Each row represents possible challenges for t-fold accepting transcripts. Now select k rows from q^t rows, which compose a k-by-t sub-matrix. The special soundness error κ is the probability that no columns has all pairwise distinct elements. We denote B as the set of all such sub-matrices. We have

$$|B| \leq (q^k - \frac{q!}{(q-k)!})^t.$$

The inequality is because the right side takes repeated rows into account. Let S be the set of all possible sub-matrices chosen from accepting q^t rows, then

$$|S| = \frac{q^t!}{(q^t-k)!}.$$

Then the special soundness error κ can be computed as:

$$\begin{aligned}
\kappa = & \frac{|B|}{|S|} \\
& \leq \frac{(q^k - \frac{q!}{(q-k)!})^t}{\frac{q^t!}{(q^t-k)!}} \\
& \leq \frac{(q^k - (q^k - kq^{k-1}))^t}{q^{tk} - kq^{t(k-1)}} \\
& = \frac{(kq^{k-1})^t}{q^{tk} - kq^{t(k-1)}} \\
& = \frac{k^t}{q^t - k} \\
& = \frac{1}{(\frac{q}{k})^t - \frac{k}{k^t}}.
\end{aligned}$$

Since $k < q$, $\kappa \leq \frac{1}{(\frac{q}{k})^t - \frac{k}{k^t}}$ decays exponentially with the growth of t. By setting t slightly larger than $\frac{\lambda}{\log \frac{q}{k}}$, the special soundness error $\kappa \leq 2^{-\lambda}$ which is negligible in λ.

It remains to prove the case of $k = q$ which often happens when q is small. We have

$$\begin{aligned}\kappa = & \frac{|B|}{|S|} \\ & \leq \frac{(q^q - q!)^t}{\frac{q^t!}{(q^t-q)!}} \\ & \leq \frac{(q^q - (\frac{q}{e})^q)^t}{(\frac{q^t}{e})^q} \\ & = e^q(1 - \frac{1}{e^q})^t \\ & \leq e^{q - \frac{t}{e^q}}.\end{aligned}$$

By setting t slightly larger than $(\lambda+q)e^q$, the special soundness error $\kappa \leq 2^{-\lambda}$ which is negligible in λ. □

Next we prove that statistical k-special soundness still implies knowledge soundness.

Theorem 7. *Let Π be a statistical k-special sound protocol with challenge space $\mathcal{C}$ and special soundness error κ. Then Π is knowledge sound with knowledge error $\kappa_{\mathsf{KS}} = (k-1)/q + \kappa$, where $q = |\mathcal{C}|$ denotes the size of challenge space $\mathcal{C}$.*

Proof. The proof idea is straightforward. The knowledge extractor $\mathcal{E}_{\mathsf{KS}}$ can be seen as a combination of the following procedures:

1. Given the cheating prover P^* the knowledge extractor invokes a tree builder to extract sufficient accepting transcripts;
2. Given sufficiently many accepting transcripts, the knowledge extractor runs corresponding extractor for special soundness to extract a witness.

The tree builder in Step 1 is exactly the same as any of the classical results for perfect k-special soundness [7,8,19]. The probability for failing to extract sufficient amount of accepting transcripts is bounded by $(k-1)/q$. In Step 2, the corresponding extractor has a probability κ of failure. Thus, the total knowledge error $\kappa_{\mathsf{KS}} \leq (k-1)/q + \kappa$, which completes the proof. □

5.2 Statistical Special Soundness for Multi-round Protocols

We now analyze the previous results in the multi-round setting as presented in Definition 12. We stress that parallel repetition of multi-round protocols suffer from exactly the same issue as the composition of multi-round protocols. The tree of transcripts does not guarantee that pairwise distinct challenges can be obtained at the same position among different rounds, leading to a failed extraction by the extractor.

