NIZKs with an Untrusted CRS: Security in the Face of Parameter Subversion

Mihir Bellare[1(✉)], Georg Fuchsbauer[2], and Alessandra Scafuro[3]

[1] Department of Computer Science and Engineering, University of California, San Diego, San Diego, USA
mihir@eng.ucsd.edu
[2] Inria, Ecole Normale Supérieure, CNRS and PSL Research University, Paris, France
georg.fuchsbauer@ens.fr
[3] Department of Computer Science, North Carolina State University, Raleigh, USA
ascafur@ncsu.edu

Abstract. Motivated by the subversion of "trusted" public parameters in mass-surveillance activities, this paper studies the security of NIZKs in the presence of a maliciously chosen common reference string. We provide definitions for subversion soundness, subversion witness indistinguishability and subversion zero knowledge. We then provide both negative and positive results, showing that certain combinations of goals are unachievable but giving protocols to achieve other combinations.

1 Introduction

The summer of 2013 brought shocking news of mass surveillance being conducted by the NSA and its counter-parts in other countries. The documents revealed new ways in which the adversary compromises security, ways not covered by standard models and definitions in cryptography. This opens up a new research agenda, namely to formalize security goals that defend against these novel attacks, and study the achievability of these goals. This agenda is being pursued along several fronts. The front we pursue here is *parameter subversion*, namely the compromise of security by the malicious creation of supposedly trusted public parameters for cryptographic systems. The representative example is the Dual EC random number generator (RNG).

Dual EC. Dual EC is an NSA-designed, elliptic-curve-based random number generator, standardized as NIST SP 800-90 and ANSI X9.82. BLN [14] say that its story is "one of the most interesting in modern cryptography." The RNG includes two points P, Q on an elliptic curve that function as public parameters for the algorithm. At the Crypto 2007 rump session, Shumow and Ferguson noted that anyone who knew the discrete logarithm of P to base Q, meaning a scalar s such that $P = sQ$, could predict generator outputs. In a Wired Magazine article the same year, Schneier warned against Dual EC because it "just might contain

J.H. Cheon and T. Takagi (Eds.): ASIACRYPT 2016, Part II, LNCS 10032, pp. 777–804, 2016.
DOI: 10.1007/978-3-662-53890-6_26

a backdoor for the NSA." The NSA's response was that they had "generated P, Q in a secure, classified way." But the Snowden revelations (documents from project Bullrun and SIGINT) show that Dual EC was part of a systematic NSA effort to subvert standards. And in 2014, CNEGLRBMSF [24] showed the practical effectiveness of the subversion by demonstrating how the backdoor could be exploited to break TLS.

Two things are remarkable. The first is that the "trusted" public parameters were in fact subverted. The second is the effort put into ensuring that the subverted parameters were standardized and used. NSA-based pressure and lobbying not only lead to Dual EC remaining a US standard but even to its being in an international standard, ISO 18031:2005. In 2013 Reuters reported that the NSA paid RSA corporation $10 million to make Dual EC the default method for random number generation in their BSafe library.

Cryptography resistant to parameter subversion. The lesson to take away is that a cryptographic system that relies on public parameters assumed to have been honestly generated, say by some "trusted" party, is at great practical risk from the possibility that the parameters were in fact maliciously generated with intent to subvert security of their use. We suggest that in response we should develop cryptography that is resistant to parameter subversion. This means that it should provide its usual security with trusted parameters, but retain as much security as possible when the parameters are maliciously generated.

Parameters arise in many places in cryptography, but a prominent one that springs to mind are non-interactive zero-knowledge (NIZK) systems, where the common reference string (CRS) is assumed to be honestly generated. NIZKs are not only important in their own right but used in a wide variety of applications, so their security under parameter subversion has far-reaching effects. This paper provides a treatment of resistance to parameter subversion for NIZKs, with definitions, negative results and positive results.

NIZKs. Non-interactive zero-knowledge systems originate with BFM [17] and BDMP [16] and have since seen an explosion in constructions and applications. The Groth-Sahai framework for efficient NIZKs [44] is widely utilized and we are seeing not only efficient NIZKs but also their implementation in systems [12, 13, 31, 39, 44]. Structure-preserving cryptography [1, 2, 40] was developed to allow these NIZKs to be used for efficient applications.

The NIZK model postulates a common reference string (CRS) that has been honestly generated according to some distribution. The pragmatics of how this is done receives little explicit attention. Some early works talk of using digits of π and others speak whimsically of "a random string in the sky," but for the most part the understanding is that a trusted party will generate, and make public, the CRS. In light of the above, however, we must be concerned that the CRS is in fact maliciously generated. This is the issue addressed by our work.

An immediate avenue of attack that may come to mind is the following. NIZK security requires that there is a simulator that generates a simulated CRS (indistinguishable from the honest one) together with a trapdoor allowing the simulator to generate proofs without knowing the witness. What if the subvertor

generates the CRS via the simulator, so that it knows the trapdoor? Since this CRS is indistinguishable from an honestly generated one, the subversion will not be detected. Now, what does the subvertor gain? This seems to depend on the particular system and its properties. For example, the subvertor may be able to generate proofs of *false* statements and violate soundness. In some cases the trapdoor permits extraction of witnesses from honest proofs, in which case the subvertor would be able to violate zero knowledge. What we see here is that features built into the standard notions and constructions of NIZKs turn out to be potential liabilities in the face of subversion. Put another way, current NIZKs have the possibility of subversion effectively built into the security requirement because the simulator works by "subverting" the CRS.

Two remarks with regard to the above. (1) First, if it is unclear what is going on, or what conclusion to draw, there is a good reason, namely that we are trying to think or talk about what subversion does in the absence of a clear understanding of the subversion-resistance goal, effectively jumping the gun. To be able to effectively assess security we first need precise definitions of the new goal(s) underlying resistance to CRS subversion. Providing such definitions is the first contribution of this paper. (2) Second, while the above discussion may lead one to be pessimistic, we will see that in fact a surprising amount of security can be retained even under a maliciously generated CRS.

NIZK SECURITY, NOW. To discuss the new goals in subversion-resistant NIZKs we first back up to recall the standard goals in the current model where the CRS is trusted and assumed to be honestly generated. We distinguish three standard goals for a non-interactive (NI) system Π relative to an **NP** relation R defining the language $L(\mathsf{R}) \in \mathbf{NP}$. The formalizations are recalled in Sect. 4.

SND: (Soundness) It is hard for an adversary, given an honestly generated *crs*, to find an $x \notin L(\mathsf{R})$ together with a valid proof π (meaning one that the verification algorithm $\Pi.\mathsf{V}$ accepts) for x relative to *crs*.

WI: (Witness indistinguishability) Assuming *crs* is honestly generated, an adversary can't tell under which of two valid witnesses an honest proof (i.e., generated by the prover algorithm $\Pi.\mathsf{P}$ under *crs*) for an instance x was created, and this even holds for multiple, adaptively chosen instances depending on *crs*.

