

Acyclicity Programming for Sigma-Protocols

Masayuki Abe¹, Miguel Ambrona¹, Andrej Bogdanov², Miyako Ohkubo³, and Alon Rosen⁴

¹NTT Secure Platform Laboratories

{abe.masayuki.cp, miguel.ambrona.fu}@hco.ntt.co.jp

²Chinese University of Hong Kong, andrejb@cse.cuhk.edu.hk

³Security Fundamentals Laboratory, CSR, NICT, m.ohkubo@nict.go.jp

⁴Herzliya Interdisciplinary Center, alon.rosen@idc.ac.il

Abstract. Cramer, Damgård, and Schoenmakers (CDS) built a proof system to demonstrate the possession of subsets of witnesses for a given collection of statements that belong to a prescribed access structure \mathcal{P} by composing so-called sigma-protocols for each atomic statement. Their verifier complexity is linear in the size of the monotone span program representation of \mathcal{P} .

We propose an alternative method for combining sigma-protocols into a single non-interactive system for a compound statement in the random oracle model. In contrast to CDS, our verifier complexity is linear in the size of the *acyclicity program* representation of \mathcal{P} , a complete model of monotone computation introduced in this work. We show that the acyclicity program size of a predicate is never larger than its de Morgan formula size and it is polynomially incomparable to its monotone span program size. We additionally present an extension of our proof system, with verifier complexity linear in the *monotone circuit size* of \mathcal{P} , in the common reference string model.

Finally, considering the types of statement that naturally reduce to acyclicity programming, we discuss several applications of our new methods to protecting privacy in cryptocurrency and social networks.

Keywords: sigma-protocols, zero-knowledge proofs, random oracles

1 Introduction

The feasibility of proving every NP-statement in zero-knowledge was first established by Goldreich, Micali, and Wigderson [GMW86]. For decades, such generic protocols have been perceived as mere proof-of-concept results, in part because they involve costly NP reductions and also require multiple independent repetitions for the sake of soundness error reduction.

More recently, we have witnessed a widespread push towards the efficient realization of general-purpose zero-knowledge proofs, with varying degrees of practicality [Gro16, AHIV17, BCG⁺17, GKM⁺18, BBB⁺18, BCR⁺19, CFQ19, MBKM19, LMR19]. Despite significant progress, the efforts are still ongoing,

with performance a far cry from traditional cryptographic constructs such as public-key cryptography.

In contrast, special purpose protocols for specific hard-on-average languages within NP have seen wide practical deployment, most notably in the form of standardized digital signatures and anonymous credentials. An abstraction that is at the heart of this approach is so-called *sigma-protocols* [Cra97]. A sigma-protocol is a three-move interactive protocol between a prover and a verifier where the prover sends an initial message, a , to the verifier who replies with a random challenge, e , and the prover then responds with an answer, z , based on which the verifier accepts or rejects.

Sigma-protocols often enjoy low soundness error by design, resulting in high efficiency relative to their generic counterparts. On the other hand, their zero-knowledge property is typically only guaranteed against honest verifiers. Using the Fiat-Shamir (FS) heuristic [FS87], they can be made non-interactive and secure against malicious verifiers. This conversion generates the challenge e by applying a hash function to the initial message a .

1.1 Sigma-Protocol Composition

Importantly, sigma-protocols can be used to prove compound NP relations, obtained by composing several sigma-protocols for “atomic” statements. Cramer, Damgård, and Schoenmakers (CDS) showed that sigma-protocols could be generically composed [CDS94]. Namely, given a collection of sigma-protocols for NP statements x_1, \dots, x_n , one can obtain a zero-knowledge proof of knowledge for any subset of statements that belong to a prescribed access structure \mathcal{P} .

The idea in CDS composition is to secret-share the challenge e according to the access structure \mathcal{P} and then use the shares as challenges in the corresponding sigma-protocols for each of the atomic statements. It works for any \mathcal{P} recognized by a monotone span program [CDM00]. It results in sigma-protocols that can be made non-interactive, using the FS heuristic, and proved secure in the programmable random oracle model [BR93]. Due to its simplicity and generality, the CDS paradigm is popular both in theory and in practice.

Another composition technique, which is more closely related to this work, was introduced in [AOS02] and recently revisited in [FHJ20]. The idea, inspired by the ring signatures of [RST01], is to sequentially generate a challenge e_i to a statement x_i by hashing the first message a_i from the roof for the statement x_{i-1} . This sequence ends when the first challenge e_1 is generated from a_n . In [FHJ20], the method is referred to as a *sequential composition* and turned into a signature scheme whose unforgeability is proven in the non-programmable random oracle model (NPROM) [Nie02] (see also [Lin15, CPSV16]). It is in contrast to CDS+FS, whose signature counterpart is shown to be not provable in the NPROM in a black-box manner.

1.2 Our contributions

We propose a new method for combining sigma-protocols for statements x_1, \dots, x_n into a non-interactive zero-knowledge proof of knowledge of witnesses w_i for subsets of the statements that belong to a prescribed access structure \mathcal{P} (Section 3). The complexity of the proof system is dictated by the *acyclicity program* (ACP) size of \mathcal{P} , a new model of monotone computation introduced in this work.

In contrast, CDS composition complexity depends on the monotone span program size of \mathcal{P} . Relying on monotone complexity theory results, we show the existence of function families whose monotone span program size is inherently superpolynomial in their acyclicity program size (Section 4). Acyclicity programs can be polynomially simulated by monotone circuits, so there is also a superpolynomial gap in the other direction [BGW99]. Thus the two composition methods are of incomparable complexity.

We prove the knowledge soundness of our acyclicity programming proof system in the programmable random oracle model (Theorem 1). We also obtain regular soundness in the non-programmable random oracle model (Theorem 2).

In Section 5 we describe a variant of the proof system whose complexity is determined by the monotone circuit size of \mathcal{P} . Monotone circuit size is at most polynomial in acyclicity program size but maybe superpolynomially smaller. The knowledge soundness of this protocol requires that the prover sample pairs of instances, out of which he can certify knowledge of exactly one. This is possible assuming the hardness of discrete logarithms.

Composition via acyclicity programming. Our starting point is the sequential composition method of [AOS02, FHJ20]. They represent a disjunction of statements $x_1 \vee \dots \vee x_n$ by a cycle on the nodes $1, \dots, n$ with the interpretation that the cycle is “satisfied” if the cycle can be traversed starting at any node labeled by a true statement with a known witness. An acyclicity program extends these semantics to general directed graphs that may contain multiple and overlapping cycles. The nodes of an acyclicity program are labeled by statements, allowing repetitions. A satisfying set of witnesses corresponds to a set of nodes that cover every cycle.

Composition for monotone formulas. In Section 4.2 we show how to convert any monotone formula describing an access structure into an acyclicity program of the same size. This reduction from monotone formulas enjoys concrete efficiency and thus may be of practical interest. For relatively simple representations such as those in disjunctive normal form over distinct instances, compound proofs using our ACP method tend to be shorter than those by CDS+FS as challenges are not necessarily included in the proofs.

Composition for monotone circuits. In Section 5 we address a more general composition, represented by monotone circuits consisting of AND and THRESHOLD gates. This method compiles sigma-protocols into non-interactive systems with a common reference string (CRS), assuming the added property that the CRS

allows oblivious sampling of instances from a hard language having a sigma-protocol. We note that oblivious sampling is often possible without CRS in the random oracle model, e.g. [GPS08], which is used in our basic ACP composition.

Unlike their non-monotone counterparts, monotone circuits are not universal for efficient computation of monotone functions [Raz85, AB87, Tar88]. Non-monotone circuits on n inputs, however, can be turned into monotone circuits on their $2n$ literals (variables and their negations) while at most doubling their size. Therefore, assuming the additional availability of sigma-protocols for proving *ignorance* [BM90, DK18] of the base statements, our composition method is essentially universal. Succinct proofs of ignorance are in particular known for discrete logarithms under a CRS and are also used in our circuit protocol construction. (The CRS we consider is simply a group element with an unknown discrete logarithm.)

Other advantages over related works. CDS composition essentially relies on the fact that the composed sigma-protocols are 2-special sound, i.e., two colliding transcripts allow to extract a witness with probability one efficiently. Our construction can be used to compose more general κ -special sound protocols since the challenges to every run of the underlying protocol can be directly re-assigned in the programmable random oracle model, and we can rewind the execution until κ valid transcripts are obtained. This allows the composition of a broader class of sigma-protocols, including, for instance, Stern’s protocol [Ste96], which is 3-special sound and is used in lattice/code-based constructions [LLNW16, LLNW17, NTWZ19, FLWL20]. Other examples are protocols for proofs of binary secrets [BCC⁺15] which are also 3-special sound and useful for proving the correctness of inputs to a circuit.

A recent work [AAB⁺20] enhances the CDS composition paradigm in the random oracle model, yielding shorter proofs and weakening the soundness requirements of the constituent sigma-protocols. Like CDS itself, their complexity is governed by the access structure’s span program size and is therefore of incomparable complexity to the results in this work.


1.3 Technical overview

Let $(\mathcal{C}_1, \mathcal{Z}_1, \mathcal{V}_1)$ and $(\mathcal{C}_2, \mathcal{Z}_2, \mathcal{V}_2)$ be sigma-protocols for relations R_1 and R_2 respectively. In the following, we show how to build a non-interactive zero-knowledge argument system (in the random oracle model) for the disjunctive combination of the above relations, i.e., for the language:

$$L_{R_1} \vee L_{R_2} := \{(x_1, x_2) \mid \exists w : (x_1, w) \in R_1 \vee (x_2, w) \in R_2\} .$$

Let H be a random oracle that maps arbitrarily large inputs into $\{0, 1\}^\lambda$ (w.l.o.g., the challenge space of both sigma-protocols). A proof on instance (x_1, x_2) will consist of two accepting transcripts (a_1, e_1, z_1) and (a_2, e_2, z_2) for the respective sigma-protocols, where the challenges are computed as:

$$e_1 = H(a_2) \quad \text{and} \quad e_2 = H(a_1)$$



and where $H(a)$ is an abbreviation of $H(x_1 \| x_2 \| i \| a)$.¹ Note that an honest prover in possession of a valid witness, say w_1 such that $(x_1, w_1) \in R_1$, can produce both transcripts by selecting $a_1 \leftarrow \mathcal{C}_1(x_1, w_1; r_1)$, computing $e_2 = H(a_1)$, running the simulator of the second sigma-protocol, obtaining $(a_2, z_2) \leftarrow \mathcal{S}_2(x_2, e_2)$ and using w_1 to complete the transcript for the first protocol on challenge $e_1 = H(a_2)$, producing an accepting z_1 . The soundness of the protocol can be proven based on the soundness of the underlying sigma-protocols in the random oracle model, since the structure of our scheme and the fact that H is a random oracle guarantee that either e_1 will be uniformly chosen after having fixed a_1 or that e_2 will be uniformly chosen after having fixed a_2 . Finally, zero-knowledge can be argued by defining a simulator with the ability to program the oracle H .

Observe that this technique can be generalized to combine in disjunction more than two statements, e.g., for combining three sigma-protocols we can define challenges as $e_1 = H(a_3)$, $e_2 = H(a_1)$ and $e_3 = H(a_2)$; but it can also be used to prove more complex compound statements. For example, given a monotone formula f in conjunctive normal form (CNF), consider a directed graph (where nodes are labelled as statements) that satisfies the following property: for every disjunctive clause in f there exists a cycle in the graph consisting of nodes associated to all the statements in the clause, and there are not other cycles that correspond to a clause not implied by the original formula. We can build a proof system for f where proofs consist of an accepting sigma-protocol

¹ To highlight the role of a_1 and a_2 , we use this shorthand in this section. We however remind readers the importance of including instances x_1 and x_2 as pointed out in [BPW12, BDG20].

transcript for every node in the graph, and where for every node i , the challenge e_i is computed based on all nodes j_1, \dots, j_k that have an edge pointing to i , that is, $e_i = H(a_{j_1} \parallel \dots \parallel a_{j_k})$. See Figure 1 for two examples. Graph 1a corresponds to the already mentioned disjunctive predicate of three statements. On the other hand, graph 1b corresponds to the formula $(1 \vee 2) \wedge (1 \vee 3) \wedge (1 \vee 4) \wedge (2 \vee 3 \vee 4)$; there exist a cycle in the graph between all nodes that appear in the same clause.

It turns out that every graph induces a proof system of a certain monotone compound predicate over the base statements. We refer to Section 3 for more details about how to characterize the predicate associated with a graph.

1.4 Potential Applications

An immediate application of our novel composition method is a ring signature scheme [RST01] where an ACP access structure can describe the admissible group of signers. It extends the result of [BSS02] where admissible signers are described by a monotone formula, since ACP can simulate monotone formulas (Section 4.2). Unforgeability of the signature scheme could be proven in the programmable random oracle model following standard techniques.