Theorem 8. *The t-fold parallel repetition Π^t of a perfect $(k_1, \ldots, k_\mu)$-special sound protocol Π with challenge space $\mathcal{C}_1 \times \cdots \times \mathcal{C}_\mu$ is a statistical $(k_1, \ldots, k_\mu)$-special sound protocol with special soundness error κ decaying exponentially with the growth of t.*

Proof. Similarly, the extractor is given a t-fold $(k_1, \ldots, k_\mu)$-tree of accepting transcripts. Each t-fold transcript has at least one challenge different from other transcripts (in the same fold) in each round. The extractor expects that if t is large enough, then with high probability in one certain fold all the challenges are pairwise distinct. We can view the game as follows:

Suppose a t-column matrix $C = C_1||C_2|| \ldots ||C_\mu$ concatenated by rows. The elements in C_i are chosen from $\mathcal{C}_i$. Without loss of generality we assume $|\mathcal{C}_i| = q_i$. In matrix C_i each row is different from any other row (at least at one column). There are q_i^t such rows. Each row represents possible challenges for t-fold accepting transcripts. Now select k_i rows from q_i^t rows, which compose a k_i-by-t sub-matrix. The special soundness error κ is the probability that for all matrices $C_i (i \in [\mu])$ no columns have all pairwise distinct elements. We denote B as the set of all such sub-matrices. We have

$$|B| \leq (\prod_i q_i^{k_i} - \prod_i \frac{q_i!}{(q_i - k_i)!})^t.$$

The inequality is because the right side takes repeated rows into account. Let S be the set of all possible sub-matrices chosen from all accepting rows, then

$$|S| = \prod_i \frac{q_i^t!}{(q_i^t - k_i)!}.$$

Then the special soundness error κ can be computed as:

$$\begin{aligned}
\kappa = & \frac{|B|}{|S|} \\
\leq & \frac{(\prod_i q_i^{k_i} - \prod_i \frac{q_i!}{(q_i - k_i)!})^t}{\prod_i \frac{q_i^t!}{(q_i^t - k_i)!}} \\
\leq & \frac{(\prod_i q_i^{k_i} - \prod_i (q_i^{k_i} - k_i q_i^{k_i - 1}))^t}{\prod_i (q_i^{tk_i} - k_i q_i^{t(k_i - 1)})} \\
= & \frac{(\prod_i q_i - \prod_i (q_i - k_i))^t}{\prod_i (q_i^t - k_i)} \\
= & \frac{(\prod_i \frac{q_i}{k_i} - \prod_i (\frac{q_i}{k_i} - 1))^t}{\prod_i ((\frac{q_i}{k_i})^t - \frac{k_i}{k_i^t})}.
\end{aligned}$$

By setting t slightly larger than $\frac{\lambda}{\sum_i \log \frac{q_i}{k_i}}$, the special soundness error $\kappa \leq 2^{-\lambda}$ which is negligible in λ.

The case where there exists some $k_i = q_i$ can be directly combined from the above proof and the case of $k = q$ in Theorem 6.

Specifically, for the case of $k_i = q_i$ for all $i \in [\mu]$, we have

$$\begin{aligned}
\kappa = & \frac{|B|}{|S|} \\
& \leq \frac{(\prod_i q_i^{q_i} - \prod_i q_i!)^t}{\prod_i \frac{q_i^t!}{(q_i^t - q_i)!}} \\
& \leq \frac{(\prod_i q_i^{q_i} - \prod_i (\frac{q_i}{e})^{q_i})^t}{\prod_i (\frac{q_i^t}{e})^{q_i}} \\
& = e^{\sum_i q_i} (1 - \frac{1}{e^{\sum_i q_i}})^t \\
& \leq e^{\sum_i q_i - \frac{t}{e^{\sum_i q_i}}}.
\end{aligned}$$

By setting t slightly larger than $(\lambda + \sum_i q_i) e^{\sum_i q_i}$, the special soundness error $\kappa \leq 2^{-\lambda}$ which is negligible in λ. □

Finally, similar to statistical k-special soundness, we prove that statistical $(k_1, \ldots, k_\mu)$-special soundness implies knowledge soundness.

Theorem 9. *Let Π be a $(2\mu + 1)$-round statistical $(k_1, \ldots, k_\mu)$-special sound protocol with challenge space $\mathcal{C}_1 \times \mathcal{C}_2 \times \cdots \times \mathcal{C}_\mu$ and special soundness error κ. Then Π is knowledge sound with knowledge error $\kappa_{\mathsf{KS}} = 1 - \prod_{i=1}^{\mu} \frac{q_i - k_i + 1}{q_i} + \kappa$, where $q_i = |\mathcal{C}_i|$ denotes the size of challenge space $\mathcal{C}_i$.*

Proof. Similar to Theorem 7, the proof idea is straightforward. The knowledge extractor $\mathcal{E}_{\mathsf{KS}}$ can be seen as a combination of the following procedures:

1. Given the cheating prover P^* the knowledge extractor invokes a tree builder to extract sufficient accepting transcripts;
2. Given sufficiently many accepting transcripts, the knowledge extractor runs corresponding extractor for special soundness to extract a witness.