ZK: (Zero-knowledge) There is a simulator $\Pi.\mathsf{Sim.crs}$ returning a simulated CRS crs_0 and associated trapdoor *std*, and an accomplice simulator $\Pi.\mathsf{Sim.pf}$ taking an instance $x \in L(\mathsf{R})$ and *std* and returning a proof, such that an adversary given crs_b cannot tell whether a proof it receives was created honestly (with the honest prover algorithm, an honest crs_1 and a witness; the $b = 1$ case) or via $\Pi.\mathsf{Sim.pf}$ (the $b = 0$ case). Moreover this holds even for multiple, adaptively chosen instances depending on crs_b.

NIZK SECURITY UNDER SUBVERSION. The key change in our model is that the adversary generates the CRS. It can retain, via its coins r, some kind of "back-

door" related to this CRS. In Sect. 4 we formalize the following goals:

S-SND: (Subversion soundness) It is hard for the adversary to generate a (malicious) CRS *crs* together with an instance $x \notin L(\mathsf{R})$ and a valid proof π for x relative to *crs*. (The goal of the subvertor here is to create a CRS that allows it to give proofs of false statements.)

S-WI: (Subversion witness indistinguishability) Even if the adversary creates *crs* maliciously and retains the corresponding coins r, it can't tell under which of two valid witnesses an honest proof (meaning one generated by the prover algorithm $\Pi.\mathsf{P}$ under the subverted *crs*) for an instance x was created, and moreover this holds even for multiple, adaptively chosen instances depending on *crs*.

S-ZK: (Subversion zero knowledge) For any adversary X creating a malicious CRS crs_1 using coins r_1, there is a simulator $\mathsf{S.crs}$ returning not only a simulated CRS crs_0 and associated trapdoor *std but also simulated coins* r_0, and an accomplice simulator $\mathsf{S.pf}$ taking an instance $x \in L(\mathsf{R})$ and *std* and returning a proof, such that an adversary A given crs_b, r_b cannot tell whether a proof it receives was created honestly (with $\Pi.\mathsf{P}$ using crs_1 and a witness; the $b = 1$ case) or via $\mathsf{S.pf}$ (the $b = 0$ case). Moreover this holds even for multiple, adaptively chosen instances depending on crs_b, r_b.

The right side of Fig. 1 may help situate the notions. It shows the obvious relations: S-X implies X; ZK implies WI and S-ZK implies S-WI.

Achievability. Is subversion resistance achievable? This question first needs to be meaningfully posed. The subversion resistance goals are easy to achieve *in isolation*. For example, S-SND is achieved for any **NP** relation by having the prover send the witness, but this is not ZK. S-ZK is achieved by having the prover send the empty string as the proof and having the verifier always accept, but this is not SND. Such trivial constructions are un-interesting. The interesting question is whether meaningful combinations of the goals are simultaneously achievable. A pragmatic viewpoint is that we already have systems achieving SND+WI+ZK. We want to "upgrade" these to get some resistance to subversion. While retaining SND, WI and ZK, what can be added from the list S-SND, S-WI, S-ZK? Can we have them all? Are things so bad that we can have none? We will be able to completely categorize what is achievable and what is not and will see that the truth is somewhere between these extremes and on the whole the news is perhaps more positive than we might have expected. Our core results are summarized in the table on the left side of Fig. 1. In any row, we are considering simultaneously achieving the notions indicated by the bullets. The last column indicates whether or not it is possible. We now discuss these results, beginning with the negative result of the first row.

Negative result. We first ask whether we can achieve S-SND (soundness for a malicious CRS) while retaining what we have now, namely SND, WI and ZK. Result **N** (the first row of Fig. 1) indicates that we cannot. It says that

	Standard			Subversion resistant			Achievable?
	SND	ZK	WI	S-SND	S-ZK	S-WI	
N		•		•			✗ Thm. 1
P1	•	•	•		•	•	✓ Thm. 3
P2	•		•	•		•	✓ Thm. 5
P3	•	•	•			•	✓ Thm. 6

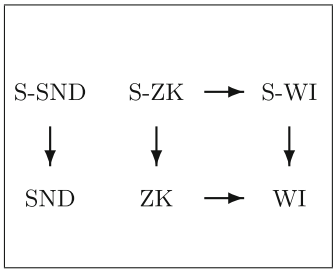

Fig. 1. Left: Achievability chart showing our negative result **N** and positive results **P1**, **P2**, **P3**. In a row we refer to simultaneously achieving all selected notions. **Right:** Relations.

there is no NI system that achieves both ZK and S-SND. (More precisely, this is only possible for trivial **NP**-relations, i.e., where verifiers can check if $x \in L(\mathsf{R})$ themselves.) We stress that ZK here is the standard notion where the CRS is honest. We are not asking for S-ZK but only to retain ZK. The proof of Theorem 1 establishing this uses the paradigm of GO [36] of using the simulator to break soundness.

Positive results. Figure 1 lists three positive results that we discuss in turn:

P1: The most desirable target is S-ZK. By result **N** it cannot be achieved in combination with S-SND. The next best thing would be to get it in combination with SND. We show in Theorem 3 that this is possible. Since S-ZK implies ZK, S-WI and WI, this yields result **P1** of the table of Fig. 1, showing we can simultaneously achieve all notions but S-SND. Theorem 3 is based on a knowledge-of-exponent assumption (KEA) in a group equipped with a bilinear map. The assumption is certainly strong, but (1) this is to be expected since our goal implies certain forms of 2-move interactive ZK that have themselves only been achieved under extractability assumptions [15], (2) similar assumptions have been made before [39], and (3) unlike other knowledge assumptions [15], our assumption is not ruled out assuming indistinguishability obfuscation. See the beginning of Sect. 6.1 for a high-level description of the ideas of our construction.

P2: The question left open by **P1** is whether there is some meaningful way to achieve S-SND. (It is the one item missing in row **P1**.) We know from result **N** that we cannot do this in combination with ZK. Result **P2** of the table of Fig. 1 says that we can do the best possible given this limitation. Namely we can simultaneously achieve both S-SND and S-WI (and thus SND and WI). Theorem 5 establishing this is under a standard assumption, namely the decision-linear assumption (DLin). It follows easily from the existence of a SND and WI NI system with trivial CRS under DLin [42] and the observation (Lemma 4) that any such system is obviously also S-SND and S-WI.

P3: Result **P3** of the Fig. 1 represents "hedging." The system has the desired properties (SND, WI, ZK) under an honest CRS. When the CRS is maliciously chosen, it does not break completely; it retains witness indistinguishability in the form of S-WI. In practice this offers quite a bit of protection. Our hedging construction combines a PRG with a zap. (A zap is a 2-move witness-indistinguishable interactive protocol [30].)

Result **P3** may seem redundant; isn't it implied by **P1**? (Indeed it selects a strict subset of the notions selected by **P1**.) While **P1** uses strong (extractability) assumptions, **P3** is established in Theorem 6 under the minimal assumption that some SND+WI+ZK NI system exists. Our hedging thus adds no extra assumptions. This is because a zap can be built from any SND+ZK NI system [30].

Full achievability picture. The broad question we have asked is, which combinations of the six notions SND, WI, ZK, S-SND, S-WI, S-ZK are simultaneously achievable? Fig. 1 looks at four combinations. But there are in principle 2^6 combinations about which one could ask. In the full version [6] we go systematically over *all* combinations and evaluate achievability. We are able to give the answer in all cases. Briefly, Fig. 1 covers the interesting cases, which is why we have focused on those here, and other cases are dealt with relatively easily.