As we have already mentioned, the complexity of our new model of computation is incomparable to the complexity of monotone span programs. There exist specific predicates that support superpolynomially more compact ACP representations. We describe some scenarios where the structure of ACPs could yield substantial improvements in communication and computation.

Our ACP proof system is especially suitable for predicates related to *disconnecting a graph*. One can prove knowledge of witnesses associated to certain nodes in an *acyclic* graph that represent a *cut* from a node s to a node t (by artificially adding an edge between t and s and arguing the absence of cycles in the new graph). Such a proof would hide which subset of nodes has been used among all *cuts* that disconnect s from t . Even when the original graph is not acyclic, in Proposition 2 we show that the separation of s from t can be enforced with ACP at a cubic cost in the size of the graph, thereby retaining superpolynomial advantage over monotone span programs.

While the focus of our work is theoretical, we envision potential applications relating to certifying complex statements about social or financial activities. For example, consider a cryptocurrency system with a *transaction graph* whose nodes represent public keys, and whose (directed) edges represent money flows. Our proof system could apply to the following types of claims:

- *Proof of possession of white money.* A user may want to certify that the money associated with a given node (logically, owned by them) has been transferred (possibly not directly) from a set of whitelisted organizations (such as banks or well-reputed companies). To preserve the organizations' pseudonymity (not disclosing the nodes in the network controlled by them) the organizations can certify the user's integrity by proving knowledge of valid credentials associated with nodes that form a cut from the cryptocurrency

genesis node to the node of interest, without revealing any information about the actual cut.

- *Proof of self-transaction.* A common practice in pseudonymous networks to increase anonymity is to transfer assets to freshly generated identities. Repeating such self-transactions makes the graph complex, and it is believed to be good for privacy. The mechanism described above would allow a user to certify a statement of the form “*this node belongs to me, and all of its money was once owned by myself or by the subset of parties who are colluding with me to create this proof*”, while retaining privacy on the origin nodes.

Besides scenarios where the graph represents transactions, we expect that our proof system may find more applications in other settings like social networks. Other possible applications based on the expressivity of our new approach, i.e., arguing knowledge of witnesses that make a graph acyclic, are:

- *Resolving circular software dependencies.* A user may be willing to pay money to a company for resolving a circular dependency issue in their updated software. In the paradigm of *fair exchange*, the user wants to pay after the problem is solved, but the company wants the money upfront. They can leverage the approach devised by Maxwell based on contingent payment [Max11]. This approach requires the company first proving in zero-knowledge that they can resolve the circular dependency, a statement that could be easily expressed with an ACP program.
- *st-unreachability in ℓ or fewer steps.* As a generalization of the *directed st-cut* predicate, our method can be used to prove unreachability between two nodes in a directed graph by considering paths of limited length ℓ . This could be useful on applications where the graph represents a map or, as above, in graphs of transactions where the transactions’ lifetime matters. This application would require reducing the original graph G with n vertices into a graph H with $n\ell$ vertices so that G has a path from s to t of length at most ℓ if and only if H has a cycle. Such reduction could be devised as in the proof of Proposition 2 if distances were tracked as a separate parameter.

2 Preliminaries

For a finite set S , we write $a \stackrel{\$}{\leftarrow} S$ to denote that a is uniformly sampled from S . We denote the security parameter by $\lambda \in \mathbb{N}$. Given two functions $f, g : \mathbb{N} \rightarrow [0, 1]$, 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 $\{m, m+1, \dots, n\}$. We denote by $[n]$ the range $[1, n]$. By \mathbb{N}^* we denote the space of arbitrarily-long sequences of numbers in \mathbb{N} . We use multiplicative notation for groups. 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 $a \leftarrow \mathcal{A}(x)$; $e \xleftarrow{\$} \{0, 1\}^\lambda$; $z \leftarrow \mathcal{A}(e)$ 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). Let $R : \mathcal{X} \times \mathcal{W} \rightarrow \{0, 1\}$ be a binary relation defined over a set of instances \mathcal{X} and a set of witnesses \mathcal{W} . We denote by L_R be the language defined as $L_R := \{x \in \mathcal{X} \mid \exists w \in \mathcal{W} : R(x, w) = 1\}$.

2.1 Sigma-protocols

A sigma-protocol, introduced in [Cra97], is a public-coin interactive proof system that consists of only three data transfers between the prover and the verifier, and it satisfies a weaker notion of zero-knowledge and special soundness. In this work we employ a generalized definition of special soundness as in [AAB⁺20], where $\kappa \geq 2$ colliding transcripts with distinct (uniformly sampled) challenge lead to efficient witness extraction with probability $1 - \epsilon$ for a negligible ϵ . Compared to the original special soundness, where $\kappa = 2$ and $\epsilon = 0$, it captures a broader class of schemes, e.g. the parallel version of Stern's protocol, where an exponential number (but still negligible compared to the size of the challenge space) of colliding transcripts can be prepared without knowing the witness.

Definition 1 (Sigma-protocol). *A sigma-protocol for R is a triple of PPT algorithms $(\mathcal{C}, \mathcal{Z}, \mathcal{V})$ associated to the following execution:*

- i) $a \leftarrow \mathcal{C}(x, w; r)$: given $(x, w) \in \mathcal{X} \times \mathcal{W}$, outputs an initial message.*
- ii) $e \xleftarrow{\$} \{0, 1\}^\lambda$: the verifier's challenge is uniformly sampled.*
- iii) $z \leftarrow \mathcal{Z}(x, w, r, e)$: the prover answers to the challenge.*
- iv) $1/0 \leftarrow \mathcal{V}(x, a, e, z)$: the verifier returns 1 (acceptance) or 0 (rejection).*

A sigma-protocol must satisfy completeness, special soundness, and special honest verifier zero-knowledge defined as follows.

Perfect completeness. *For every $\lambda \geq 1$ and every pair $(x, w) \in R$,*

$$\Pr \left[e \xleftarrow{\$} \{0, 1\}^\lambda; a \leftarrow \mathcal{C}(x, w; r) : \mathcal{V}(x, a, e, z) = 1 \right] = 1.$$

κ -special soundness. *There exists a predicate $\Phi : (\{0, 1\}^\lambda)^\kappa \rightarrow \{0, 1\}$ and there exists a deterministic polynomial-time extractor \mathcal{E} such that:*

- On input $(x, a, \{e_1, z_1\}, \dots, \{e_\kappa, z_\kappa\})$ such that $\mathcal{V}(x, a, e_i, z_i) = 1$ for all $i \in [\kappa]$ and $\Phi(e_1, \dots, e_\kappa) = 1$, \mathcal{E} outputs a witness w s.t. $R(x, w) = 1$.*
- $\Pr \left[e_1, \dots, e_\kappa \xleftarrow{\$} \{0, 1\}^\lambda : \Phi(e_1, \dots, e_\kappa) = 0 \right]$ is negligible in λ . (This probability is referred to as the special-soundness error.)*

When Φ only checks for the challenges being different, the protocol is said to be perfectly κ -special sound.

Special honest verifier zero-knowledge. *There exists a probabilistic polynomial-time algorithm \mathcal{S} that, for every stateful PPT distinguisher \mathcal{D} ,*

$$\begin{aligned} & \Pr [(x, w, e) \leftarrow \mathcal{D}(1^\lambda); a \leftarrow \mathcal{C}(x, w; r); z \leftarrow \mathcal{Z}(x, w, r, e) : \mathcal{D}(a, z) = 1] \\ \approx & \Pr [(x, w, e) \leftarrow \mathcal{D}(1^\lambda); (a, z) \leftarrow \mathcal{S}(x, e) : \mathcal{D}(a, z) = 1] \end{aligned}$$

where \mathcal{D} must output values such that $(x, w) \in R$ and $e \in \{0, 1\}^\lambda$.

The following standard soundness is implied by κ -special soundness, for any constant κ .

Definition 2 (Soundness). *A sigma-protocol for L_R is said to be sound with error ϵ_{snd} if, for any PPT stateful adversary \mathcal{A} and for any $x \notin L_R$, it holds $\Pr[a \leftarrow \mathcal{A}(x); e \xleftarrow{\$} \{0, 1\}^\lambda; z \leftarrow \mathcal{A}(e) : \mathcal{V}(x, a, e, z) = 1] < \epsilon_{\text{snd}}$. We say the protocol is sound if ϵ_{snd} is negligible in λ .*

If there existed an adversary \mathcal{A} against the soundness game, winning with non-negligible probability, by the Forking Lemma (Lemma 1), there would exist another algorithm \mathcal{B} producing κ accepting transcripts with the same first message with non-negligible probability (there would be a probability loss with κ in the exponent, but κ is constant). Because the challenges in those transcripts would be sampled uniformly and independently, as the soundness game allows, $\Phi(e_1, \dots, e_\kappa)$ would be 1 with overwhelming probability. On the other hand, \mathcal{E} cannot extract a valid witness for x , since such a witness does not exist. This contradicts κ -special soundness.

For an example of our generalized version of special soundness, think of the parallel version of Stern's protocol, with challenge space $\{0, 1, 2\}^\lambda$. It achieves 3-special soundness under our definition, for condition $\Phi(e, e', e'')$ defined as $\exists i \in [\lambda]$ such that $\{e_i, e'_i, e''_i\} = \{0, 1, 2\}$, where the equality is between (unordered) sets. In this case, the special soundness error would be $(7/9)^\lambda$.

2.2 Non-Interactive Arguments

We define non-interactive argument systems in the random oracle model.

Definition 3 (Non-Interactive Argument System). *A non-interactive argument system for relation R in the random oracle model is a pair of polynomial-time oracle algorithms (Prove, Verify) that, for random oracle H :*

- $\text{Prove}^H(x, w) \rightarrow \pi$ is a probabilistic algorithm that takes an instance x and a witness w as input, and outputs a proof π .
- $\text{Verify}^H(x, \pi) \rightarrow 0/1$ is a deterministic algorithm that takes x and π , and outputs either 1 or 0 representing acceptance or rejection, respectively.

It is correct if, for every sufficiently large $\lambda \in \mathbb{N}$, all H , all $(x, w) \in R$, $\text{Verify}^H(x, \text{Prove}^H(x, w))$ outputs 1 except with negligible probability in λ .

Definition 4 (Zero-Knowledge). A non-interactive argument system (Prove, Verify) for R is zero-knowledge in the random oracle model if there exists a probabilistic polynomial-time algorithm **Simulator** that for every probabilistic polynomial-time distinguisher \mathcal{D} , the following difference is negligible in λ :

$$\Pr[1 \leftarrow \mathcal{D}^{H, \mathcal{P}}(1^\lambda)] - \Pr[1 \leftarrow \mathcal{D}^{\mathcal{S}^H}(1^\lambda)] ,$$

where H is a random oracle that returns an independently and uniformly chosen value in an appropriate domain for every distinct input. \mathcal{P} is a prover oracle that, on input $(x, w) \in R$, outputs $\text{Prove}^H(x, w)$. \mathcal{S} is a simulator oracle that takes two separate types of inputs: one is any string forwarded to **Prove** directly, and the other is $(x, w) \in R$ that only x is given to **Simulator**, and returns the output.

Definition 5 (Witness Indistinguishability). A non-interactive argument system (Prove, Verify) for R is witness indistinguishable if for every probabilistic polynomial-time adversary \mathcal{A} , any polynomial-size string h , any x, w_1, w_2 with $R(x, w_1) = R(x, w_2) = 1$, and any random oracle H , the following difference is negligible in λ :

$$\Pr[1 \leftarrow \mathcal{A}(x, w_1, \text{Prove}^H(x, w_1), h)] - \Pr[1 \leftarrow \mathcal{A}(x, w_1, \text{Prove}^H(x, w_2), h)] .$$

Definition 6 (Soundness). A non-interactive argument system (Prove, Verify) for L_R is sound if for any PPT oracle algorithm \mathcal{A} , any $x \notin L_R$, and a random oracle H , $\Pr[\pi \leftarrow \mathcal{A}^H(x) : 1 = \text{Verify}^H(x, \pi)]$ is negligible in λ . The probability is taken over the coins of \mathcal{A} and H .

The following definition models the fact that provers who create a valid proof must know a witness for the statement being proven. This is enforced by the existence of an extractor (that can invoke the prover in an oracle manner) that extracts a witness whenever the prover produces a valid proof.

Definition 7 (Knowledge Soundness). A non-interactive argument system (Prove, Verify) for R is knowledge sound with knowledge error ϵ in the random oracle model if there exists an expected polynomial-time algorithm \mathcal{E} such that for any probabilistic polynomial-time (potentially cheating) prover P^* , it holds that:

$$\Pr \left[\begin{array}{l} (x, \pi) \leftarrow P^{*H}(1^\lambda) \\ w \leftarrow \mathcal{E}^{H, P^*}(x, \pi) \end{array} : \text{Verify}^H(x, \pi) = 1 \wedge R(x, w) = 0 \right] \leq \epsilon(\lambda) ,$$

where the probability is taken over the coins of P^* , \mathcal{E} , and random oracle H . The system is said to be knowledge sound if ϵ is negligible in λ .