The tree builder in Step 1 is exactly the same as any of the classical results for perfect $(k_1, \ldots, k_\mu)$-special soundness [7,8]. The probability for failing to extract sufficient amount of accepting transcripts is bounded by $1 - \prod_{i=1}^{\mu} \frac{q_i - k_i + 1}{q_i}$. In Step 2, the corresponding extractor has a probability κ of failure. Thus, the total knowledge error $\kappa_{\mathsf{KS}} \leq 1 - \prod_{i=1}^{\mu} \frac{q_i - k_i + 1}{q_i} + \kappa$, which completes the proof. □

6 Conclusion

In this paper we present multi-round and k-special sound adaptation of CDS composition that features soundness preserving, offline simulation, mNC^1 access structure, and round preserving in the plain model. We also study about the complexity of adversarial structures and extension of challenge space for k-special sound protocols.

Some interesting open problems remain:

- Can we have a multi-round composition preserving zero-knowledge, and soundness at the same time as the original CDS does?
- If the above is not possible, can we achieve the same goal with post-quantum security? It requires a quantum-secure dual-mode commitment scheme with demanded properties and associated Σ-protocols. Note that our composition is already sound against unbound, including quantum, prover even though the mentioned building blocks are instantiated in the classical setting.
- Can we go beyond $\Gamma \in \mathsf{mNC}^1$, e.g., MSP, without paying too much cost? Since we have shown that computing Γ' from Γ is hard, the construction must eliminate proofs for Γ' requiring a novel approach.

Acknowledgements. We are grateful to Siyao Guo for useful advice and anonymous reviewers from Eurocrypt 2024 and Crypto 2024 for helpful comments. Work supported by European Research Council (ERC) under the EU's Horizon 2020 research and innovation programme (Grant agreement No. 101019547) and Cariplo CRYPTONOMEX grant.

A Dual-Mode Commitment Scheme from DDH

This section presents an instantiation of a dual-mode commitment scheme associated with a Σ-protocol for proving correct generation of a binding key. We first describe the dual-mode commitment scheme taken from [38] with a modification that the prover generates the public parameters, $(\mathbb{G}, q, g)$, and chooses a random generator h by itself while they are given as a common reference string in [38]. This results in losing the original computational binding property in the hiding mode, which we do not need in our construction. It is noted that it does not spoil the computational hiding property in the binding mode as well as the mode indistinguishability where the adversary is the verifier. We instantiate $\mathsf{dmCom} = \{\mathsf{CGen}_{\mathsf{BIND}}, \mathsf{CGen}_{\mathsf{HIDE}}, \mathsf{Com}, \mathsf{Eqv}\}$ as follows:

- $\mathsf{CGen}_{\mathsf{BIND}}(1^\lambda) \to \mathsf{ck}$: Run $(\mathbb{G}, q, g) \leftarrow \mathcal{G}(1^\lambda)$. Choose $h \leftarrow\!\!\$\, \mathbb{G}$, and $(\rho_1, \rho_2) \leftarrow\!\!\$\, \mathbb{Z}_q^2$. Compute $u = g^{\rho_1}$ and $v = g^{\rho_2}$, and output $\mathsf{ck} := (\mathbb{G}, q, g, h, u, v)$.
- $\mathsf{CGen}_{\mathsf{HIDE}}(1^\lambda) \to (\mathsf{ck}, \mathsf{td})$: As above, except for choosing $\rho \leftarrow\!\!\$\, \mathbb{Z}_q$, and computing $u = g^\rho$, and $v = h^\rho$. Output $\mathsf{ck} := (\mathbb{G}, q, g, h, u, v)$, and $\mathsf{td} := \rho$.
- $\mathsf{Com}_{\mathsf{ck}}(m; r) \to \mathsf{com}$: Choose $r \leftarrow\!\!\$\, \mathbb{Z}_q$ and compute $a = g^r/u^m$ and $b = h^r/v^m$. Output $\mathsf{com} := (a, b)$.
- $\mathsf{Eqv}_{\mathsf{td}}(\mathsf{com}, \mathsf{ck}, m', r', m) \to r$: Take ρ from td and output $r := (r' - \rho m') + \rho m$.