Other notions. We have been selective rather than exhaustive with regard to which notions to consider in this setting, focusing on the basic soundness, witness indistinguishability and zero knowledge. There are many other notions in this area that could be considered including robustness, simulation soundness and extractability [26,28,38,41] but it seems fairly apparent that these stronger notions will be subject to commensurately strong negative results with regard to security under CRS subversion. For example, extractability asks that the simulator can create a CRS such that, with a trapdoor it withholds, it can extract the witness from a valid proof. But if so, a subvertor can create the CRS like the simulator so that it has the trapdoor and can also extract the witness.

2 Discussion and Related Work

Relation to 2-move protocols. There is a natural connection between NI systems and 2-move interactive protocols in which NI system Π corresponds to the protocol $\mathsf{2MV}$ in which the verifier first sends the CRS and the prover sends the proof in the second move. We can then think of the following correspondence of notions for Π and $\mathsf{2MV}$: S-WI $\leftrightarrow$ ZAP; ZK $\leftrightarrow$ honest-verifier ZK; S-ZK $\leftrightarrow$ full (cheating-verifier) ZK. This analogy provides intuition and insight and opens up connections we exploit for both positive and negative results, but one must be wary that the analogy is not fully accurate in either direction. We look separately at this for negative and positive results.

On the negative side, many forms of 2-move ZK are impossible [4,36]. This does not directly imply that S-ZK is impossible because S-ZK does not imply these particular forms of 2-move ZK. For example, S-ZK does not incorporate

auxiliary inputs and thus does not imply auxiliary-input 2-move ZK, so the fact that the latter is ruled out [36] does not mean the former is ruled out. (Why does our definition of S-ZK not incorporate auxiliary inputs? One reason was exactly to avoid the impossibility results. But also, an important reason to introduce auxiliary inputs in the interactive case was to be able to prove that ZK for multiple instances is provided, by sequential composition. But our S-ZK formulation already and directly requires security for multiple, adaptively chosen instances, removing the main motivation for auxiliary inputs.)

On the positive side, some forms of 2-move ZK are possible [4,5,15,50]. A natural question is whether one can obtain S-ZK+SND (the goal of **P1**) from them by the obvious transformation, namely to make the verifier's move the CRS. Unfortunately, this does not in general achieve S-ZK. In particular the simulation requirement for S-ZK is stronger than for ZK because the simulated CRS must be produced upfront without knowing the instance, and then the simulator must be able to adaptively produce simulated proofs for multiple instances.

So 2-move ZK as claimed and proven by [4,5,15] does not directly yield S-ZK. The next natural question is whether the protocols of these papers can, nonetheless, be directly shown to have the stronger properties needed to obtain S-ZK. This appears to be the case for the protocols of [4,15,50], because the verifier's first message does not depend on the instance. Starting from BLV [4], the assumption would be that Micali's conjecture [48] (there exist CS proofs or two-round universal arguments) is true. Starting from BCPR [15], the assumption would be the existence of privately verifiable P-delegation, 1-hop FHE, and a complexity-leveraging commitment scheme. In this light, we have chosen to present our knowledge of exponent based **P1** construction as a concrete, self-contained illustration of one simple route to S-ZK+SND from a plausible assumption, but other routes are possible. We do note that BLV [4] themselves view their assumption as so strong that they hesitate to call their result a positive one, instead referring to it as "a negative result on negative results."

BP [5] build one-message ZK arguments, but the simulation is super polynomial time. (This is also true of the construction of Pass [50].) These would thus yield S-ZK with super-polynomial-time simulation. But we require simulation for S-ZK to be polynomial time. This is in keeping with the intuition behind zero-knowledge that the entity running the verifier in the protocol should be able to run the simulator to produce a similar view.

Finally, in the bare public-key model of [21], Wee [56] constructs a weak non-uniform non-interactive zero-knowledge argument. This can be turned into a NI system by using the verifier's public key as the CRS. However this form of ZK allows a super-polynomial simulator whose size depends on the size of the distinguisher and the distinguishing gap, and this is weaker than S-ZK. Also Wee's [56] construction is only proved for one instance, while in S-ZK we require security for multiple, adaptively-chosen instances.

Context. Resistance of NIZKs to parameter subversion may not be of *immediate* practical relevance but we believe it is an important long-term consideration for this technology. The foundational tradition has always had as its stated goal

to model and capture realistic, practical attacks and then investigate theoretically whether or not security can be achieved. Parameter subversion is such a realistic attack not previously considered, and it leads us to revisit the foundations of NIZKs to bring it into the picture. We are seeing large efforts in the creation of efficient NIZKs and their implementation in systems towards eventual applications [11–13,31,39,44]. For security, parameter subversion must be kept in mind from the start.

A standard suggestion to protect against CRS subversion is to generate the CRS via a multi-party computation protocol so that no particular party controls the outcome. This is pursued in [11]. The effectiveness and practicality of this solution are not very clear. What parties would perform this task, and why can we trust *any* of them? The Snowden revelations indicate that corporations cooperate with the NSA toward subversion, either willingly or due to court orders. NIZKs with built-in resistance to subversion, as we define and achieve, provide greater protection.

One might note that in some applications, such as the use of NIZKs for signatures [7,23,28] and IND-CCA encryption [29,49], users can pick their own CRS and be confident of its quality. However this blows up key sizes and increases system complexity. It would be more convenient if there were a single, global CRS, in which case resistance to subversion matters.

CPs [22] study UC-secure computation in a model where the CRS is drawn from a distribution that is adversarially chosen subject to several restrictions, including that it has high min-entropy and is efficiently sampleable via an algorithm known to the simulator. They do not consider NIZKs, and in their model the CRS is not chosen fully maliciously, with no restrictions, as in our model. GO [41] studied the "multi-CRS" model where the adversary can substitute t out of m CRSs, GGJS [33] consider replacing a single trusted setup in UC with multiple, untrusted ones and KKZZ [46] consider distributing the setup for UC-secure multi-party computation. Concern with trust in a CRS is exhibited in the context of elections by KZZ [47], who have the CRS generated by the election authority using the voter's coins.

Algorithm-substitution attacks, studied in [3,9], are another form of subversion, going back to the broader framework of kleptography [57,58]. Back-doored blockciphers were studied in [51–53]. DGGJR [27] provide a formal treatment of back-dooring of PRGs in response to the Dual EC debacle. The cliptography framework [54] aims to capture many forms of subversion.

3 Notation

The empty string is denoted by ε. If x is a (binary) string then $|x|$ is its length. If S is a finite set then $|S|$ denotes its size and $s \leftarrow\!\!\$\, S$ denotes picking an element uniformly from S and assigning it to s. We denote by $\lambda \in \mathbb{N}$ the security parameter and by 1^λ its unary representation. Algorithms are randomized unless otherwise indicated. "PT" stands for "polynomial time", whether for randomized or deterministic algorithms. By $y \leftarrow A(x_1, \ldots; r)$ we denote the operation

of running A on inputs $x_1, \ldots$ and coins r and letting y denote the output. By $y \leftarrow\!\!\$\ A(x_1, \ldots)$, we denote letting $y \leftarrow A(x_1, \ldots; r)$ for random r. We denote by $[A(x_1, \ldots)]$ the set of points that have positive probability of being output by A on inputs $x_1, \ldots$ Adversaries are algorithms. Complexity is uniform throughout: scheme algorithms and adversaries are Turing Machines, not circuit families.