2.3 Graphs

A directed graph is a tuple $G = (V, E)$ where V is a set of elements called vertices and E is a set of ordered pairs, $E \subseteq V \times V$, called *directed edges* or *arrows*. Throughout the paper, we use N to denote the size of V . Given an edge $e = (u, v)$,

u is called the *head* of e , denoted by $head(e)$ and v is called the *tail* of e , denoted by $tail(e)$. We usually assume an ordering on the vertices V and denote by v_ℓ (or simply ℓ when it is clear from the context) the ℓ -th vertex, for every $\ell \in [N]$. For a vertex $v \in V$, we denote by $pred(v)$ the set of *predecessors* of v , that is, $pred(v) := \{u \in V \mid (u, v) \in E\}$. Given a subset of vertices $U \subset V$, we define the subgraph $G(U)$ as the graph (U, E_U) , where $E_U := \{(u, v) \in E \mid u \in U \wedge v \in U\}$. A *cycle* in G is a finite sequence of edges (e_1, \dots, e_ℓ) satisfying $tail(e_i) = head(e_{i+1})$ for every $i \in [\ell-1]$ and $tail(e_\ell) = head(e_1)$. A graph that contains no cycles is called *acyclic*.

3 ACP composition

3.1 Construction

In this section we describe and analyze our ACP composition for zero-knowledge composition of sigma-protocols. The common input to the prover and verifier consists of n statements x_1, \dots, x_n and a monotone set system \mathcal{P} over base set $[n]$ represented by a monotone acyclicity program, which we will define shortly. The prover wants to convince the verifier that he knows valid witnesses $w_i, \forall i \in S$ for some set S in \mathcal{P} . We assume the availability of special-sound, honest verifier zero-knowledge sigma-protocols $\Sigma_i = (\mathcal{C}_i, \mathcal{Z}_i, \mathcal{V}_i, \mathcal{S}_i)$ for verifying that w_i is a witness for x_i .

Definition 8 (Monotone acyclicity program). A monotone acyclicity program A is a directed graph G whose nodes are labeled by the variables $1, \dots, n$, allowing repetitions. We use $var(j) \in [n]$ to denote the label of node $j \in [N]$, and $pred(j)$ to denote the set of nodes pointing to j . A set $S \subseteq [n]$ is accepted by A if every (directed) cycle in G contains a node j such that $var(j) \in S$.

Equivalently, a set $S \subseteq [n]$ is accepted by A if the subgraph $G(\bar{S})$ induced by the complement set of nodes $\{j \mid var(j) \notin S\}$ contains no cycles. The sets accepted by A form a monotone set system $\Pi(A)$. We will measure the size of A by the number of nodes N and, as a complexity parameter of secondary interest, the number of edges M .

We now describe the ACP verifier (Figure 2). A proof π consists of N parts $(a_1, e_1, z_1), \dots, (a_N, e_N, z_N)$, where (a_j, e_j, z_j) has the format of a transcript for $\Sigma_{var(j)}$. The verifier accepts if all transcripts are accepting. The key idea is to enforce a specific choice of challenges e_j that allows the prover to pass verification only when knowing a sufficiently large set of witnesses S for which the graph $G(\bar{S})$ is acyclic. To accomplish this, we require that e_j be the value of the random oracle H evaluated at all a_p for $p \in pred(j)$ (ordered canonically).

As an example, consider the acyclicity program in Figure 1a, where the challenges e_j must satisfy the constraints $e_1 = H(\mathbf{x}||1||a_3)$, $e_2 = H(\mathbf{x}||2||a_1)$ and $e_3 = H(\mathbf{x}||3||a_2)$. A prover that knows a witness, say w_1 , can pass verification by first choosing $a_1 \leftarrow \mathcal{C}_1(x_1, w_1; r_1)$, then simulating (a_2, z_2) given e_2 , then simulating (a_3, z_3) given e_3 , and finally completing the proof of knowledge

Inputs: $\mathbf{x} = x_1, \dots, x_n$, proof π , and acyclicity program A (with N nodes).
 Random oracle: H .

1. Parse π as $((a_1, e_1, z_1), \dots, (a_N, e_N, z_N))$.
2. Verify that $e_j = H(\mathbf{x} \| j \| \{a_p \mid p \in \text{pred}_A(j)\})$ for all $j \in [N]$.
3. Accept iff $\mathcal{V}_i(x_i, a_j, e_j, z_j)$ accepts for all $j \in [N]$, where $i = \text{var}(j)$.

Fig. 2: The ACP Verifier.

Inputs: $\mathbf{x} = x_1, \dots, x_n$, witnesses $w_i : i \in S$, acyclicity program A (with N nodes).
 Precondition: A accepts S .
 Random oracle: H .

1. Compute $a_j \leftarrow \mathcal{C}_j(x_i, w_i; r_i)$ for all $j \in [N] : i = \text{var}(j) \in S$.
2. Compute a topological sort $j(1), \dots, j(t_{\max})$ of the vertices in $G(\bar{S})$.
3. For t ranging from 1 to t_{\max} , do:
 - (a) Set $e_{j(t)} := H(\mathbf{x} \| j(t) \| \{a_p \mid p \in \text{pred}(j(t))\})$.
 - (b) Compute $(a_{j(t)}, z_{j(t)}) \leftarrow \mathcal{S}_{j(t)}(x_i, e_{j(t)})$, where $i = \text{var}(j(t))$.
4. For all j such that $\text{var}(j) \in S$, do:
 - (a) Set $e_j := H(\mathbf{x} \| j \| \{a_p \mid p \in \text{pred}(j)\})$.
 - (b) Compute $z_j \leftarrow \mathcal{Z}_j(x_i, w_i, r_j, e_j)$, for all $j \in [N] : i = \text{var}(j) \in S$.
5. Output $\pi := ((a_1, e_1, z_1), \dots, (a_N, e_N, z_N))$.

Fig. 3: The ACP Prover.

$\mathcal{Z}_1(x_1, w_1, r_1, e_1)$ for w_1 upon challenge e_1 . This strategy can be generalized to an arbitrary acyclicity program A , as long as the complement of the set of known witnesses is acyclic, or equivalently the set of known witnesses is accepting for A . The general prover is given in Figure 3.

The communication complexity of this protocol is the sum of the communication complexities of the sigma-protocols $\Sigma_{\text{var}(1)}, \dots, \Sigma_{\text{var}(N)}$. In fact, the challenges e_1, \dots, e_N do not have to be sent explicitly as the ACP verifier can turn step 2 from a verification into a calculation. The running time of the verifier is $O(Nt_{\mathcal{V}} + M)$ where $t_{\mathcal{V}}$ is the maximum running time of the verifiers \mathcal{V}_i , assuming H can be evaluated in linear time.

3.2 Security

The *composition* of witness relations R_i with the monotone set system \mathcal{P} is the witness relation $R_{\mathcal{P}}$ in which $(w_i : i \in S)$ is a witness for $\mathbf{x} = (x_1, \dots, x_n)$ if $S \in \mathcal{P}$ and w_i is an R_i -witness for x_i for all $i \in S$.

Theorem 1. *If $\Sigma_i = (\mathcal{C}_i, \mathcal{Z}_i, \mathcal{V}_i, \mathcal{S}_i)$ is a sigma-protocol for R_i for all $i \in [n]$ over the common challenge space $\{0, 1\}^\lambda$, \mathcal{P} is the set system computed by acyclicity program A , and H is a programmable random oracle, then ACP is a non-interactive zero-knowledge argument of knowledge for $R_{\mathcal{P}}$.*

We now prove completeness, zero-knowledge, and knowledge soundness of ACP.

Completeness. Assume that the prover knows witnesses $w_i : i \in S$ for some accepting set S of A . By definition, the graph $G(\bar{S})$ is acyclic and admits a topological sort.² Therefore all the inputs to H in step 3a are well-defined. By the HVZK property of Σ_i for $i = \text{var}(j) \notin S$, all the HVZK simulators \mathcal{S}_i run in step 3b (on random challenges) yield accepting transcripts. By the completeness of Σ_i for $i = \text{var}(j) \in S$, all the transcripts computed in steps 1 and 4 are also accepting. It follows that the ACP prover passes verification with overwhelming probability. The completeness error of ACP is at most the sum of the larger one among the completeness and simulation errors of the input protocols Σ_i .

Zero-knowledge. The zero-knowledge simulator runs independent copies of the simulators $\mathcal{S}_i(x_i, e_j)$, for all $j \in [N]$ where $i = \text{var}(j)$ and e_j is uniformly sampled, obtaining transcripts (a_j, e_j, z_j) . It then outputs their concatenation as the proof π . The random oracle H is programmed so that e_j equals $H(\mathbf{x} \| j \| \{a_p | p \in \text{pred}(j)\})$. The produced simulated proof is distributed as a genuine proof owing to the indistinguishability of the output by \mathcal{S}_i from regular transcripts of Σ_i . By a hybrid argument, the simulation error is at most the sum of the simulation errors of the input protocols Σ_i . Perfect zero-knowledge is achieved if every Σ_i is perfect special honest verifier zero-knowledge.

Knowledge soundness. Let P^* be a potentially cheating prover that produces statement $\mathbf{x} = (x_1, \dots, x_n)$ and a valid proof π with probability $\delta + \binom{q}{2} 2^\lambda$, for some $0 < \delta \leq 1$. P^* runs in time t , has query complexity q and is given oracle answers $\mathbf{h} = (h_1, \dots, h_q) \in (\{0, 1\}^\lambda)^q$ to its q queries. We will show that if all the Σ_i protocols are κ -special sound with knowledge extractor \mathcal{E}_i with (negligible) error $\epsilon_i \leq (\delta/2)^\kappa / 2q^{\kappa-1}$, then our ACP extractor, defined in Figure 4, produces a valid witness for $R_{\mathcal{P}}$ with probability at least $\delta/2$, in expected time $O(\kappa q n t / \delta)$.

We will assume that (1) all prover queries to the random oracle H are distinct, and (2) if $(\mathbf{x}, \pi) \leftarrow P^*(\mathbf{h})$ for $\pi = ((a_i, e_i, z_i) : 1 \leq i \leq N)$ then P^* made all queries of the form $(\mathbf{x} \| j \| \{a_p : p \in \text{pred}(j)\})$. Assumption (1) and (2) are without loss of generality as these constraints can be enforced while preserving prover efficiency. Assumption (2) ensures that all oracle queries of the ACP verifier must have also been made by P^* . We say that the output (\mathbf{x}, π) of $P^*(\mathbf{h})$ is *valid* if the ACP verifier accepts (\mathbf{x}, π) when interacting with any oracle that answers the k -th query of P^* by h_k , for all $k \in [q]$.³

² A topological sort is a linear ordering of the vertices of the graph satisfying that for every edge (u, v) , u comes before v in the ordering.

³ This is a slight abuse of terminology for the sake of readability as validity is not determined by \mathbf{x} and π only, but may also depend on \mathbf{h} .

Input: Acyclicity program A with graph G .
Oracle access to: P^* .

1. Set $S = \emptyset$.
2. While $G(\mathcal{S})$ has a directed cycle C , repeat the following:
 3. Keep sampling random \mathbf{h}^1 until $P^*(\mathbf{h}^1) = (\mathbf{x}^1, \pi^1)$ is valid.
 4. Let $j^* = \text{last}(\mathbf{h}^1; C, \mathbf{x}^1)$ and $i = \text{var}(j^*)$.
 5. Remember $x_i = \mathbf{x}_{j^*}^1$ and $(a_{j^*}^1, e_{j^*}^1, z_{j^*}^1) = (j^*\text{-th component of } \pi^1)$.
 6. Set $r = 2$ and repeat the following until $r = \kappa$ but at most $4\kappa q/\delta$ times:
 7. Sample \mathbf{h}^r at random conditioned on $h_k^r = h_k^1$ for $1 \leq k < \text{det}(\mathbf{h}^1; C, \mathbf{x}^1)$.
 8. If $P^*(\mathbf{h}^r) = (\mathbf{x}^r, \pi^r)$ is valid and $\text{det}(\mathbf{h}^r; C, \mathbf{x}^r) = \text{det}(\mathbf{h}^1; C, \mathbf{x}^1)$,
 9. Remember $(e_{j^*}^r, z_{j^*}^r) = (\text{last 2 terms in } j^*\text{-th component of } \pi^r)$.
 10. Increment r .
 11. If $r = \kappa$ and $\mathcal{E}_i(x_i, a_{j^*}^1, \{e_{j^*}^1, z_{j^*}^1\}, \dots, \{e_{j^*}^r, z_{j^*}^r\})$ outputs a valid witness w_i ,
 12. Remember x_i, w_i and add i to S .
13. Output the statement-witness pairs $(x_i, w_i : i \in S)$.

Fig. 4: The ACP Extractor.