Theorem 10. *The above* dmCom *is a dual-mode commitment scheme. It is perfectly binding in the binding mode, and equivocal in the hiding mode. It is mode indistinguishable and computationally hiding in the binding mode under the decision Diffie-Hellman assumption relative to* $\mathcal{G}$.

Proof. For a commitment, (a, b), let $\tilde{a} := \log_g a$, $\tilde{b} := \log_h b$ and $x := \log_g h$. We first show that dmCom is perfectly binding in the binding mode. From $a = g^r/u^m$

and $b = h^r/v^m$, we have $\tilde{a} = r + \rho_1 m$ and $\tilde{b} = r + \rho_2 m$. Since $\rho_1 \neq \rho_2$, they determine $m \in \mathbb{Z}_q$ uniquely by $m = (\tilde{a} - \tilde{b})/(\rho_1 - \rho_2)$. To show that commitments are computationally hiding in the binding mode under the DDH assumption, consider two distributions, $C(m)$ and R, over variables (g, h, u, v, a, b) that (g, h, u, v) is generated by $\mathsf{CGen}_{\mathsf{BIND}}(1^\lambda)$, and (a, b) for $C(m)$ is generated by $\mathsf{Com}_{\mathsf{ck}}(m)$ with $\mathsf{ck} := (\mathbb{G}, q, g, h, u, v)$, and (a, b) for R is uniform over $\mathbb{G}^2$. We then set up the following hybrid. For a given DDH question $Q := (g, g^r, h, h^{r'})$ and message m, let $(u, v) \leftarrow\!\$\, \mathbb{G}^2$ and $(a, b) := ((g^r)/u^m, (h^{r'})/v^m)$. If Q is taken from the DDH distribution where $r = r'$, (g, h, u, v, a, b) is in $C(m)$. On the other hand, if Q is taken from the random distribution where $r \neq r'$, (g, h, u, v, a, b) is in R. Thus, distinguishing $C(m)$ and R is hard under the DDH assumption. For $m'(\neq m)$, $C(m')$ and R are indistinguishable for the same reason. Thus, for any m and m', two distributions $C(m)$ and $C(m')$ are indistinguishable under the DDH assumption.

We next show that dmCom is perfectly hiding in the hiding mode. Since $u = g^\rho$ and $v = h^\rho$, we have $\tilde{a} = \tilde{b} = r - \rho m$. Thus, for any m', there exists r' that satisfies $\tilde{a} = \tilde{b} = r' - \rho m'$. Therefore, (r, m) is perfectly hiding.

Mode indistinguishability is directly from the DDH assumption. The quadruple (g, h, u, v) in a hiding key is in the DDH distribution, and the one in a binding key is in the uniform distribution. Thus, they are indistinguishable under the DDH assumption.

Finally, equivocality is verified by inspecting that $r' = (r - \rho m) + \rho m'$ satisfies $a = g^r/u^m = g^{r'}/u^{m'}$ and $b = h^r/v^m = h^{r'}/v^{m'}$.

□

We present a Σ-protocol Π^{bind} for proving that $\mathsf{ck} := (\mathbb{G}, q, g, h, u, v)$ is a binding key. We assume that group generator $\mathcal{G}$ is transparent, i.e., $(\mathbb{G}, q, g) \in \mathcal{G}(1^\lambda)$ can be verified publicly, and focus on the relation among (g, h, u, v). Concretely, the proof is for the following relation:

$$R^{\mathsf{bind}} := \{(\rho_1, \rho_2) \,|\, u = g^{\rho_1} \wedge v = h^{\rho_2} \wedge \rho_1 \neq \rho_2\}.$$

Inequality $\rho_1 \neq \rho_2$ is shown by checking $d_1^{(\rho_1 - \rho_2)} \neq 1_{\mathbb{G}}$ for a random generator d_1.