For our security definitions and some proofs we use the code-based game playing framework of [10]. A game G (e.g. Fig. 2) usually depends on some scheme and executes one or more adversaries. It defines oracles for the adversaries as procedures. The game eventually returns a boolean. We let $\Pr[\mathrm{G}]$ denote the probability that G returns true.

4 Security of NIZKs Under CRS Subversion

We first recall and discuss standard notions of NIZK security in the setting used until now where the CRS is trusted. We then formulate new notions of NIZK security in the setting where the CRS is subverted, starting with the syntax.

4.1 NP Relations and NI Systems

NP relations. Proofs pertain to membership in an **NP** language defined by an **NP** relation, and we begin with the latter. Suppose $\mathsf{R}\colon\ \{0,1\}^* \times \{0,1\}^* \to \{\mathsf{true}, \mathsf{false}\}$. For $x \in \{0,1\}^*$ we let $\mathsf{R}(x) = \{\, w \,:\, \mathsf{R}(x,w) = \mathsf{true} \,\}$ be the *witness set* of x. We say that R is an **NP** relation if it is PT and there is a polynomial $\mathsf{R.wl}\colon\ \mathbb{N} \to \mathbb{N}$ called the maximum witness length such that every w in $\mathsf{R}(x)$ has length at most $\mathsf{R.wl}(|x|)$ for all $x \in \{0,1\}^*$. We let $L(\mathsf{R}) = \{\, x \,:\, \mathsf{R}(x) \neq \emptyset \,\}$ be the *language* associated to R. The fact that R is an **NP** relation means that $L(\mathsf{R}) \in \mathbf{NP}$. We now go on to security properties, first giving formal definitions and then discussions.

NI systems. A non-interactive (NI) system specifies the syntax of the proof system. We can then consider various security attributes, including soundness, zero knowledge and witness indistinguishability. Formally, a NI system Π for R specifies the following PT algorithms. Via $crs \leftarrow\!\!\$\ \Pi.\mathsf{Pg}(1^\lambda)$ one generates a common reference string crs. Via $\pi \leftarrow\!\!\$\ \Pi.\mathsf{P}(1^\lambda, crs, x, w)$ the honest prover, given x and $w \in \mathsf{R}(x)$, generates a proof π that $x \in L(\mathsf{R})$. Via $d \leftarrow \Pi.\mathsf{V}(1^\lambda, crs, x, \pi)$ a verifier can produce a decision $d \in \{\mathsf{true}, \mathsf{false}\}$ indicating whether π is a valid proof that $x \in L(\mathsf{R})$. We require (perfect) completeness, namely $\Pi.\mathsf{V}(1^\lambda, crs, x, \Pi.\mathsf{P}(1^\lambda, crs, x, w)) = \mathsf{true}$ for all $\lambda \in \mathbb{N}$, all $crs \in [\Pi.\mathsf{Pg}(\lambda)]$, all $x \in L(\mathsf{R})$ and all $w \in \mathsf{R}(x)$. We also require that $\Pi.\mathsf{V}$ returns false if any of its arguments is $\perp$.

4.2 Notions for Honest CRS: SND, WI and ZK

Soundness. Soundness asks that it be hard to create a valid proof for $x \notin L(\mathsf{R})$. Formally, we say that Π is sound for R, abbreviated SND, if $\mathbf{Adv}^{\mathrm{snd}}_{\Pi,\mathsf{R},\mathsf{A}}(\cdot)$ is

negligible for all PT adversaries A, where $\mathbf{Adv}^{\mathrm{snd}}_{\Pi,\mathsf{R},\mathsf{A}}(\lambda) = \Pr[\mathrm{SND}_{\Pi,\mathsf{R},\mathsf{A}}(\lambda)]$ and game SND is specified in Fig. 2. This is a computational soundness requirement as opposed to a statistical one, as is sufficient for applications.

WI. This notion [32] requires that a PT adversary, which chooses two witnesses, cannot tell which one was used to create a proof. Formally, we say that Π is witness-indistinguishable (WI) for R, if $\mathbf{Adv}^{\mathrm{wi}}_{\Pi,\mathsf{R},\mathsf{A}}(\cdot)$ is negligible for all PT adversaries A, where $\mathbf{Adv}^{\mathrm{wi}}_{\Pi,\mathsf{R},\mathsf{A}}(\lambda) = 2\Pr[\mathrm{WI}_{\Pi,\mathsf{R},\mathsf{A}}(\lambda)] - 1$ and game WI is specified in Fig. 2. In this game, an adversary A can request a proof for x under one of two witnesses w_0, w_1. It is returned an honestly generated proof under w_b where b is the challenge bit. It can adaptively request and obtain many such proofs before outputting a guess b' for b. The game returns true if this guess is correct.

ZK. We say that Π is zero-knowledge for R, abbreviated ZK, if Π specifies additional PT algorithms $\Pi.\mathsf{Sim.crs}$ and $\Pi.\mathsf{Sim.pf}$ such that $\mathbf{Adv}^{\mathrm{zk}}_{\Pi,\mathsf{R},\mathsf{A}}(\cdot)$ is negligible for all PT adversaries A, where $\mathbf{Adv}^{\mathrm{zk}}_{\Pi,\mathsf{R},\mathsf{A}}(\lambda) = 2\Pr[\mathrm{ZK}_{\Pi,\mathsf{R},\mathsf{A}}(\lambda)] - 1$ and game ZK is specified in Fig. 2. Adversary A can adaptively request proofs by supplying an instance and a valid witness for it. The proof is produced either by the honest prover using the witness, or by the proof simulator $\Pi.\mathsf{Sim.pf}$ using a trapdoor *std*. The adversary outputs a guess b' as to whether the proofs were real or simulated.

Discussion. The classical definitions of soundness and zero knowledge for proof systems [37] were in what we will call the complexity-theoretic style. The soundness condition said that for all $x \notin L(\mathsf{R})$, the probability that a dishonest prover could convince the honest verifier to accept was low. Zero knowledge, similarly, looked at distributions associated to a fixed $x \in L(\mathsf{R})$ and then at ensembles over x. The first definition for NIZK was similar [16]. But over time, NIZK definitions have adapted to what we call a cryptographic style [26,43]. This is the style we use because it seems more prevalent now and it works better for applications. Here x is not quantified but chosen by an adversary. The definitions directly capture proofs for multiple, related statements. All adversaries are PT, meaning all metrics are computational.

One consequence of the complexity-theoretic style was a need for non-uniform complexity for adversaries and assumptions [35,37]. In [34] Goldreich made a case for uniform complexity. The cryptographic style we adopt is in this vein, and in our setting all complexity (adversaries, algorithms, assumptions) is uniform.