We will also assume that (3) all oracle answers are distinct. Assumption (3) incurs a loss of $\binom{q}{2}2^{-\lambda}$ in the probability of $P^*(\mathbf{h})$ being valid (where \mathbf{h} is now a sequence q uniformly sampled distinct values from $\{0, 1\}^\lambda$), so from here on we will assume that success probability of P^* is at least δ .

The ACP extractor greedily collects witnesses from every cycle C of A . To do so, it seeks to find a query index Q such that the Q -th query is the one that determines the commitments and challenges along the cycle C . By forking at this query Q , which is prefixed by $\mathbf{x} \parallel j$ for some node j , a witness for $i = \text{var}(j)$ can be extracted. We define “the query that determines commitments and challenges along C ” as follows. Assume $(\mathbf{x}, \pi) \leftarrow P^*(\mathbf{h})$ is a valid statement-proof pair, all challenges e_j are represented in the oracle answers \mathbf{h} , and all entries h_k in \mathbf{h} are distinct (thus, every challenge e_j is uniquely represented as $h_{k(j)}$ in \mathbf{h}). Let $\text{det}(\mathbf{h}; C, \mathbf{x})$ be the smallest k_C such that $k(j) \leq k_C$ for all nodes j in cycle C , and $\text{last}(\mathbf{h}; C, \mathbf{x})$ be the node j_C so that the k_C -th query is prefixed by $\mathbf{x} \parallel j_C$. We argue that if $(\mathbf{x}, \pi) \leftarrow P^*(\mathbf{h})$ is valid then the first $k_C = \text{det}(\mathbf{h}; C, \mathbf{x})$ entries of \mathbf{h} determine the commitment a_{j_C} in π in the following sense:

Claim 1 *Assume P^* is deterministic. If \mathbf{h} and \mathbf{h}' are two sequences, each of which has pairwise distinct entries, $(\mathbf{x}, \pi) := P^*(\mathbf{h})$ and $(\mathbf{x}', \pi') := P^*(\mathbf{h}')$ are both valid, $k_C = \text{det}(\mathbf{h}; C, \mathbf{x}) = \text{det}(\mathbf{h}'; C, \mathbf{x}')$, and $h_k = h'_k$ for all $1 \leq k < k_C$, then $\mathbf{x} = \mathbf{x}'$ and $a_{j_C} = a'_{j_C}$.⁴*

⁴ For $\pi = ((a_1, e_1, z_1), \dots, (a_N, e_N, z_N))$ and $\pi' = ((a'_1, e'_1, z'_1), \dots, (a'_N, e'_N, z'_N))$.

Proof. As $P^*(\mathbf{h})$ is valid, for every $j \in C$ there is a unique query step $k(j)$ such that the query $(\mathbf{x} \| j \| \{a_p : p \in \text{pred}(j)\})$ must have been made by $P^*(\mathbf{h})$ at step $k(j)$ and answered by $h_{k(j)} = e_j$. The same occurs analogously for $P^*(\mathbf{h}')$, denote such query step by $k'(j)$ for all $j \in C$, in this case.

Since the k -th answers h_k and h'_k are identical for all $1 \leq k < k_C$ and P^* is deterministic, the k -th queries of $P^*(\mathbf{h})$ and $P^*(\mathbf{h}')$ must be identical for all $1 \leq k \leq k_C$, and so $\text{last}(\mathbf{h}; C, \mathbf{x}) = \text{last}(\mathbf{h}'; C, \mathbf{x}') =: j_C$. By definition of det , we have $k(j), k'(j) < k_C$ for all $j \in C \setminus \{j_C\}$, so $k(j) = k'(j)$ and therefore $e_j = h_{k(j)} = h'_{k'(j)} = e'_j$ for all $j \in C \setminus \{j_C\}$.

In particular, $e_{j^+} = e'_{j^+}$ for the successor j^+ of j_C in C . Consequently, the corresponding query $k(j^+) = k'(j^+)$ made by P^* is then of the form $(\mathbf{x} \| j^+ \| \{a_p | p \in \text{pred}(j^+)\}) = (\mathbf{x}' \| j^+ \| \{a'_p | p \in \text{pred}(j^+)\})$. It follows that $\mathbf{x} = \mathbf{x}'$ and, as $j_C \in \text{pred}(j^+)$, also $a_{j_C} = a'_{j_C}$. \square

Claim 2 *Assume $P^*(\mathbf{h})$ is valid with probability δ (over the internal coins of P^* and the sampling of \mathbf{h} : q uniformly chosen distinct values from $\{0, 1\}^\lambda$) and the special soundness error of all Σ_i is at most $(\delta/2)^\kappa / 4q^{\kappa-1}$. The ACP extractor outputs a witness for $R_{\mathcal{P}}$ in expected time $O(\kappa q n t / \delta)$ with probability at least $\delta/2$ over the randomness of P^* .*

Proof. The ACP extractor interacts in a black-box manner with a P^* whose internal randomness τ has been fixed (after being sampled from the appropriate distribution). We first argue that, with reasonable probability, fixing P^* 's internal randomness results in a “decent” prover algorithm. More concretely, let X be the random variable that maps τ to $\Pr[P^* \text{ succeeds} | \tau]$ we have that for every $\alpha \in [0, 1]$, $\Pr[X \geq \alpha] \geq \delta - \alpha$, since

$$\begin{aligned} \delta &= \Pr[P^* \text{ succeeds}] \\ &= \Pr[P^* \text{ succeeds} | X < \alpha] \Pr[X < \alpha] + \Pr[P^* \text{ succeeds} | X \geq \alpha] \Pr[X \geq \alpha] \\ &\leq \alpha + \Pr[X \geq \alpha] \ . \end{aligned}$$

In particular, for $\alpha = \delta/2$, we have $\Pr[X \geq \delta/2] \geq \delta/2$. Consequently, with probability $\delta/2$, fixing P^* 's internal randomness will result in a deterministic prover with success probability (over the choice of \mathbf{h}) of at least $\delta/2$. From now on, we assume P^* has its randomness fixed.

Fix an arbitrary directed cycle C in G . Let Valid_r be the event that $P^*(\mathbf{h}^r)$ is valid and Match_r be the event $\text{det}(\mathbf{h}^r; C, \mathbf{x}^r) = \text{det}(\mathbf{h}^1; C, \mathbf{x}^1)$. By assumption, $\Pr[\text{Valid}_1] \geq \delta/2$. Applying the Forking Lemma 1 below (with $\kappa = 2$, $F(\mathbf{h}) = \text{det}(\mathbf{h}; C, \mathbf{x}^1)$, and $A = \text{“}P^*(\mathbf{h}) \text{ is valid”}$), we get

$$\Pr[\text{Valid}_r \cap \text{Match}_r | \text{Valid}_1] \geq \frac{\delta}{2q} \ .$$

(By Claim 1, $\text{Valid}_r \cap \text{Valid}_1 \cap \text{Match}_r$ implies that $\mathbf{x}^r = \mathbf{x}^1$ and $a_{j_C}^r = a_{j_C}^1$.) The event that the ACP extractor reaches step 9 in the r -th iteration of loop 6 (given that it completed step 3) is precisely $\text{Valid}_r \cap \text{Match}_r$ (conditioned on

$Valid_1$). If T_r is the number of executions of loop 6 with a given value of r , then $T_r \mid Valid_1$ is a geometric random variable with success parameter at least $\delta/2q$ so $E[T_r \mid Valid_1] \leq q/\delta$ and, by linearity of expectation,

$$E[T_1 + \dots + T_{\kappa-1} \mid Valid_1] = E[T_1 \mid Valid_1] + \dots + E[T_{\kappa-1} \mid Valid_1] \leq \frac{2\kappa q}{\delta} .$$

Let $Reach$ be the event that r reaches value κ in step 11. By Markov's inequality,

$$\Pr[Reach] = \Pr[T_1 + \dots + T_{\kappa-1} \leq 4\kappa q/\delta \mid Valid_1] \geq \frac{1}{2} .$$

We now upper bound the probability of the event $Fail$ that \mathcal{E}_i fails to extract a witness w_i for x_i from these transcripts, given $Reach$. The event $Reach$ is the same as $Valid_1 \cap \dots \cap Valid_r \cap Fork$, where $Fork = Match_1 \cap \dots \cap Match_r$ is the event “ $det(\mathbf{h}^r; C, \mathbf{x}^1) = det(\mathbf{h}^1; C, \mathbf{x}^1)$ for all r and $\mathbf{h}^1, \dots, \mathbf{h}^\kappa$ match in the first $det(\mathbf{h}^1; C, \mathbf{x}^1)$ positions”. Therefore:

$$\Pr[Fail \mid Reach] \leq \Pr[\neg A \mid Reach] \leq \frac{\Pr[\neg A \mid Fork]}{\Pr[Valid_1 \cap \dots \cap Valid_r \mid Fork]} ,$$

where A is the set of admissible transcripts for \mathcal{E}_i . Conditioned on $Fork$, by Claim 1, for $i = var(j_C)$, the i -th component x_i^r of the input \mathbf{x}^r is equal to x_i^1 and the commitment $a_{j_C}^r$ is equal to $a_{j_C}^1$, for all r . The challenges $e_{j_C}^r$ are however independent of one another (and distribute uniformly) because the sequences \mathbf{h}^r fork in position $k_C = det(\mathbf{h}^1; C, \mathbf{x}^1)$. So the probability of $\neg A$ given $Fork$ is exactly the κ -special soundness error ϵ_i of \mathcal{E}_i : $\Pr[\neg A \mid Fork] = \epsilon_i$. On the other hand, by the Forking Lemma 1,

$$\Pr[Valid_1 \cap \dots \cap Valid_r \mid Fork] \geq \frac{(\delta/2)^\kappa}{q^{\kappa-1}} .$$

By our assumption on ϵ_i , we conclude that $\Pr[Fail \mid Reach] \leq 1/2$.

In conclusion, in any given cycle C there exists a node j_C such that loop 2 succeeds in extracting a solution $(x_i, w_i) \in R_i$ for $i = var(j_C)$ with probability at least

$$\Pr[Reach \cap \neg Fail] = \Pr[Reach](1 - \Pr[Fail \mid Reach]) \geq \frac{1}{2} \cdot \frac{1}{2} \geq \frac{1}{4} .$$

It follows that the expected number of iterations to extract a solution to $R_{\mathcal{P}}$, namely a solution to $R_{var(j_C)}$ for some j_C in every cycle, is at most $4n$. As each loop 2 iteration has expected running time $O(\kappa qt/\delta)$, the total expected running time of the extractor is $O(\kappa qnt/\delta)$ as desired. \square

By Claim 2, the ACP extractor succeeds with probability at least $1/2\delta$ over the combined randomness of the extractor and P^* . To extract a witness with probability $1 - \delta$ we run the ACP extractor independently $O((1/\delta) \log(1/\delta))$ times (with different fixed randomness for P^*), thereby establishing knowledge soundness of ACP.

It remains to establish the following version of the Forking Lemma, first introduced by Pointcheval and Stern [PS00] and later generalized by Bellare and Neven [BN06], here restated and proved in our notation.

Lemma 1 (Forking Lemma). *Let F be a possibly randomized function from q -term sequences to the set $\{1, \dots, q\}$ and A be a subset of q -term sequences. The probability that $\mathbf{h}^1, \dots, \mathbf{h}^\kappa$ are all in A and $F(\mathbf{h}^1) = \dots = F(\mathbf{h}^\kappa)$, conditioned on $\mathbf{h}^1, \dots, \mathbf{h}^\kappa$ being identical in the first $F(\mathbf{h}^1)$ entries and independent otherwise, is at least $\delta(A)^\kappa / q^{\kappa-1}$, where $\delta(A)$ is the measure of A .*

Proof. Say $\mathbf{h}^1, \dots, \mathbf{h}^\kappa$ fork at f if they are independent conditioned on their first $f - 1$ entries being identical.

$$\begin{aligned}
& \Pr[\mathbf{h}^1, \dots, \mathbf{h}^\kappa \in A \text{ and } F(\mathbf{h}^1) = \dots = F(\mathbf{h}^\kappa) \mid \mathbf{h}^1, \dots, \mathbf{h}^\kappa \text{ fork at } F(\mathbf{h}^1)] \\
&= \sum_{f=1}^q \Pr[\mathbf{h}^r \in A \text{ and } F(\mathbf{h}^r) = f \text{ for all } r \leq \kappa \mid \mathbf{h}^1, \dots, \mathbf{h}^\kappa \text{ fork at } f] \\
&= \sum_{f=1}^q \mathbb{E}[\Pr[\mathbf{h} \in A \text{ and } F(\mathbf{h}) = f \mid h_1, \dots, h_{f-1}]^\kappa] \\
&\geq \sum_{f=1}^q \Pr[\mathbf{h} \in A \text{ and } F(\mathbf{h}) = f]^\kappa, \\
&\geq q \cdot \left(\sum_{f=1}^q \frac{1}{q} \cdot \Pr[\mathbf{h} \in A \text{ and } F(\mathbf{h}) = f] \right)^\kappa \\
&= \frac{\Pr[\mathbf{h} \in A]^\kappa}{q^{\kappa-1}},
\end{aligned}$$

where the last two inequalities are both applications of Jensen's inequality. \square

3.3 Security in the non-programmable random oracle model

Sigma-protocols with the Fiat-Shamir transform provide soundness when the hash function in the transformation is modeled as a non-programmable random oracle. Here we show that an analogue statement holds for ACP as well. Roughly, in the NPROM, the same random oracle is given to every algorithm including *the reduction* built in the security proof. Thus, the reduction does not have the advantage of being able to program the input-output correspondence of the random oracle. For a more formal definition we refer to [Nie02].