[Π^{bind} : Proof of Binding Key]

- Prover's private input is (ρ_1, ρ_2), and the common input to the prover and verifier is $\mathsf{ck} := (\mathbb{G}, q, g, h, u, v)$.
- Prover computes $d_1 \leftarrow\!\$\, \mathbb{G}$, $d_2 := d_1^{(\rho_1 - \rho_2)}$, $(r_1, r_2) \leftarrow\!\$\, \mathbb{Z}_q^2$, $a_1 := g^{r_1}$, $a_2 := h^{r_2}$, $a_3 := d_1^{r_1 - r_2}$, and sends $(a_1, a_2, a_3, d_1, d_2)$ to the verifier.
- Verifier selects $c \leftarrow\!\$\, \mathbb{Z}_q$ and send it to Prover.
- Prover computes $z_1 = r_1 - c\,\rho_1$, $z_2 = r_2 - c\,\rho_2$, and send (z_1, z_2) to Verifier
- Verifier accepts if $a_1 = g^{z_1} u^c$, $a_2 = h^{z_2} v^c$, $a_3 = d_1^{z_1 - z_2} d_2^c$, $d_1 \neq 1_{\mathbb{G}}$, and $d_2 \neq 1_{\mathbb{G}}$. Reject, otherwise.

Theorem 11. *Π^{bind} is a three-round public-coin proof protocol for relation R^{bind}. It is complete, 2-special sound, and honest verifier zero-knowledge.*

Proof. We focus on soundness and zero-knowledge properties since others are direct from the construction. To prove 2-special soundness, we construct an extractor, $\mathcal{E}$, as follows. Given a colliding transcript $(\mathsf{ck}, a_1, a_2, a_3, d_1, d_2, c, z_1, z_2, c', z_1', z_2')$ that satisfies

$$a_1 = g^{z_1} u^c = g^{z_1'} u^{c'}, \tag{2}$$

$$a_2 = h^{z_2} v^c = h^{z_2'} v^{c'}, \tag{3}$$

$$a_3 = d_1^{z_1 - z_2} d_2^c = d_1^{z_1' - z_2'} d_2^{c'}, \tag{4}$$

$c \neq c'$, $d_1 \neq 1$, and $d_2 \neq 1$, extractor $\mathcal{E}$ computes

$$\rho_1 := \frac{z_1' - z_1}{c - c'}, \quad \rho_2 := \frac{z_2' - z_2}{c - c'},$$

and outputs (ρ_1, ρ_2).

To verify that $\mathcal{E}$ is correct, observe that, from (2) and (3):

$$z_1 + \log_g u \cdot c = z_1' + \log_g u \cdot c' \Rightarrow \log_g u = \frac{z_1' - z_1}{c - c'} = \rho_1, \text{ and}$$

$$z_2 + \log_h v \cdot c = z_2' + \log_h v \cdot c' \Rightarrow \log_h v = \frac{z_2' - z_2}{c - c'} = \rho_2.$$

Also, from (4),

$$(z_1 - z_2) + c \log_{d_1} d_2 = (z_1' - z_2') + c' \log_{d_1} d_2 \Rightarrow \log_{d_1} d_2 = \frac{z_1' - z_1}{c - c'} - \frac{z_2' - z_2}{c - c'}$$

$$\Rightarrow \log_{d_1} d_2 = \log_g u - \log_h v = \rho_1 - \rho_2.$$

Since $d_1 \neq 1$ and $d_2 \neq 1$, $\log_{d_1} d_2 \neq 0$. Thus, we have $\rho_1 - \rho_2 \neq 0$, proving $\rho_1 \neq \rho_2$.

To prove honest verifier zero-knowledge, we construct a simulator, $\mathcal{S}$, as follows. Given instance $\mathsf{ck} = (\mathbb{G}, g, h, u, v)$, $\mathcal{S}$ picks c, z_1, z_2 uniformly and independently of $\mathbb{Z}_q$. It also chooses d_1 and d_2 uniformly and independently of $\mathbb{G}$. It then computes $\tilde{a}_1 := g^{\tilde{z}_1} u^c$, $\tilde{a}_2 := h^{\tilde{z}_2} v^c$, and $\tilde{a}_3 := d_1^{\tilde{z}_1 - \tilde{z}_2} d_2^c$. Finally, it outputs $(a_1, a_2, a_3, d_1, d_2, c, z_1, z_2)$.