4.3 Notions for Subverted CRS: S-SND, S-WI and S-ZK

A core assumption in NIZKs is that the CRS is honestly generated. In light of subversion of parameters in other contexts as part of the mass-surveillance revelations, we ask what would happen if the CRS were maliciously generated. We will define subversion-resistance analogues S-SND, S-WI and S-ZK of the SND, WI, ZK goals above. The key difference is that the CRS is selected by an adversary rather than via the CRS-generation algorithm $\Pi.\mathsf{Pg}$ prescribed by Π.

GAME $\mathrm{SND}_{\Pi,\mathsf{R},\mathsf{A}}(\lambda)$	GAME $\mathrm{S\text{-}SND}_{\Pi,\mathsf{R},\mathsf{A}}(\lambda)$
$crs \leftarrow\!\!\$\ \Pi.\mathsf{Pg}(1^\lambda)$ $(x,\pi) \leftarrow\!\!\$\ \mathsf{A}(1^\lambda, crs)$ Return $(x \notin L(\mathsf{R})$ and $\Pi.\mathsf{V}(1^\lambda, crs, x, \pi))$	$(crs, x, \pi) \leftarrow\!\!\$\ \mathsf{A}(1^\lambda)$ Return $(x \notin L(\mathsf{R})$ and $\Pi.\mathsf{V}(1^\lambda, crs, x, \pi))$

GAME $\mathrm{WI}_{\Pi,\mathsf{R},\mathsf{A}}(\lambda)$	GAME $\mathrm{S\text{-}WI}_{\Pi,\mathsf{R},\mathsf{A}}(\lambda)$
$b \leftarrow\!\!\$\ \{0,1\}$ $crs \leftarrow\!\!\$\ \Pi.\mathsf{Pg}(1^\lambda)$ $b' \leftarrow\!\!\$\ \mathsf{A}^{\mathrm{PROVE}}(1^\lambda, crs)$ Return $(b = b')$	$b \leftarrow\!\!\$\ \{0,1\}$ $(crs, st) \leftarrow\!\!\$\ \mathsf{A}(1^\lambda)$ $b' \leftarrow\!\!\$\ \mathsf{A}^{\mathrm{PROVE}}(1^\lambda, crs, st)$ Return $(b = b')$
PROVE(x, w_0, w_1) If $\mathsf{R}(x, w_0) = \mathsf{false}$ or $\mathsf{R}(x, w_1) = \mathsf{false}$ then Return $\perp$ $\pi \leftarrow\!\!\$\ \Pi.\mathsf{P}(1^\lambda, crs, x, w_b)$ Return π	PROVE(x, w_0, w_1) If $\mathsf{R}(x, w_0) = \mathsf{false}$ or $\mathsf{R}(x, w_1) = \mathsf{false}$ then Return $\perp$ $\pi \leftarrow\!\!\$\ \Pi.\mathsf{P}(1^\lambda, crs, x, w_b)$ Return π

GAME $\mathrm{ZK}_{\Pi,\mathsf{R},\mathsf{A}}(\lambda)$	GAME $\mathrm{S\text{-}ZK}_{\Pi,\mathsf{R},\mathsf{X},\mathsf{S},\mathsf{A}}(\lambda)$
$b \leftarrow\!\!\$\ \{0,1\}$ $crs_1 \leftarrow\!\!\$\ \Pi.\mathsf{Pg}(1^\lambda)$ $(crs_0, std) \leftarrow\!\!\$\ \Pi.\mathsf{Sim.crs}(1^\lambda)$ $b' \leftarrow\!\!\$\ \mathsf{A}^{\mathrm{PROVE}}(1^\lambda, crs_b)$ Return $(b = b')$	$b \leftarrow\!\!\$\ \{0,1\}$ $r_1 \leftarrow\!\!\$\ \{0,1\}^{\mathsf{X.rl}(\lambda)}$; $crs_1 \leftarrow \mathsf{X}(1^\lambda; r_1)$ $(crs_0, r_0, std) \leftarrow\!\!\$\ \mathsf{S.crs}(1^\lambda)$ $b' \leftarrow\!\!\$\ \mathsf{A}^{\mathrm{PROVE}}(1^\lambda, crs_b, r_b)$ Return $(b = b')$
PROVE(x, w) If $\mathsf{R}(x, w) = \mathsf{false}$ then Return $\perp$ If $b = 1$ then $\pi \leftarrow\!\!\$\ \Pi.\mathsf{P}(1^\lambda, crs_1, x, w)$ Else $\pi \leftarrow\!\!\$\ \Pi.\mathsf{Sim.pf}(1^\lambda, crs_0, std, x)$ Return π	PROVE(x, w) If $\mathsf{R}(x, w) = \mathsf{false}$ then Return $\perp$ If $b = 1$ then $\pi \leftarrow\!\!\$\ \Pi.\mathsf{P}(1^\lambda, crs_1, x, w)$ Else $\pi \leftarrow\!\!\$\ \mathsf{S.pf}(1^\lambda, crs_0, std, x)$ Return π

Fig. 2. Games defining standard (left) and subversion (right) security of NI system Π. Top to bottom: Soundness, witness indistinguishability, zero knowledge.

SUBVERSION SOUNDNESS. Subversion soundness asks that if a subvertor creates a CRS in any way it likes, it will still be unable to prove false statements under that CRS. Formally, we say that Π is subversion-sound (abbreviated S-SND) for R if $\mathbf{Adv}^{\text{s-snd}}_{\Pi,\mathsf{R},\mathsf{A}}(\cdot)$ is negligible for all PT adversaries A, where $\mathbf{Adv}^{\text{s-snd}}_{\Pi,\mathsf{R},\mathsf{A}}(\lambda) = \Pr[\mathrm{S\text{-}SND}_{\Pi,\mathsf{R},\mathsf{A}}(\lambda)]$ and game S-SND is specified in Fig. 2. Compared to the honest-CRS game SND to the left of it, the adversary now not only generates x and π, but itself supplies crs, modeling a malicious choice of the latter.

SUBVERSION WI. Subversion WI asks that if a subvertor creates a CRS in any way it likes then it will still be unable to tell which of two witnesses was used to create a proof, even given both witnesses. Formally, we say that Π is subversion

witness-indistinguishable (S-WI) for R if $\mathbf{Adv}^{\text{s-wi}}_{\Pi,\mathsf{R},\mathsf{A}}(\cdot)$ is negligible for all PT adversaries A, where $\mathbf{Adv}^{\text{s-wi}}_{\Pi,\mathsf{R},\mathsf{A}}(\lambda) = 2\Pr[\text{S-WI}_{\Pi,\mathsf{R},\mathsf{A}}(\lambda)] - 1$ and game S-WI is specified in Fig. 2. Compared to the honest-CRS game WI, the CRS *crs* is now generated by the adversary in a first stage, along with state information *st* passed to its second stage. In the latter, via its PROVE oracle, it adaptively obtains proofs for instances of its choice under a challenge witness, and outputs a guess b' for the challenge b. The state can contain the coins of A or any trapdoor associated to *crs* that A chooses to put there helping its distinguishing task.