Theorem 2. *If $\Sigma_i = (C_i, \mathcal{Z}_i, \mathcal{V}_i, \mathcal{S}_i)$ is a sigma-protocol for L_{R_i} for all $i \in [n]$ over the common challenge space $\{0, 1\}^\lambda$, \mathcal{P} is the set system computed by acyclicity program A , then ACP is a non-interactive witness indistinguishable argument for $L_{R_{\mathcal{P}}}$ in the non-programmable random oracle model.*

Proof. Witness indistinguishability is shown by a hybrid argument. Given two distinct sets of witnesses, \mathbf{w}_1 and \mathbf{w}_2 for \mathbf{x} , we build a hybrid starting from a proof obtained by running the ACP prover for input $(\mathbf{x}, \mathbf{w}_1, A)$. Adding a witness from $\mathbf{w}_2 \setminus \mathbf{w}_1$ one by one, it reaches to a proof for $(\mathbf{x}, \mathbf{w}_1 \cup \mathbf{w}_2, A)$ as input to the prover. Then, removing a witness belonging to $\mathbf{w}_1 \setminus \mathbf{w}_2$ one by one, it reaches to a proof for $(\mathbf{x}, \mathbf{w}_2, A)$. Any noticeable change of probability of a distinguisher outputting 1 at any point of the hybrid reduces to violating the special honest verifier zero-knowledge property of the corresponding Σ_i . Suppose that the gap happens between hybrids with respect to witness set \mathbf{x}^* and $\mathbf{x}^* \cup \{w_{i^*}\}$. The reduction executes the prover algorithm with \mathbf{x}^* using the zero-knowledge challenger that, given e_{i^*} , returns (a_{i^*}, z_{i^*}) , instead of zero-knowledge simulator in step 3(b) in

Figure 3. Thus the proof distributes in the same way as prover with \mathbf{x}^* if (a_{i^*}, z_{i^*}) is a simulated one while it is the same as prover with $\mathbf{x}^* \cup \{w_{i^*}\}$ if (a_{i^*}, z_{i^*}) is created using w_{i^*} . Accordingly, the witness indistinguishability error is at most $2N\epsilon_{zk}$ where ϵ_{zk} is the largest zero-knowledge error among Σ_i .

For soundness, suppose that there is an adversary \mathcal{A} that outputs, with probability $\epsilon_{\mathcal{A}}$, an instance $\mathbf{x} := (x_1, \dots, x_n)$ together with a valid proof on \mathbf{x} , $\pi := ((a_1, e_1, z_1), \dots, (a_N, e_N, z_N))$ that satisfies:

- for every qualified set S , there exists i that $x_i \notin L_{R_i}$,
- $e_j = H(\mathbf{x} \| j \| \{a_p \mid p \in \text{pred}(j)\})$ for all $j \in [N]$, and
- $\mathcal{V}_i(x_i, a_j, e_j, z_j)$ accepts for all $j \in [N]$, where $i = \text{var}(j)$.

Let $e_j \leftarrow H_j$ denote the event that $H(\mathbf{x} \| j \| \{a_p \mid p \in \text{pred}(j)\})$ is evaluated and e_j is returned. Let $H_k \leftarrow a_j$ denote the event that $\mathbf{x} \| k \| \{a_p \mid p \in \text{pred}(k)\}$ satisfying $a_j \in \text{pred}(k)$ is queried to H . Over the sequence of q distinct queries from \mathcal{A} to the random oracle, we define function τ that, given an event E , outputs $\ell \in [q]$ such that the ℓ -th query to the random oracle made event E happened in the view of \mathcal{A} . (Define the output of τ as 0 if E did not occur at any query).

Suppose that there exist $j, k \in [N]$ that either $e_j \leftarrow H_j$ or $H_k \leftarrow a_j$ do not happen in the view of \mathcal{A} . Then the probability that the randomly assigned hash output (drawn in verification for the first time) satisfies the corresponding relation with respect to the already fixed $x_{\text{var}(j)}$, a_j and z_j , is bound by the soundness error of Σ_j . Denoting the maximum soundness error by ϵ_{snd} and taking union bound for all q possible queries, the probability for this event is upper-bounded by $q \cdot \epsilon_{\text{snd}}$.

Now, consider the case where for all $j, k \in [N]$, both events $e_j \leftarrow H_j$ and $H_k \leftarrow a_j$ happen in the view of \mathcal{A} . For every cycle C in graph G , let (j^*, k^*) be the arrow in C that satisfies $\tau(H_{k^*} \leftarrow a_{j^*}) = \min(\{\tau(H_k \leftarrow a_j) \mid (j, k) \in C\})$. Observe that $\tau(H_k \leftarrow a_{j^*}) = \tau(e_{j^*} \leftarrow H_{j^*})$ for some $k \neq k^*$ and therefore, $\tau(H_{k^*} \leftarrow a_{j^*}) < \tau(e_{j^*} \leftarrow H_{j^*})$. (Equivalently, a_{j^*} is fixed before e_{j^*} in the view of the adversary.) Let J be the set of such index j^* over all cycles in G . Then, by definition, $S := \{\text{var}(j) \mid j \in J\}$ is a qualified set. To see this, suppose that there exists $i \in S$ with $x_i \notin L_{R_i}$ and let j be such that $i = \text{var}(j) (\in J)$. For fixed x_i and a_j , the probability that, e_j is sampled uniformly at random (by H_j) and the adversary outputs z_j passing the verification is bounded by the soundness error ϵ_{snd} of Σ_i . Consequently, the probability that \mathcal{A} produces a valid proof where there exists $i \in S$ that $x_i \notin L_{R_i}$ is upper-bound by $q \cdot \epsilon_{\text{snd}}$.

Accumulating the above bounds, we have $\epsilon_{\mathcal{A}} < 2q \cdot \epsilon_{\text{snd}}$, which concludes the proof of soundness in NPRM. \square

4 The expressive power of acyclicity programs

The communication and verification efficiencies of ACP composition grow linearly in the monotone acyclicity program size of the set system \mathcal{P} . To this end, in this section we study the complexity of acyclicity programs. In Section 4.1 we show

that monotone acyclicity program size is polynomially equivalent to monotone branching program size [GS92].

As a corollary, leveraging known results from monotone complexity theory in Section 4.1 we conclude that there exist families whose monotone span program size grows superpolynomially in their monotone acyclicity program size. As the efficiency of the Cramer, Damgård, Schoenmakers sigma-protocol composition [CDS94] (CDS composition) is dictated by the former parameter, we obtain concrete examples of set systems for which ACP composition is asymptotically more efficient than CDS composition.

In Section 4.2 we demonstrate a simulation of de Morgan formulas by acyclicity programs of the same size. This is a consequence of known simulations of formulas by branching programs.

In this section, it will be useful to view (monotone) set systems \mathcal{P} over $[n]$ as (monotone) functions $\mathcal{P}: \{0, 1\}^n \rightarrow \{0, 1\}$ via the usual identification of a set by its indicator vector. Recall that the *monotone dual* \mathcal{P}^\dagger of \mathcal{P} is the function $\neg\mathcal{P}(\neg x_1, \dots, \neg x_n)$. An acyclicity program for \mathcal{P} can then be interpreted as a “cyclicity program” for \mathcal{P}^\dagger :

Definition 9 (Monotone cyclicity program). *A monotone cyclicity program is a directed graph whose nodes are labeled by variables x_1, \dots, x_n . The program accepts a given input if the subgraph induced by the true-valued nodes has a cycle.*

The size of a cyclicity program is the number of nodes. We may, without affecting the size, allow for nodes labeled by the constant `true`. (These nodes can be bypassed by contracting all their incoming and outgoing edge pairs.)

4.1 Polynomial equivalence with branching programs

A *monotone branching program* is also a directed acyclic graph with distinguished start and accept nodes whose vertices are labeled by variables or the constant 1. The program accepts a given input if the subgraph induced by the 1-valued vertices has a path from the start state to the accept state.

Let $mCP(\mathcal{P})$, $mBP(\mathcal{P})$ be the size of the smallest monotone cyclicity program, smallest monotone branching program for \mathcal{P} . The efficiency of the ACP verifier for the composed relation $R_{\mathcal{P}}$ in Section 1 is proportional to $mCP(\mathcal{P})$.

The following proposition shows that the complexity of our proof system is polynomially related to the monotone branching program size of the monotone dual \mathcal{P}^\dagger of the composition predicate. As monotone branching programs can polynomially simulate monotone formulas [GS92], the complexity is also upper-bounded by the formula size $mF(\mathcal{P}^\dagger)$ of \mathcal{P}^\dagger . We elaborate on this simulation in the next section.

Proposition 1. *For every monotone \mathcal{P} , $mBP(\mathcal{P}) \leq mCP(\mathcal{P}) \leq mBP(\mathcal{P})^3 + 2$.*

On the other hand, the prover complexity in the NIZK composition framework of Cramer, Damgård, and Schoenmakers [CDS94] is proportional to the monotone span program complexity $mSP_{\mathbb{F}}(\mathcal{P}^\dagger)$ times some polylogarithmic factor in \mathbb{F} , over

any finite field \mathbb{F} .⁵ Let stC denote the directed st -connectivity family of functions, Pitassi and Robere [PR18] showed that:

$$mSP_{\mathbb{F}}(\text{stC}) = mBP(\text{stC})^{\Omega(mBP(\text{stC}))} ,$$

for every \mathbb{F} . The monotone dual stC^\dagger is a family of predicates on n inputs for which ACP composition has complexity linear in n , but CDS composition requires $n^{\Omega(\log n)}$ communication for any potential implementation of stC^\dagger by a monotone span program (over any finite field).

In the other direction, it is known that:

$$mSP_{\mathbb{F}}(\mathcal{P}) \leq mF(\mathcal{P}) \leq mBP(\mathcal{P})^{O(mBP(\mathcal{P}))} ,$$

for every \mathcal{P} [GS92] so the quasipolynomial separation is optimal.

Proof of Proposition 1. In the course of proving Proposition 1 we establish the equivalence of several natural directed graph problems with respect to suitable reductions. A *labeled graph* is a directed graph whose nodes are labeled by variables x_1, \dots, x_n or the constant 1. A *projection reduction* from graph property P to graph property Q is an algorithm that converts a labeled graph G into another labeled graph H with the same set of labels such that for any assignment to the variables, the subgraph induced by the 1-labels of G has property P if and only if the subgraph induced by the 1-labels of H has property Q . We consider the following graph properties:

- **CYCLICITY:** the graph has a (directed) cycle.
- **st -CONNECTIVITY:** the graph has a (directed) path from vertex s to vertex t .
- **ACYCLIC st -CONNECTIVITY:** st -connectivity under the promise that the graph is acyclic.

Proposition 2. CYCLICITY, st -CONNECTIVITY, and ACYCLIC st -CONNECTIVITY are polynomial-time equivalent with respect to projection reductions.

Proof. ACYCLIC st -CONNECTIVITY reduces to CYCLICITY by adding a back edge from t to s and trivially reduces to st -CONNECTIVITY.

In the other direction, we view CYCLICITY and st -CONNECTIVITY as special cases of the following property PATHS: Given a set of triplets (s_i, t_i, ℓ_i) does there exist a path from s_i to t_i of length ℓ_i for some i ? In CYCLICITY $s_i = t_i$ ranges over all vertices and ℓ_i ranges over all values from 1 to n . In st -CONNECTIVITY $s_i = s, t_i = t$, and ℓ_i ranges over all values from 0 to $n - 1$.

The reduction from PATHS to ACYCLIC st -CONNECTIVITY consists of building the reachability graph for the natural guess-and-verify nondeterministic logspace algorithm for PATHS. Let V be the vertex set of the PATHS instance G . For every pair $(v, w) \in V \times V$ and ℓ between 0 and n create a node (v, w, ℓ) in the ACYCLIC

⁵ Both mF and mSP [KW93] are in fact invariant under taking monotone duals, but mBP is not [GS92].

st -CONNECTIVITY instance H with the same label as w and two special nodes s and t with label 1. Also, for every edge (w, w') in G , every $v \in V \setminus \{w'\}$ and all $\ell \geq 1$, create an edge between node $(v, w, \ell - 1)$ node (v, w', ℓ) in H . Finally, connect s to $(s_i, s_i, 0)$ and connect (s_i, t_i, ℓ_i) to t for all i in H . By construction, H is acyclic and its st -paths are all of the form

$$s, (s_i, v_0, 0), (s_i, v_1, 1), \dots, (s_i, v_{\ell_i}, \ell_i), t ,$$

where v_0, \dots, v_{ℓ_i} is a path from s_i to t_i in G . □

As the reduction from ACYCLIC st -CONNECTIVITY to CYCLICITY preserves the number of vertices and the reverse reduction maps an n -vertex graph into an $n^3 + 2$ -vertex graph (after shortcutting the bottom-layer vertices) we obtain Proposition 1.