We show that the above output from the simulator is computationally indistinguishable from that observed in a real protocol run. We construct a hybrid as follows. Given DDH question $Q := (g, A, B, C)$ relative to $\mathcal{G}(1^n)$, we set $u = A$, $v = h^{\rho_2}$, $d_1 = B$ and $d_2 = CB^{-\rho_2}$, and create other variables as prescribed by the simulator. This implies $\rho_1 = \log_g A$. If Q is in the DDH distribution, $C = B^{\rho_1}$. We then have $d_2 = B^{\rho_1 - \rho_2}$. Thus, the transcript is in the same distribution as the real prover outputs. On the other hand, if Q is in a random distribution, d_2 distributes uniformly as in the simulated transcript due to the randomness of C. Accordingly, the simulated and real transcripts are indistinguishable if the DDH assumption holds for $\mathcal{G}(1^n)$. This concludes the proof of honest verifier computational zero-knowledge of Π^{bind}.

□

References

1. Abe, M., Ambrona, M., Bogdanov, A., Ohkubo, M., Rosen, A.: Non-interactive composition of sigma-protocols via share-then-hash. In: Moriai, S., Wang, H. (eds.) ASIACRYPT 2020. LNCS, vol. 12493, pp. 749–773. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-64840-4_25
2. Abe, M., Ambrona, M., Bogdanov, A., Ohkubo, M., Rosen, A.: Acyclicity programming for sigma-protocols. In: Nissim, K., Waters, B. (eds.) TCC 2021. LNCS, vol. 13042, pp. 435–465. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-90459-3_15
3. Abe, M., Ohkubo, M., Suzuki, K.: 1-out-of-n signatures from a variety of keys. In: Zheng, Y. (ed.) ASIACRYPT 2002. LNCS, vol. 2501, pp. 415–432. Springer, Heidelberg (2002). https://doi.org/10.1007/3-540-36178-2_26
4. Ajtai, M., Komlós, J., Szemerédi, E.: An $o(nlogn)$ sorting network. In: STOC '83: Proceedings of the Fifteenth Annual ACM Symposium on Theory of Computing, pp. 1–9. ACM Press, New York (1983)
5. Ames, S., Hazay, C., Ishai, Y., Venkitasubramaniam, M.: Ligero: lightweight sublinear arguments without a trusted setup. In: Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, CCS 2017, Dallas, TX, USA, October 30 - November 03, 2017, pp. 2087–2104. ACM (2017)
6. Attema, T., Cramer, R., Fehr, S.: Compressing proofs of k-Out-Of-n partial knowledge. In: Malkin, T., Peikert, C. (eds.) CRYPTO 2021. LNCS, vol. 12828, pp. 65–91. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-84259-8_3
7. Attema, T., Cramer, R., Kohl, L.: A compressed σ-protocol theory for lattices. In: Malkin, T., Peikert, C. (eds.) CRYPTO 2021. LNCS, vol. 12826, pp. 549–579. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-84245-1_19
8. Attema, T., Fehr, S.: Parallel repetition of $(k_1, \ldots, k_\mu)$-special-sound multi-round interactive proofs. In: Dodis, Y., Shrimpton, T. (eds.) CRYPTO 2022, Part I. LNCS, vol. 13507, pp. 415–443. Springer, Cham (2022). https://doi.org/10.1007/978-3-031-15802-5_15
9. Attema, T., Fehr, S., Resch, N.: Generalized special-sound interactive proofs and their knowledge soundness. In: Rothblum, G., Wee, H. (eds.) TCC 2023, Part III. LNCS, vol. 14371, pp. 424–454. Springer, Cham (2023). https://doi.org/10.1007/978-3-031-48621-0_15
10. Avitabile, G., Botta, V., Friolo, D., Visconti, I.: Efficient proofs of knowledge for threshold relations. In: Atluri, V., Di Pietro, R., Jensen, C.D., Meng, W. (eds.) ESORICS 2022, Part III. LNCS, vol. 13556, pp. 42–62. Springer, Cham (2022). https://doi.org/10.1007/978-3-031-17143-7_3
11. Baum, C., Malozemoff, A.J., Rosen, M.B., Scholl, P.: Mac'n'Cheese: zero-knowledge proofs for boolean and arithmetic circuits with nested disjunctions. In: Malkin, T., Peikert, C. (eds.) CRYPTO 2021. LNCS, vol. 12828, pp. 92–122. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-84259-8_4
12. Ben-Sasson, E., Bentov, I., Horesh, Y., Riabzev, M.: Fast reed-solomon interactive oracle proofs of proximity. In: 45th International Colloquium on Automata, Languages, and Programming, ICALP 2018, July 9-13, 2018, Prague, Czech Republic. LIPIcs, vol. 107, pp. 14:1–14:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2018)
13. Benaloh, J., Leichter, J.: Generalized secret sharing and monotone functions. In: Goldwasser, S. (ed.) CRYPTO 1988. LNCS, vol. 403, pp. 27–35. Springer, New York (1990). https://doi.org/10.1007/0-387-34799-2_3