SUBVERSION ZK. Subversion ZK asks that for any CRS subvertor X creating a CRS in any way it likes there is a simulator able to produce the full view of the CRS subvertor, including its coins and proofs corresponding to adaptively chosen instances, without knowing the witnesses. Formally, a simulator S for X specifies PT algorithms $\mathsf{S.crs}$ and $\mathsf{S.pf}$. Now consider game S-ZK of Fig. 2 associated to $\Pi, \mathsf{R}, \mathsf{X}, \mathsf{S}$ and an adversary A. We let $\mathbf{Adv}^{\text{s-zk}}_{\Pi,\mathsf{R},\mathsf{X},\mathsf{S},\mathsf{A}}(\lambda) = 2\Pr[\text{S-ZK}_{\Pi,\mathsf{R},\mathsf{X},\mathsf{S},\mathsf{A}}(\lambda)] - 1$. We say that Π is subversion zero-knowledge (S-ZK) for R if for all PT CRS subvertors X there is a PT simulator S such that for all PT A the function $\mathbf{Adv}^{\text{s-zk}}_{\Pi,\mathsf{R},\mathsf{X},\mathsf{S},\mathsf{A}}(\cdot)$ is negligible.

In this game, if the challenge bit b is 1 then the CRS crs_1 is generated via X with the coins r_1 made explicit. Otherwise, if $b = 0$, the first stage $\mathsf{S.crs}$ of the simulator is run to produce simulated versions crs_0, r_0 not only of the CRS but *also of the coins of* X. Alongside, $\mathsf{S.crs}$ produces a simulation trapdoor *std* as in ZK to allow its second stage to simulate proofs. Now, A gets to request its PROVE oracle for proofs of instances of its choice. If $b = 1$, these are produced by the honest prover with the given witness; but if $b = 0$, they are produced via the second stage $\mathsf{S.pf}$ of the simulator using the simulation trapdoor *std* and no witness. Adversary A produces its guess b' and wins of $b' = b$.

The definition reflects that X here is like a cheating verifier in classical ZK [37]. The simulator thus needs to produce its coins as well as the transcript of its interaction with its oracle. But also, to reflect the ZK requirement of *non-interactive* systems above, more is required, namely that the simulator must first produce the simulated CRS and coins, and then, in its second stage, be able to produce simulated proofs. The definition is thus quite demanding. Note that the simulator can depend (in a non-blackbox way) on X, but not on A. The latter is important to ensure that S-ZK implies ZK.

4.4 2-Move Protocols

We will have many occasions to refer to and use 2-move interactive protocols, so we fix a syntax for them. A 2-move protocol $\mathsf{2MV}$ for **NP** relation R specifies PT algorithms $\mathsf{2MV.V}$, $\mathsf{2MV.P}$, $\mathsf{2MV.D}$. Via $(m_1, st) \leftarrow\!\!\$\ \mathsf{2MV.V}(1^\lambda, x)$ the honest verifier generates the first move message m_1 on input x, retaining associated state information *st*. Via $m_2 \leftarrow\!\!\$\ \mathsf{2MV.P}(1^\lambda, x, w, m_1)$ the honest prover generates a reply computed from x, a witness $w \in \mathsf{R}(x)$ and the first move message m_1. Deterministic decision algorithm $\mathsf{2MV.D}$ takes x, m_1, m_2, st and returns a boolean decision. Security notions will be discussed as needed.

5 Negative Result: ZK and S-SND Are Not Compatible

All the different forms of subversion security (S-SND, S-WI, S-ZK) are easy to achieve in isolation. For example sending the witness as the proof achieves S-SND (but this is not ZK). Having the verification algorithm always accept and sending the empty string as the proof achieves S-ZK (but not SND). These kinds of results are not interesting. We want to study the simultaneous achievability of meaningful combinations of the notions, meaning some kind of soundness together with some kind of zero knowledge or witness indistinguishability.

We already have NI systems that are SND+ZK and we do not want to degrade this. If now the CRS is subverted, what more can we have without losing the initial properties? The first question we ask is, can we up the ante for soundness, meaning add S-SND? That is, we want subversion soundness while retaining ZK. We will show that this is not possible.

An impossibility result in this domain means no NI system satisfying the conditions exists unless the relation R is trivial. Roughly, trivial means that the verification algorithm can decide membership in $L(\mathsf{R})$ on its own. Impossibility results of this type begin with Goldreich and Oren (GO) [36]. Their definition of R being trivial was simple, namely that it is in **BPP**. This will not suffice here, so we begin with a more precise definition of relation triviality and an explanation of why it is needed.

$\underline{\text{GAME } \mathrm{DEC}_{\mathsf{IG},\mathsf{R},\mathsf{M}}(\lambda)}$
$(x, w) \leftarrow\!\!\$\ \mathsf{IG}(1^\lambda)$; $d_1 \leftarrow \mathsf{R}(x, w)$
If $(x \in L(\mathsf{R})$ and $d_1 = \mathsf{false})$ then return false
$d_0 \leftarrow\!\!\$\ \mathsf{M}(1^\lambda, x)$; return $(d_0 \neq d_1)$

Fig. 3. Game defining language triviality

Relation triviality. The definition of a relation R being trivial if $L(\mathsf{R}) \in$ **BPP** works when the formulations of ZK and soundness are in the complexity-theoretic style, meaning the conditions refer to universally quantified inputs. As discussed in Sect. 4.2 however, our formulations, following modern treatments of NI systems in the literature, are in the cryptographic style, which is better suited for applications. Here the only instances that come into play are those that can be generated by PT algorithms, and the only positive instances that come into play are those generated with witnesses. In this setting, **BPP** will not work as a definition of triviality because membership in standard complexity classes like **BPP** refers to arbitrary inputs, not merely ones that one can generate in PT. For our purposes we thus give a definition of a language (actually an **NP** relation) being trivial, which can be seen as defining a cryptographic version of **BPP**.

Let R be an **NP** relation. An *instance generator* is a PT algorithm that on input 1^λ returns a pair (x, w). Here x is a challenge instance that may or may

not be in $L(\mathsf{R})$, and w should be in $\mathsf{R}(x)$ if $x \in L(\mathsf{R})$. Let M be an algorithm (decision procedure) taking $1^\lambda, x$ and returning a boolean representing whether or not it thinks x is in $L(\mathsf{R})$. Consider game DEC of Fig. 3 associated to $\mathsf{IG}, \mathsf{R}, \mathsf{M}$ and let $\mathbf{Adv}^{\mathrm{dec}}_{\mathsf{IG},\mathsf{R},\mathsf{M}}(\lambda) = \Pr[\mathrm{DEC}_{\mathsf{IG},\mathsf{R},\mathsf{M}}(\lambda)]$. We say that algorithm M decides R if for every PT IG the function $\mathbf{Adv}^{\mathrm{dec}}_{\mathsf{IG},\mathsf{R},\mathsf{M}}(\cdot)$ is negligible. We say that R is trivial if there is a PT algorithm M that decides R. Intuitively, in game DEC, think of IG as an adversary trying to make M fail. The game returns true when IG succeeds, meaning M returns the wrong decision. A technical point is that if IG generates a positive instance x, the game forces it to lose if the witness w is not valid. Thus we are asking that M is able to decide membership in PT for instances that can be efficiently generated with valid witnesses if the instance is positive. But this does not mean it can decide membership on all instances. Thus if $L(\mathsf{R}) \in \mathbf{BPP}$ then R is certainly trivial, but the converse need not be true.