4.2 Acyclicity programs for monotone formulas

A monotone (*de Morgan*) formula for a n -variate function \mathcal{P} is a tree whose internal nodes are labeled by AND/OR and whose leaves are labeled by inputs, allowing repetitions. The size of a formula is the number of leaves. Monotone formulas naturally compute monotone functions. The monotone formula size remains invariant under duality by de Morgan's laws.

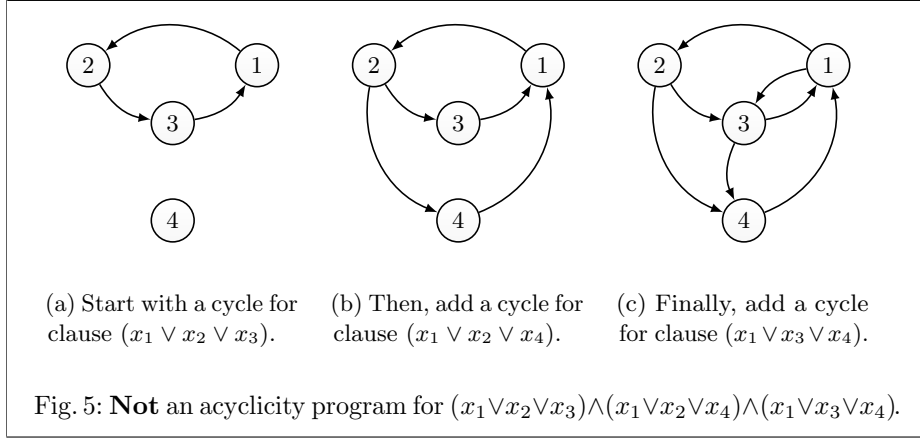
By a standard simulation of (monotone) formulas by (monotone) branching programs (e.g., [GS92]) it is known that every monotone formula can be simulated by a branching program of the same size, i.e., $mF(\mathcal{P}) \geq mBP(\mathcal{P})$. It follows that both $mF(\mathcal{P})$ and $mF(\mathcal{P}^\dagger)$ are lower-bounded by $mCP(\mathcal{P})$, i.e. a formula (and its dual) can be computed by an acyclicity program of the same size:

Proposition 3. *A formula for \mathcal{P} can be converted to an acyclicity program for \mathcal{P} of the same size in at most quadratic time in the size.*

Before we describe this simulation, we discuss an alternative seemingly intuitive approach that **does not** work. To be specific, let us look at the 4-variate CNF formula $(x_1 \vee x_2 \vee x_3) \wedge (x_1 \vee x_2 \vee x_4) \wedge (x_1 \vee x_3 \vee x_4)$. It may be tempting to “convert” this formula to the acyclicity program in Figure 5c, computed by greedily including a cycle for each of the clauses. The two are, however not equivalent, because the assignment $x_1 x_2 x_3 x_4 = \mathbf{f t f t}$ (\mathbf{t}, \mathbf{f} stand for *true, false*) is satisfying for the CNF but not for the acyclicity program. (Note that an unintended cycle has been created between nodes 1 and 3, unavoidably.)

To prove Proposition 3 we first convert the formula to a branching program, add a back edge from the accept state to the start state, and then short-circuit all states labeled by the constant 1. The branching program is constructed inductively. A variable x_i is represented by the branching program $\mathbf{s}(x_i) \rightarrow i \rightarrow \mathbf{a}(x_i)$. Their AND and their OR are represented by the branching programs:

$$\begin{array}{c} \mathbf{s} \rightarrow A \rightarrow \mathbf{t}(\mathbf{rue}) \rightarrow B \rightarrow \mathbf{a} \end{array} \quad \begin{array}{ccc} & \mathbf{s} & \longrightarrow & A \\ & \downarrow & & \downarrow \\ & B & \longrightarrow & \mathbf{a} \end{array}$$



In the AND construction, the accept state of A is fused with the start state of B into the **true** constant. To illustrate the transformation, consider the formula:

$$\mathcal{F} = x_1 \vee (x_{21} \wedge (x_{221} \vee x_{222} \vee x_{223}) \wedge x_{23}) \vee x_3 \vee x_4 .$$

To convert it into an acyclicity program, we work with its dual:

$$\mathcal{F}^\dagger = x_1 \wedge (x_{21} \vee (x_{221} \wedge x_{222} \wedge x_{223}) \vee x_{23}) \wedge x_3 \wedge x_4 .$$

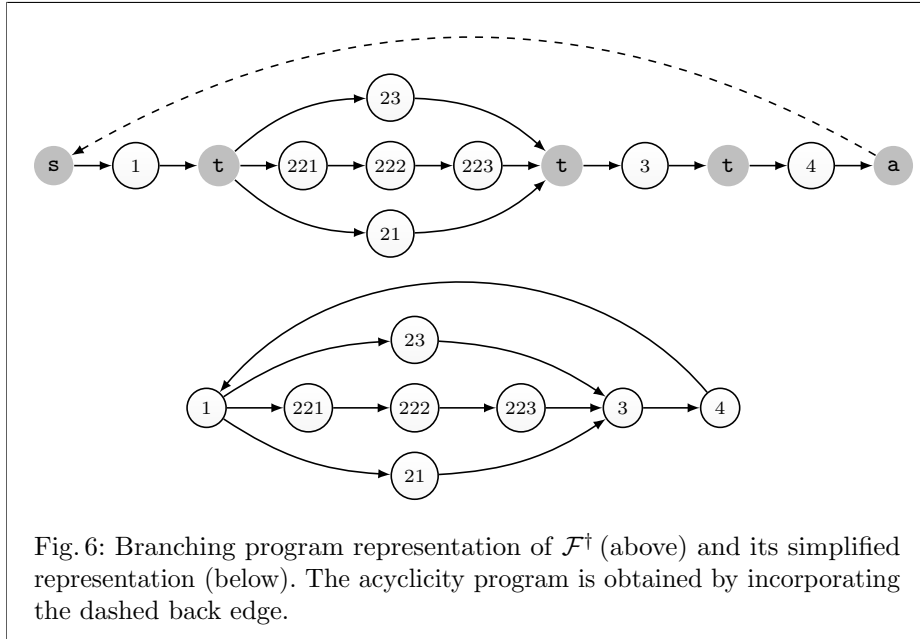
Figure 6 (above) shows the branching program representation of \mathcal{F}^\dagger including the constant states. The cyclicity program representation of \mathcal{F}^\dagger , which is the same as the acyclicity program representation of \mathcal{F} is then obtained by adding a back edge from the accept state to the start state. Figure 6 (below) shows the simplification that results from short-cutting the constant states.

We note that this reduction from monotone formulas makes it easy to turn the witness indistinguishability property into zero-knowledge in the common reference string model by following the techniques from [FLS90, CPSV16]: select a random instance x from a hard language as a common reference string and prove the extended statement “ $(x \in L) \vee$ (the original statement)”.

5 Composition for Predicates Represented by Circuits

In this section we present a sigma-protocol composition scheme with respect to predicates represented by monotone circuits. Unlike Theorem 1, this scheme relies on additional computational assumptions, specifically, we propose an instantiation based on the hardness of computing discrete logarithms.

A direct simulation of monotone circuits by acyclicity programs is inherently inefficient, since $mBP(\text{stC}^\dagger) = mBP(\text{stC})^{\Omega(mBP(\text{stC}))}$, as shown by Grigni and Sipser [GS92]. By Proposition 1 it follows that the acyclicity program size of stC



is superpolynomial in its branching program size and therefore in its monotone circuit size.

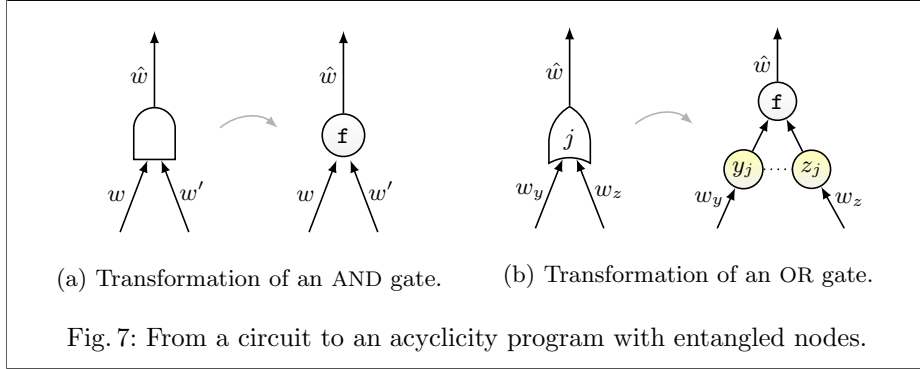
Instead, we emulate a circuit C on n inputs x_1, \dots, x_n by an acyclicity program A that, in addition to these n inputs as nodes, contains paired auxiliary nodes $(y_1, z_1), \dots, (y_m, z_m)$. Accepting inputs to C yield accepting inputs to A in which at most one of the two nodes in each pair (y_i, z_i) is set to true (and vice versa). Each auxiliary pair represents a pair of sigma-protocols for proving knowledge of at most one out of two auxiliary witnesses. Such protocols can be constructed and proved secure in the common random string model from the discrete logarithm assumption.

For simplicity, we apply the construction to monotone circuits in the AND/OR basis, but we note that the method can be extended to other threshold gates.

Definition 10 (Boolean circuit). *An AND/OR boolean circuit is a directed acyclic graph with n sources, and other nodes of in-degree 2, called internal gates. Sources are labeled by inputs x_1, \dots, x_n , all other nodes are labeled by AND/OR.*

The size of the circuit is the total number of nodes. We consider circuits with a single output node, which compute a monotone function from $\{0, 1\}^n$ to $\{0, 1\}$: after assigning the input nodes, the internal gates can be evaluated in any order consistent with the underlying graph, determining in a unique output value.

An *entangled assignment* to a pair of boolean variables (y, z) is an assignment in which at most one of the values is true.



Proposition 4. *Given a circuit C with n sources x_1, \dots, x_n and m OR nodes there exists an acyclicity program A with $n+2m$ nodes $x_1, \dots, x_n, y_1, z_1, \dots, y_m, z_m$ such that x_1, \dots, x_n satisfies C if and only if there exists an entangled assignment to all (y_j, z_j) so that $\mathbf{x}, \mathbf{y}, \mathbf{z}$ satisfies A . Moreover, A (and \mathbf{y}, \mathbf{z}) can be computed from C (and \mathbf{x}) in time at most $n + O(m^2)$.*

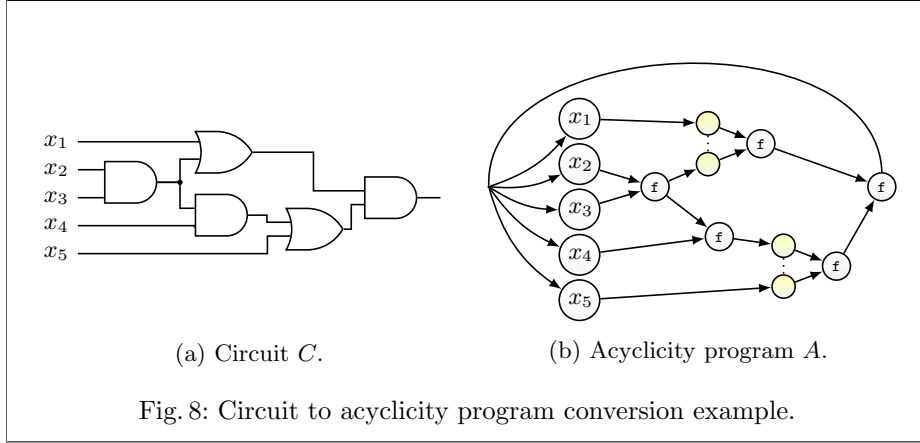
Proof. We replace each AND gate by a **false** node and the j -th OR gate by a gadget consisting of (y_j, z_j) and a **false** node connected as in Figure 7. We add back edges from the output node to all n input nodes. See Figure 8 for an example of this transformation.

If we ignore the back edges, we need to argue that satisfying assignments to C can be extended to cover all source-to-sink paths in A and vice versa. First suppose \mathbf{x} is a satisfying assignment for C . For each true OR gate j with input wires w_y, w_z , set y_j to true if w_y is false and set z_j to true if w_z is false. (The values to the false OR gates can be set arbitrarily.) We now argue that all source-sink paths in A are covered by induction on circuit size. The claim is vacuous for a circuit with one node. Since the assignment is satisfying, the top gate top evaluates to true. If top is AND, then by the inductive hypothesis the paths in both subcircuits are covered. If top is OR, then the true subcircuit paths are covered by the inductive hypothesis, while the false subcircuit paths (if any) are covered by our definition of y_j and z_j .