14. Bootle, J., Cerulli, A., Chaidos, P., Ghadafi, E., Groth, J., Petit, C.: Short accountable ring signatures based on DDH. In: Pernul, G., Ryan, P.Y.A., Weippl, E. (eds.) ESORICS 2015. LNCS, vol. 9326, pp. 243–265. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-24174-6_13
15. Bünz, B., Bootle, J., Boneh, D., Poelstra, A., Wuille, P., Maxwell, G.: Bulletproofs: short proofs for confidential transactions and more. In: 2018 IEEE Symposium on Security and Privacy, SP 2018, Proceedings, 21-23 May 2018, San Francisco, California, USA, pp. 315–334. IEEE Computer Society (2018)
16. Bünz, B., Fisch, B., Szepieniec, A.: Transparent SNARKs from DARK compilers. In: Canteaut, A., Ishai, Y. (eds.) EUROCRYPT 2020. LNCS, vol. 12105, pp. 677–706. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-45721-1_24
17. Ciampi, M., Persiano, G., Scafuro, A., Siniscalchi, L., Visconti, I.: Improved OR-composition of sigma-protocols. In: Kushilevitz, E., Malkin, T. (eds.) TCC 2016. LNCS, vol. 9563, pp. 112–141. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-49099-0_5
18. Ciampi, M., Persiano, G., Scafuro, A., Siniscalchi, L., Visconti, I.: Online/offline OR composition of sigma protocols. In: Fischlin, M., Coron, J.-S. (eds.) EUROCRYPT 2016. LNCS, vol. 9666, pp. 63–92. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-49896-5_3
19. Cramer, R.: Modular Design of Secure yet Practical Cryptographic Protocols. Ph.D. thesis, University of Amsterdam (1997)
20. Cramer, R., Damgård, I., MacKenzie, P.: Efficient zero-knowledge proofs of knowledge without intractability assumptions. In: Imai, H., Zheng, Y. (eds.) PKC 2000. LNCS, vol. 1751, pp. 354–372. Springer, Heidelberg (2000). https://doi.org/10.1007/978-3-540-46588-1_24
21. Cramer, R., Damgård, I., Maurer, U.: General secure multi-party computation from any linear secret-sharing scheme. In: Preneel, B. (ed.) EUROCRYPT 2000. LNCS, vol. 1807, pp. 316–334. Springer, Heidelberg (2000). https://doi.org/10.1007/3-540-45539-6_22
22. Cramer, R., Damgård, I., Schoenmakers, B.: Proofs 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). https://doi.org/10.1007/3-540-48658-5_19
23. Don, J., Fehr, S., Majenz, C., Schaffner, C.: Efficient nizks and signatures from commit-and-open protocols in the QROM. In: CRYPTO 2022, Part II. LNCS, vol. 13508, pp. 729–757. Springer, Cham (2022). https://doi.org/10.1007/978-3-031-15979-4_25
24. Don, J., Fehr, S., Majenz, C., Schaffner, C.: Online-extractability in the quantum random-oracle model. In: Dunkelman, O., Dziembowski, S. (eds.) EUROCRYPT 2022, Part III. LNCS, vol. 13277, pp. 677–706. Springer, Cham (2022). https://doi.org/10.1007/978-3-031-07082-2_24
25. Feng, H., Liu, J., Wu, Q., Li, Y.-N.: Traceable ring signatures with post-quantum security. In: Jarecki, S. (ed.) CT-RSA 2020. LNCS, vol. 12006, pp. 442–468. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-40186-3_19
26. Fischlin, M., Harasser, P., Janson, C.: Signatures from sequential-OR proofs. In: Canteaut, A., Ishai, Y. (eds.) EUROCRYPT 2020. LNCS, vol. 12107, pp. 212–244. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-45727-3_8
27. Fouque, P., Georgescu, A., Qian, C., Roux-Langlois, A., Wen, W.: A generic transform from multi-round interactive proof to NIZK. In: Boldyreva, A., Kolesnikov, V. (eds.) PKC 2023, Part II. LNCS, vol. 13941, pp. 461–481. Springer, Cham (2023). https://doi.org/10.1007/978-3-031-31371-4_16