Result. We show that ZK and subversion soundness (S-SND) cannot co-exist, meaning only trivial relations will have NI systems with both attributes. We stress that we are not asking here for subversion ZK but just plain ZK.

Theorem 1. *Let Π be a NI system satisfying zero knowledge (ZK) and subversion soundness (S-SND) for an* **NP** *relation R. Then R is trivial.*

The proof follows the basic paradigm of GO [36]. We use the simulator to build a cheating prover that violates soundness. In our case this works if soundness holds relative to a simulated CRS, but S-SND guarantees this.

Proof. (Theorem 1). Define the following decision procedure M:

$$\begin{array}{l}\underline{\text{Algorithm } \mathsf{M}(1^\lambda, x)} \\ (crs_0, std_0) \leftarrow\!\!\$\ \Pi.\mathsf{Sim.crs}(1^\lambda)\ ;\ \pi \leftarrow\!\!\$\ \Pi.\mathsf{Sim.pf}(1^\lambda, crs_0, std_0, x) \\ \text{Return } \Pi.\mathsf{V}(1^\lambda, crs_0, x, \pi)\end{array}$$

Thus, to decide if $x \in L(\mathsf{R})$, algorithm M runs the simulator to get a simulated CRS and simulation trapdoor, uses the latter to generate a simulated proof, and decides that $x \in L(\mathsf{R})$ if this proof is valid. Let IG be any PT instance generator. We will show below that $\mathbf{Adv}^{\mathrm{dec}}_{\mathsf{IG},\mathsf{R},\mathsf{M}}(\cdot)$ is negligible. This shows that R is trivial.

To show $\mathbf{Adv}^{\mathrm{dec}}_{\mathsf{IG},\mathsf{R},\mathsf{M}}(\cdot)$ is negligible, below we will define PT adversaries A, B such that

$$\mathbf{Adv}^{\mathrm{dec}}_{\mathsf{IG},\mathsf{R},\mathsf{M}}(\lambda) \leq \mathbf{Adv}^{\mathrm{zk}}_{\Pi,\mathsf{R},\mathsf{A}}(\lambda) + \mathbf{Adv}^{\text{s-snd}}_{\Pi,\mathsf{R},\mathsf{B}}(\lambda) \tag{1}$$

for all $\lambda \in \mathbb{N}$. By assumption, Π satisfies ZK and S-SND for R, so the functions $\mathbf{Adv}^{\mathrm{zk}}_{\Pi,\mathsf{R},\mathsf{A}}(\cdot)$ and $\mathbf{Adv}^{\text{s-snd}}_{\Pi,\mathsf{R},\mathsf{B}}(\cdot)$ are both negligible. Thus Eq. (1) implies that $\mathbf{Adv}^{\mathrm{dec}}_{\mathsf{IG},\mathsf{R},\mathsf{M}}(\cdot)$ is negligible, as desired.

Consider games $\mathrm{G}_0, \mathrm{G}_1, \mathrm{G}_2$ of Fig. 4. Game G_0 is defined ignoring the box, while game $\boxed{\mathrm{G}_1}$ includes it. Games G_0 and G_1 split up the decision process depending on whether or not $x \in L(\mathsf{R})$. Game G_2 switches to a real CRS and proofs, which it can do since the instance generator provided a witness.

Game DEC returns true iff $((x \notin L(\mathsf{R}))$ AND $(d_0 = \mathsf{true}))$ OR $((x \in L(\mathsf{R}))$ AND $(d_1 = \mathsf{true})$ AND $(d_0 = \mathsf{false}))$. The first condition in the OR is when game

GAMES G_0, G_1	GAME G_2
$(x, w) \leftarrow\!\!\$\ \mathsf{IG}(1^\lambda)$; $d_1 \leftarrow \mathsf{R}(x, w)$	$(x, w) \leftarrow\!\!\$\ \mathsf{IG}(1^\lambda)$; $d_1 \leftarrow \mathsf{R}(x, w)$
$(crs, std) \leftarrow\!\!\$\ \Pi.\mathsf{Sim.crs}(1^\lambda)$	$crs \leftarrow\!\!\$\ \Pi.\mathsf{Pg}(1^\lambda)$
$\pi \leftarrow\!\!\$\ \Pi.\mathsf{Sim.pf}(1^\lambda, crs, std, x)$	$\pi \leftarrow\!\!\$\ \Pi.\mathsf{P}(1^\lambda, crs, x, w)$
$d_0 \leftarrow \Pi.\mathsf{V}(1^\lambda, crs, x, \pi)$	$d_0 \leftarrow \Pi.\mathsf{V}(1^\lambda, crs, x, \pi)$
$b \leftarrow ((x \notin L(\mathsf{R})) \wedge (d_0 = \mathsf{true}))$	$b \leftarrow ((d_1 = \mathsf{true}) \wedge (d_0 = \mathsf{false}))$
$\boxed{b \leftarrow ((d_1 = \mathsf{true}) \wedge (d_0 = \mathsf{false}))}$	Return b
Return b	

Fig. 4. Games for proof of Theorem 1

G_0 returns true. The second condition in the OR is equivalent to ((d_1 = true) AND (d_0 = false)), which is the condition under which game G_1 returns true. Furthermore the conditions are mutually exclusive. We thus have

$$\mathbf{Adv}^{\mathrm{dec}}_{\mathsf{IG},\mathsf{R},\mathsf{M}}(\lambda) = \Pr[G_0] + \Pr[G_1] = \Pr[G_0] + \Pr[G_2] + (\Pr[G_1] - \Pr[G_2]) \quad (2)$$

Notice that by completeness of Π we have

$$\Pr[G_2] = 0\ . \quad (3)$$

Now we specify the adversaries A, B as follows:

Adversary $\mathsf{A}^{\mathrm{PROVE}}(1^\lambda, crs)$	Adversary $\mathsf{B}(1^\lambda)$
$(x, w) \leftarrow\!\!\$\ \mathsf{IG}(1^\lambda)$; $d_1 \leftarrow \mathsf{R}(x, w)$	$(x, w) \leftarrow\!\!\$\ \mathsf{IG}(1^\lambda)$
$\pi \leftarrow\!\!\$\ \mathrm{PROVE}(x, w)$; $d_0 \leftarrow \Pi.\mathsf{V}(1^\lambda, crs, x, \pi)$	$(crs, std) \leftarrow\!\!\$\ \Pi.\mathsf{Sim.crs}(1^\lambda)$
If $((d_1 = \mathsf{true}) \wedge (d_0 = \mathsf{false}))$ then $b' \leftarrow 0$	$\pi \leftarrow\!\!\$\ \Pi.\mathsf{Sim.pf}(1^\lambda, crs, std, x)$
Else $b' \leftarrow 1$	Return (crs, x, π)
Return b'	

Then we have

$$\Pr[G_0] \le \mathbf{Adv}^{\text{s-snd}}_{\Pi,\mathsf{R},\mathsf{B}}(\lambda) \quad (4)$$

$$\Pr[G_1] - \Pr[G_2] \le \mathbf{Adv}^{\mathrm{zk}}_{\Pi,\mathsf{R},\mathsf{A}}(\lambda)\ . \quad (5)$$

Putting together Eqs. (2), (3), (4) and (5) we get Eq. (1). □

6 Positive Results

We already have NI systems that are SND+ZK, or SND+WI. We ask, if the CRS is subverted, what more can we have without losing the initial properties? Can we add S-ZK? In Sect. 6.1 we answer positively to this question (result **P1**), showing a protocol that is SND+S-ZK under a knowledge-of-exponent assumption (KEA) in a group equipped with a bilinear map. In light of negative result **N**, this is the best we can achieve if we want to retain ZK in presence of CRS subversion.