We prove the converse also by induction on circuit size. The claim is vacuous for a circuit with one node. Now assume \mathbf{x} is not satisfying so the top gate top evaluates to false. If top is AND, then at least one of the subcircuits is false and the paths in it cannot all be covered. If top is OR, then both of the subcircuits are false, so at least one of them cannot be covered regardless of the choice of y_j and z_j .

The false nodes in A can be short-circuited resulting in an acyclicity program of size $n + 2m$ as desired. \square

In order to instantiate the ACP compiler with the acyclicity program A from Proposition 4, we need to describe proof relations and sigma-protocols for each



entangled pair. Given two entangled inputs y and z , the prover ought to be able to prove knowledge of a witness for y or a witness for z but not both. The instances y and z should therefore be chosen in a correlated manner. In order to provide a general definition, we define the notion of *restrictive sampling*. We will then show how to build a restrictive sampling scheme from the discrete logarithm assumption.

Definition 11 (Restrictive sampling). *A restrictive sampling scheme is a triple of PPT algorithms $\mathcal{Y} = (\text{Setup}, \text{Gen}, \text{Verify})$:*

- $\text{Setup}(1^\lambda) \rightarrow pp$ is a probabilistic algorithm that outputs a set of public parameters, including an NP-relation R^* and a sigma-protocol Σ^* for it.
- $\text{Gen}(pp, b) \rightarrow (h_0, h_1, w)$ is a probabilistic algorithm that, on input pp and a bit b , outputs two instances h_0, h_1 and a witness w such that $R^*(h_b, w) = 1$.
- $\text{Verify}(pp, h_0, h_1) \rightarrow 0/1$ is a deterministic algorithm that, on input pp and two instances h_0, h_1 outputs either 1 (acceptance) or 0 (rejection).

We require a restrictive sampling scheme satisfy the following properties:

Partial knowledge. *For every PPT algorithm \mathcal{A} , the following probability is negligible in the security parameter λ :*

$$\Pr \left[\begin{array}{c} pp \leftarrow \text{Setup}(1^\lambda) \\ (h_0, h_1, w_0, w_1) \leftarrow \mathcal{A}(pp) \end{array} : \begin{array}{c} R^*(h_0, w_0) = R^*(h_1, w_1) = 1 \\ \text{Verify}(pp, h_0, h_1) = 1 \end{array} \right].$$

Witness independence. *For $b \in \{0, 1\}$, let \mathcal{D}_b be the distribution defined as:*

$$\mathcal{D}_b := (pp \leftarrow \text{Setup}(1^\lambda); (h_0, h_1, w) \leftarrow \text{Gen}(pp, b); \text{return } (pp, h_0, h_1)) .$$

We require distributions \mathcal{D}_0 and \mathcal{D}_1 be identical.

Theorem 3. For some $\lambda, n \in \mathbb{N}$, and all $i \in [n]$, let Σ_i be a sigma-protocol for relation R_i , with common challenge space $\{0, 1\}^\lambda$ and let \mathcal{Y} be a restrictive sampling scheme. Let C be a monotone AND/OR boolean circuit for the set system \mathcal{P} and let $H : \{0, 1\}^* \rightarrow \{0, 1\}^\lambda$ be a random oracle. The scheme described in Figure 9 is a NIZK proof of knowledge for $R_{\mathcal{P}}$.

Proof. *Completeness*, follows from the completeness of the ACP proof system and the “only if” part of Proposition 4.

The *zero-knowledge* simulator samples random pairs h_j^0, h_j^1 by running $\mathcal{Y}.\text{Gen}$ on input⁶ $(pp, 0)$ to simulate the first part of the view of the prover. It then runs the ACP simulator on inputs x_i, h_j^0, h_j^1 , inheriting its indistinguishability.

For *knowledge soundness*, given the *partial knowledge* property of \mathcal{Y} , we can assume that algorithms running in expected time t cannot find h_0, h_1, w_0, w_1 such that $R^*(h_0, w_0) = R^*(h_1, w_1) = 1$ and $\mathcal{Y}.\text{Verify}(pp, h_0, h_1) = 1$, where t is the running time of the ACP extractor \mathcal{E} . By soundness of the ACP extractor, if t is sufficiently large in terms of the cheating prover complexity, \mathcal{E} produces witnesses for some satisfying set S in A . Our assumption guarantees that this set cannot include a pair of entangled claims h_j^0, h_j^1 such that $\mathcal{Y}.\text{Verify}(pp, h_j^0, h_j^1) = 1$ and the corresponding witnesses w_j^0, w_j^1 satisfy $R(h_j^0, w_j^0) = R(h_j^1, w_j^1) = 1$. Consequently, the extraction of witnesses for S provided at most one witness for each pair of entangled claims, so by the knowledge soundness of \mathcal{E} and the “if” part of Proposition 4, the subset of S indexed by nodes x_1, \dots, x_n is satisfying for \mathcal{P} . \square

5.1 Restrictive sampling from the discrete logarithm assumption

Consider the following restrictive sampling scheme:

- **Setup**(1^λ): sample a λ -bits prime p and select a group $\mathbb{G} = \langle g \rangle$ of order p . Sample $\tau \xleftarrow{\$} \mathbb{Z}_p$ and let $h := g^\tau$. Define $R^*(h, x) := (h = g^x)$ and let Σ^* be the Schnorr sigma-protocol for $\text{PoK}\{(x) : R^*(h, x) = 1\}$. Return $pp := (p, \mathbb{G}, g, h)$.
- **Gen**(pp, b): sample $w \xleftarrow{\$} \mathbb{Z}_p$, set $h_b := g^w$, $h_{1-b} := hg^{-w}$. Return (h_0, h_1, w) .
- **Verify**(pp, h_0, h_1): output 1 if $h_0 h_1 = h$, output 0 otherwise.

It is not hard to see that the scheme satisfies the *partial knowledge* property, because an algorithm \mathcal{A} that, given (g, h) , outputs h_0, h_1, w_0, w_1 satisfying that $h_0 = g^{w_0}$, $h_1 = g^{w_1}$ and $h_0 h_1 = h$, can be used to compute discrete logarithms, since $w_0 + w_1$ is the discrete logarithm of h in base g .

Furthermore, the scheme satisfies the *witness independence* property, since for any $g, h \in \mathbb{G}$, the following two distributions are identical:

$$(w \xleftarrow{\$} \mathbb{Z}_p; \text{return}(g^w, hg^{-w})) \equiv (w \xleftarrow{\$} \mathbb{Z}_p; \text{return}(hg^{-w}, g^w)) .$$

⁶ The simulator may alternatively choose input $(pp, 1)$, in both cases the simulation is perfect due to the *witness independence* property of \mathcal{Y} .

Inputs: Claims x_1, \dots, x_n , circuit C represented by acyclicity program A with entangled pairs $(y_1^0, y_1^1), \dots, (y_m^0, y_m^1)$.

Common reference string: pp generated with $\mathcal{T}.\text{Setup}$, random oracle H .

Prover:

Input: Witnesses $w_i: i \in S \subseteq [n]$ and $b(j) \in \{0, 1\}$ for $j \in [m]$ such that the input set $\{x_i: i \in S\} \cup \{y_j^{b(j)}: j \in [m]\}$ is satisfying for A .

1. For each entangled pair, run $(h_j^0, h_j^1, w_j^*) \leftarrow \mathcal{T}.\text{Gen}(pp, b(j))$ and set values h_j^0, h_j^1 as statements for the claims associated to nodes $y_j^{b(j)}$ and $y_j^{1-b(j)}$, respectively.
2. Emulate the ACP prover for acyclicity program A on claims x_1, \dots, x_n with witnesses $w_i: i \in S \subseteq [n]$ and claims $h_1^0, h_1^1, \dots, h_m^0, h_m^1$ with witnesses w_j^* .

Verifier:

1. Upon receiving instances h_j^0, h_j^1 , verify that $\mathcal{T}.\text{Verify}(pp, h_j^0, h_j^1) = 1, \forall j \in [m]$.
2. Emulate the ACP verifier for acyclicity program A , claims x_i w.r.t. relation R_i , and claims h_j^0, h_j^1 w.r.t. relation R^* (defined in pp).

Fig. 9: Composition for Predicates Represented by Circuits.

5.2 Security in the non-programmable random oracle model

We can prove the security of our construction from Figure 9 in the NPRM, at the cost of having *soundness* (Definition 6) instead of *knowledge soundness*. For that, observe that the proof of Theorem 3 relies almost completely on the ACP composition security, which can be proven secure in the non-programmable random oracle model. Nevertheless, as shown in Section 3.3, in the NPRM the ACP composition achieves witness indistinguishability (and not ZK).

Our construction from Figure 9 can, however, achieve zero-knowledge if the restrictive sampling scheme is equipped with an extra algorithm that, on input pp and the randomness used to compute it (which can be seen as a trapdoor), can indeed “violate” the *partial knowledge* property. Namely, it can produce instances h_0 and h_1 and valid witnesses for both them, such that $\text{Verify}(pp, h_0, h_1) = 1$. The zero-knowledge simulator could leverage this algorithm to sample the instances associated to entangled pairs, knowing both witnesses of every pair, and then emulate the ACP prover for the acyclicity program on claims x_1, \dots, x_n (with not witnesses for them) and claims $h_1^0, h_1^1, \dots, h_m^0, h_m^1$ (with all witnesses for them).

Observe that this extra algorithm can be defined for (but it is not specific to) the instantiation of restrictive sampling based on discrete logarithm. In particular, knowing τ computed during **Setup**, one can generate $w \xleftarrow{\$} \mathbb{Z}_p$, compute $h_0 := g^w$, $h_1 := g^{\tau-w}$ and output $(h_0, h_1, w, \tau-w)$. Furthermore, observe that the above distribution on (h_0, h_1) is identical to \mathcal{D}_0 and \mathcal{D}_1 (from the *witness independence* property).

6 Concluding remarks

Constructing zero-knowledge proof systems by combining sigma-protocols in a compound statement is a powerful technique used for many applications, including anonymous credentials or ring signatures. The most famous and widely used of such composition techniques is the celebrated CDS composition [CDS94], which can be used for compound statements expressed as monotone span programs.

In this work, we have presented a novel technique for combining sigma-protocols into a single non-interactive system for a compound statement. Unlike CDS, our scheme loses the structure of sigma-protocol and it is proven secure in the random oracle model. However, our new methodology enhances the CDS composition in several flavors, including: *new expressivity* (the complexity of our system is linear in the size of the *acyclicity program* representation of the access structure, incomparable to monotone span programs), *more generality* (it is not limited to 2-special sound atomic protocols) and can often lead to *more compact* proofs (as in our circuit composition, where one single transcript is present per atomic statement). Consequently, our results arguably complement previous composition techniques.

Exploring whether our techniques can lead to more efficient zero-knowledge systems achieving post-quantum security is an appealing target for future work. Observe that the only part of the presented work that relies on classical assumptions is the instantiation of the restrictive sampling scheme. But even this primitive could be instantiated under post-quantum assumptions. For example, it could be instantiated under the Short Integer Solution (SIS) assumption [GPV08]. Note that the map $f_A(\mathbf{x}) = A\mathbf{x} \pmod{q}$ is almost surjective, so given \mathbf{y} , one could sample \mathbf{y}_0 and \mathbf{y}_1 such that $\mathbf{y}_0 + \mathbf{y}_1 = \mathbf{y}$, and know a short \mathbf{x}_b such that $A\mathbf{x}_b = \mathbf{y}_b$. For that, sample a short \mathbf{x}_b first and set $\mathbf{y}_b := A\mathbf{x}_b$, $\mathbf{y}_{1-b} := \mathbf{y} - A\mathbf{x}_b$. The *partial knowledge* property would hold because if, given \mathbf{y} , one could find short $\mathbf{x}_0, \mathbf{x}_1$ such that $A\mathbf{x}_0 + A\mathbf{x}_1 = \mathbf{y}$, then $\mathbf{x}_0 + \mathbf{x}_1$ would be a short preimage of \mathbf{y} under A . The *witness independence* is guaranteed by the fact that the distribution of $\mathbf{y} = A\mathbf{x}$ for a randomly chosen short \mathbf{x} is close to uniform.