28. Goel, A., Green, M., Hall-Andersen, M., Kaptchuk, G.: Stacking sigmas: a framework to compose Σ-protocols for disjunctions. IACR Cryptol. ePrint Arch., p. 422 (2021)
29. Goel, A., Green, M., Hall-Andersen, M., Kaptchuk, G.: Stacking sigmas: A framework to compose Σ-protocols for disjunctions. In: Dunkelman, O., Dziembowski, S. (eds.) EUROCRYPT 2022, Part II. LNCS, vol. 13276, pp. 458–487. Springer, Cham (2022). https://doi.org/10.1007/978-3-031-07085-3_16
30. Goel, A., Hall-Andersen, M., Kaptchuk, G., Spooner, N.: Speed-stacking: fast sublinear zero-knowledge proofs for disjunctions. In: EUROCRYPT 2023, Part II. LNCS, vol. 14005, pp. 347–378. Springer, Cham (2023). https://doi.org/10.1007/978-3-031-30617-4_12
31. Grigni, M., Sipser, M.: Monotone complexity. In: Proceedings of the London Mathematical Society Symposium on Boolean Function Complexity. pp. 57–75. Cambridge University Press, USA (1992)
32. Groth, J., Kohlweiss, M.: One-out-of-many proofs: or how to leak a secret and spend a coin. In: Oswald, E., Fischlin, M. (eds.) EUROCRYPT 2015. LNCS, vol. 9057, pp. 253–280. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-46803-6_9
33. Heath, D., Kolesnikov, V.: Stacked garbling for disjunctive zero-knowledge proofs. In: Canteaut, A., Ishai, Y. (eds.) EUROCRYPT 2020. LNCS, vol. 12107, pp. 569–598. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-45727-3_19
34. Ishai, Y., Kushilevitz, E., Ostrovsky, R., Sahai, A.: Zero-knowledge from secure multiparty computation. In: Proceedings of the 39th Annual ACM Symposium on Theory of Computing, San Diego, California, USA, June 11-13, 2007, pp. 21–30. ACM (2007)
35. Kattis, A.A., Panarin, K., Vlasov, A.: Redshift: transparent snarks from list polynomial commitments. In: Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security, CCS 2022, Los Angeles, CA, USA, November 7–11, 2022, pp. 1725–1737. ACM (2022)
36. Katz, J., Kolesnikov, V., Wang, X.: Improved non-interactive zero knowledge with applications to post-quantum signatures. In: Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security, CCS 2018, Toronto, ON, Canada, 15–19 October, 2018, pp. 525–537. ACM (2018)
37. Kim, A., Liang, X., Pandey, O.: A new approach to efficient non-malleable zero-knowledge. In: Chung, Y., Yung, M. (eds.) CRYPTO 2022, Part IV. LNCS, vol. 13510, pp. 389–418. Springer, Cham (2022). https://doi.org/10.1007/978-3-642-17955-6_3
38. Lindell, Y.: An efficient transform from sigma protocols to NIZK with a CRS and non-programmable random oracle. In: Dodis, Y., Nielsen, J.B. (eds.) TCC 2015. LNCS, vol. 9014, pp. 93–109. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-46494-6_5
39. Shamir, A.: How to share a secret. Commun. ACM **22**(11), 612–613 (1979)
40. Valiant, L.: Short monotone formulae for the majority function. J. Algorithms **5**(3), 363–366 (1984)
41. Wikström, D.: Special soundness revisited. Cryptology ePrint Archive, Paper 2018/1157 (2018). https://eprint.iacr.org/2018/1157

42. Wikström, D.: Special soundness in the random oracle model. Cryptology ePrint Archive, Paper 2021/1265 (2021). https://eprint.iacr.org/2021/1265
43. Zeng, G., Lai, J., Huang, Z., Wang, Y., Zheng, Z.: Dag-Σ: A dag-based sigma protocol for relations in CNF. In: Agrawal, S., Lin, D. (eds.) ASIACRYPT 2022, Part II. LNCS, vol. 13792, pp. 340–370. Springer, Cham (2022). https://doi.org/10.1007/978-3-031-22966-4_12