References

- AAB⁺20. Masayuki Abe, Miguel Ambrona, Andrej Bogdanov, Miyako Ohkubo, and Alon Rosen. Non-interactive composition of sigma-protocols via share-then-hash. In Shiho Moriai and Huaxiong Wang, editors, *Advances in Cryptology - ASIACRYPT 2020 - 26th International Conference on the Theory and Application of Cryptology and Information Security, December 7-11, 2020, Proceedings*, volume 12493 of *Lecture Notes in Computer Science*, pages 749–773. Springer, 2020.
- AB87. Noga Alon and Ravi B. Boppana. The monotone circuit complexity of boolean functions. *Combinatorica*, 7(1):1–22, 1987.
- AHIV17. Scott Ames, Carmit Hazay, Yuval Ishai, and Muthuramakrishnan Venkatasubramanian. Liger: Lightweight sublinear arguments without a trusted setup. In Bhavani M. Thuraisingham, David Evans, Tal Malkin, and

- Dongyan Xu, editors, *ACM CCS 2017*, pages 2087–2104. ACM Press, October / November 2017.
- AOS02. Masayuki Abe, Miyako Ohkubo, and Koutarou Suzuki. 1-out-of-n signatures from a variety of keys. In Yuliang Zheng, editor, *ASIACRYPT 2002*, volume 2501 of *LNCS*, pages 415–432. Springer, Heidelberg, December 2002.
- BBB⁺18. Benedikt Bünz, Jonathan Bootle, Dan Boneh, Andrew Poelstra, Pieter Wuille, and Greg Maxwell. Bulletproofs: Short proofs for confidential transactions and more. In *2018 IEEE Symposium on Security and Privacy*, pages 315–334. IEEE Computer Society Press, May 2018.
- BCC⁺15. Jonathan Bootle, Andrea Cerulli, Pyrros Chaidos, Essam Ghadafi, Jens Groth, and Christophe Petit. Short accountable ring signatures based on DDH. In Günther Pernul, Peter Y. A. Ryan, and Edgar R. Weippl, editors, *Computer Security - ESORICS 2015 - 20th European Symposium on Research in Computer Security, Vienna, Austria, September 21-25, 2015, Proceedings, Part I*, volume 9326 of *Lecture Notes in Computer Science*, pages 243–265. Springer, 2015.
- BCG⁺17. Jonathan Bootle, Andrea Cerulli, Essam Ghadafi, Jens Groth, Mohammad Hajiabadi, and Sune K. Jakobsen. Linear-time zero-knowledge proofs for arithmetic circuit satisfiability. In Tsuyoshi Takagi and Thomas Peyrin, editors, *ASIACRYPT 2017, Part III*, volume 10626 of *LNCS*, pages 336–365. Springer, Heidelberg, December 2017.
- BCR⁺19. Eli Ben-Sasson, Alessandro Chiesa, Michael Riabzev, Nicholas Spooner, Madars Virza, and Nicholas P. Ward. Aurora: Transparent succinct arguments for R1CS. In Yuval Ishai and Vincent Rijmen, editors, *EUROCRYPT 2019, Part I*, volume 11476 of *LNCS*, pages 103–128. Springer, Heidelberg, May 2019.
- BDG20. Mihir Bellare, Hannah Davis, and Felix Günther. Separate your domains: NIST PQC kems, oracle cloning and read-only indistinguishability. In Anne Canteaut and Yuval Ishai, editors, *Advances in Cryptology - EUROCRYPT 2020 - 39th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zagreb, Croatia, May 10-14, 2020, Proceedings, Part II*, volume 12106 of *Lecture Notes in Computer Science*, pages 3–32. Springer, 2020.
- BGW99. László Babai, Anna Gál, and Avi Wigderson. Superpolynomial lower bounds for monotone span programs. In *Combinatorica*, volume 19, pages 301–319, 1999.
- BM90. Mihir Bellare and Silvio Micali. Non-interactive oblivious transfer and applications. In Gilles Brassard, editor, *CRYPTO’89*, volume 435 of *LNCS*, pages 547–557. Springer, Heidelberg, August 1990.
- BN06. Mihir Bellare and Gregory Neven. Multi-signatures in the plain public-key model and a general forking lemma. In Ari Juels, Rebecca N. Wright, and Sabrina De Capitani di Vimercati, editors, *ACM CCS 2006*, pages 390–399. ACM Press, October / November 2006.
- BPW12. David Bernhard, Olivier Pereira, and Bogdan Warinschi. How not to prove yourself: Pitfalls of the Fiat-Shamir heuristic and applications to Helios. In Xiaoyun Wang and Kazue Sako, editors, *ASIACRYPT 2012*, volume 7658 of *LNCS*, pages 626–643. Springer, Heidelberg, December 2012.
- BR93. Mihir Bellare and Phillip Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In Dorothy E. Denning, Raymond Pyle, Ravi Ganesan, Ravi S. Sandhu, and Victoria Ashby, editors, *ACM CCS 93*, pages 62–73. ACM Press, November 1993.

- BSS02. Emmanuel Bresson, Jacques Stern, and Michael Szydlo. Threshold ring signatures and applications to ad-hoc groups. In Moti Yung, editor, *CRYPTO 2002*, volume 2442 of *LNCS*, pages 465–480. Springer, Heidelberg, August 2002.
- CDM00. Ronald Cramer, Ivan Damgård, and Philip D. MacKenzie. Efficient zero-knowledge proofs of knowledge without intractability assumptions. In Hideki Imai and Yuliang Zheng, editors, *PKC 2000*, volume 1751 of *LNCS*, pages 354–372. Springer, Heidelberg, January 2000.
- CDS94. Ronald Cramer, Ivan Damgård, and Berry Schoenmakers. Proofs of partial knowledge and simplified design of witness hiding protocols. In Yvo Desmedt, editor, *CRYPTO'94*, volume 839 of *LNCS*, pages 174–187. Springer, Heidelberg, August 1994.
- CFQ19. Matteo Campanelli, Dario Fiore, and Anaïs Querol. LegoSNARK: Modular design and composition of succinct zero-knowledge proofs. In Lorenzo Cavallaro, Johannes Kinder, XiaoFeng Wang, and Jonathan Katz, editors, *ACM CCS 2019*, pages 2075–2092. ACM Press, November 2019.
- CPSV16. Michele Ciampi, Giuseppe Persiano, Luisa Siniscalchi, and Ivan Visconti. A transform for NIZK almost as efficient and general as the Fiat-Shamir transform without programmable random oracles. In Eyal Kushilevitz and Tal Malkin, editors, *TCC 2016-A, Part II*, volume 9563 of *LNCS*, pages 83–111. Springer, Heidelberg, January 2016.
- Cra97. Ronald Cramer. *Modular Design of Secure yet Practical Cryptographic Protocols*. PhD thesis, University of Amsterdam, January 1997.
- DK18. Apoorvaa Deshpande and Yael Kalai. Proofs of ignorance and applications to 2-message witness hiding. Cryptology ePrint Archive, Report 2018/896, 2018.
- FHJ20. Marc Fischlin, Patrick Harasser, and Christian Janson. Signatures from sequential-OR proofs. In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020, Part III*, volume 12107 of *LNCS*, pages 212–244. Springer, Heidelberg, May 2020.
- FLS90. Uriel Feige, Dror Lapidot, and Adi Shamir. Multiple non-interactive zero knowledge proofs based on a single random string (extended abstract). In *31st FOCS*, pages 308–317. IEEE Computer Society Press, October 1990.
- FLWL20. Hanwen Feng, Jianwei Liu, Qianhong Wu, and Ya-Nan Li. Traceable ring signatures with post-quantum security. In Stanislaw Jarecki, editor, *CT-RSA 2020*, volume 12006 of *LNCS*, pages 442–468. Springer, Heidelberg, February 2020.
- FS87. Amos Fiat and Adi Shamir. How to prove yourself: Practical solutions to identification and signature problems. In Andrew M. Odlyzko, editor, *CRYPTO'86, Part III*, volume 263 of *LNCS*, pages 186–194. Springer, Heidelberg, August 1987.
- GKM⁺18. Jens Groth, Markulf Kohlweiss, Mary Maller, Sarah Meiklejohn, and Ian Miers. Updatable and universal common reference strings with applications to zk-SNARKs. In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO 2018, Part III*, volume 10993 of *LNCS*, pages 698–728. Springer, Heidelberg, August 2018.
- GMW86. Oded Goldreich, Silvio Micali, and Avi Wigderson. Proofs that yield nothing but their validity and a methodology of cryptographic protocol design (extended abstract). In *27th FOCS*, pages 174–187. IEEE Computer Society Press, October 1986.

- GPS08. Steven D. Galbraith, Kenneth G. Paterson, and Nigel P. Smart. Pairings for cryptographers. *Discret. Appl. Math.*, 156(16):3113–3121, 2008.
- GPV08. Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In Richard E. Ladner and Cynthia Dwork, editors, *40th ACM STOC*, pages 197–206. ACM Press, May 2008.
- Gro16. Jens Groth. On the size of pairing-based non-interactive arguments. In Marc Fischlin and Jean-Sébastien Coron, editors, *EUROCRYPT 2016, Part II*, volume 9666 of *LNCS*, pages 305–326. Springer, Heidelberg, May 2016.
- GS92. Michelangelo Grigni and Michael Sipser. Monotone complexity. In *Proceedings of the London Mathematical Society, Symposium on Boolean Function Complexity*, pages 57–75, 1992.
- KW93. Mauricio Karchmer and Avi Wigderson. On span programs. In *Proceedings of the Eighth Annual Structure in Complexity Theory Conference*, pages 102–111. IEEE Computer Society, 1993.
- Lin15. Yehuda Lindell. An efficient transform from sigma protocols to NIZK with a CRS and non-programmable random oracle. In Yevgeniy Dodis and Jesper Buus Nielsen, editors, *TCC 2015, Part I*, volume 9014 of *LNCS*, pages 93–109. Springer, Heidelberg, March 2015.
- LLNW16. Benoît Libert, San Ling, Khoa Nguyen, and Huaxiong Wang. Zero-knowledge arguments for lattice-based accumulators: Logarithmic-size ring signatures and group signatures without trapdoors. In Marc Fischlin and Jean-Sébastien Coron, editors, *EUROCRYPT 2016, Part II*, volume 9666 of *LNCS*, pages 1–31. Springer, Heidelberg, May 2016.
- LLNW17. Benoît Libert, San Ling, Khoa Nguyen, and Huaxiong Wang. Zero-knowledge arguments for lattice-based PRFs and applications to E-cash. In Tsuyoshi Takagi and Thomas Peyrin, editors, *ASIACRYPT 2017, Part III*, volume 10626 of *LNCS*, pages 304–335. Springer, Heidelberg, December 2017.
- LMR19. Russell W. F. Lai, Giulio Malavolta, and Viktoria Ronge. Succinct arguments for bilinear group arithmetic: Practical structure-preserving cryptography. In Lorenzo Cavallaro, Johannes Kinder, XiaoFeng Wang, and Jonathan Katz, editors, *ACM CCS 2019*, pages 2057–2074. ACM Press, November 2019.
- Max11. Gregory Maxwell. Zero knowledge contingent payment. https://en.bitcoin.it/wiki/Zero_Knowledge_Contingent_Payment, 2011.
- MBKM19. Mary Maller, Sean Bowe, Markulf Kohlweiss, and Sarah Meiklejohn. Sonic: Zero-knowledge SNARKs from linear-size universal and updatable structured reference strings. In Lorenzo Cavallaro, Johannes Kinder, XiaoFeng Wang, and Jonathan Katz, editors, *ACM CCS 2019*, pages 2111–2128. ACM Press, November 2019.
- Nie02. Jesper Buus Nielsen. Separating random oracle proofs from complexity theoretic proofs: The non-committing encryption case. In Moti Yung, editor, *CRYPTO 2002*, volume 2442 of *LNCS*, pages 111–126. Springer, Heidelberg, August 2002.
- NTWZ19. Khoa Nguyen, Hanh Tang, Huaxiong Wang, and Neng Zeng. New code-based privacy-preserving cryptographic constructions. In Steven D. Galbraith and Shiho Moriai, editors, *ASIACRYPT 2019, Part II*, volume 11922 of *LNCS*, pages 25–55. Springer, Heidelberg, December 2019.
- PR18. Toniann Pitassi and Robert Robere. Lifting nullstellensatz to monotone span programs over any field. In Ilias Diakonikolas, David Kempe, and

- Monika Henzinger, editors, *50th ACM STOC*, pages 1207–1219. ACM Press, June 2018.
- PS00. David Pointcheval and Jacques Stern. Security arguments for digital signatures and blind signatures. *Journal of Cryptology*, 13(3):361–396, June 2000.
- Raz85. A. A. Razborov. Lower bounds on monotone complexity of the logical permanent. *Mathematical notes of the Academy of Sciences of the USSR*, 37(6):485–493, 1985.
- RST01. Ronald L. Rivest, Adi Shamir, and Yael Tauman. How to leak a secret. In Colin Boyd, editor, *ASIACRYPT 2001*, volume 2248 of *LNCS*, pages 552–565. Springer, Heidelberg, December 2001.
- Ste96. Jacques Stern. A new paradigm for public key identification. *IEEE Trans. Inf. Theory*, 42(6):1757–1768, 1996.
- Tar88. É. Tardos. The gap between monotone and non-monotone circuit complexity is exponential. *Combinatorica*, 8(1):141–142, 1988.