Can we add S-SND? In light of **N**, we know that we cannot have S-SND and any form of ZK together. The best we can achieve while retaining S-SND is S-WI. In Sect. 6.2 we show that there exist NI systems that are S-SND+S-WI (result **P2**).

Result **P1** provides S-ZK but requires KEA. A natural question is, if we relax the requirement of S-ZK and aim to retain S-WI, can we achieve it from weaker assumptions? In Sect. 6.3 we show that there exists a NI system that is SND, ZK and S-WI under the weaker assumption that one-way functions and zaps exist.

6.1 Soundness and Subversion ZK

Overview. To achieve S-ZK, a simulator must be able to simulate proofs under a CRS output by a subvertor. As opposed to ZK, the simulator thus cannot embed a trapdoor in the CRS, nor can it extract one from the subvertor by rewinding, as there is no interaction with it. We will instead rely on a knowledge assumption, stating that an algorithm can only produce a certain output if it knows underlying information. This is formalized by requiring that there exists an extractor that extracts the information from the algorithm. We will use this information as the simulation trapdoor, which we can extract from a subvertor outputting a CRS. For soundness, a minimal requirement is that it is hard for the adversary to obtain the trapdoor from an honestly generated CRS.

The knowledge-of-exponent assumption (KEA) for a group $\mathbb{G}$, generated by g, states that from any algorithm which given a random element $h \leftarrow\!\!\$\ \mathbb{G}$ returns a pair of the form (g^s, h^s) one can efficiently extract s. A possible approach for a NI system is to define the CRS as a pair (g^s, h^s), for random s, and define a proof for $x \in L$ to prove that either $x \in L$ or one knows the value s in the CRS. By extracting s, the simulator in the S-ZK game can simulate proofs, while the adversary in the soundness game must supposedly use a witness for x, since it does not know s.

There are two problems with this approach: who chooses the group $\mathbb{G}$ and who chooses the element h used to prove knowledge of s? We address the first problem by letting the group $\mathbb{G}$ be part of the scheme specification. As for the choice of h, it cannot be chosen at CRS setup, since if the subvertor knows $\eta = \log_g h$, it can produce a CRS (S_1, S_2) *without* knowing s by randomly picking $S_1 \leftarrow\!\!\$\ \mathbb{G}$ and setting $S_2 \leftarrow S^\eta$. Fixing h and letting it also be also part of the scheme description is problematic, since again, what guarantees that the subvertor does not know its logarithm and can thereby break KEA? We overcome this issue by defining a new type of KEA, stating that in order to produce elements $(h = g^\eta, g^s, h^s)$, one has to *either* know s *or* η. As tuples of this form are Diffie-Hellman tuples, we call the assumption DH-KEA.

We define a CRS as a tuple $(g^{s_0}, g^{s_1}, g^{s_0 s_1})$ and let a proof for a statement x prove that either there is a witness for x or one knows s_0 or s_1. We prove knowledge by adding a ciphertext C and use a perfectly sound witness-indistinguishable NI proof ζ with trivial CRS (a.k.a. a non-interactive zap) to prove that either $x \in L$ or C encrypts s_0 or s_1. (Using linear encryption for C

and the NI system by GOS [42], both IND-CPA of C, as well as WI of ζ, follow from the decision-linear assumption (Dlin) [18].)

The sketched scheme is ZK since by encrypting the trapdoor s_0 (or s_1) proofs can be simulated, and by IND-CPA of C and WI of ζ they are indistinguishable from real ones. But we defined the CRS to allow even more: by DH-KEA, from a CRS subvertor we can *extract* either s_0 or s_1, which should yield S-ZK. Not quite, since the subvertor could simply output random group elements (S_0, S_1, S_2), from which we cannot extract. Since the GOS NI system requires a *bilinear* group, we can use its pairing to check CRS well-formedness. The prove (and verification) algorithm can then reject a malformed CRS, which together with simulatability under a well-formed CRS yields S-ZK.

Soundness intuitively holds because, by soundness of ζ, a proof for a wrong statement must contain an encryption of s_0 or s_1, which should be infeasible to obtain from an honestly generated CRS if computing discrete logarithms (DL) is hard. (Given a DL challenge S, one can randomly set S_0 or S_1 to S and with probability $\frac{1}{2}$, the proof contains an encryption of $\log S$.) To formally prove soundness, the reduction must recover s from C. We could include in the CRS a public key under which C is to be encrypted: the reduction sets up the CRS, knows the decryption key and can obtain s. Alas, this would break S-ZK: an adversary that created the CRS could also decrypt C and thereby distinguish real proofs from simulated ones.

We therefore include the linear-encryption key $pk = (g^u, g^v)$ in the proof rather than the CRS. But how would the soundness reduction then retrieve s? Could we use KEA again? Since we can only extract one of two possible logarithms, we do the following. The proof contains *two* public keys $pk_0 = (g^{u_0}, g^{v_0})$ and $pk_1 = (g^{u_1}, g^{v_1})$ and s is encrypted under both of them. Additionally, the proof contains elements $g^{u_0u_1}, g^{u_0v_1}, g^{v_0u_1}, g^{v_0v_1}$, whose consistency can be verified via the pairing. By DH-KEA, there exists an extractor which from $(g^{u_0}, g^{u_1}, g^{u_0u_1})$ extracts either u_0 or u_1, another extractor that from $(g^{u_0}, g^{v_1}, g^{u_0v_1})$ extracts u_0 or v_1, and so on. Together these four extractors either yield (u_0, v_0) or (u_1, v_1), thus one of the secret keys corresponding to pk_0 and pk_1. This way the soundness reduction can extract the value s encrypted in a proof for a false statement. At the same time we show that S-ZK still holds.

In our actual scheme we use the CDH assumption (defined below and implied by DLin) instead of DL. The reason is that CDH solutions are group elements, which can be efficiently encrypted using linear encryption. The trapdoor is then a solution to a CDH instance in the CRS. Besides 14 group elements, the most costly component of our proofs is the GOS NI proof ζ. It uses a circuit representation of the NP relation R and shows that (a) either $\mathsf{R}(x, w)$ for some w, or (b) the simulation trapdoor was encrypted (see Eq. (6)). The GOS system [42] was further developed by Groth and Sahai [44] yielding very efficient proofs for algebraic statements, and we could replace GOS by GS. As the clause (b) that we added has precisely this algebraic form, the overhead for turning a proof that is merely WI into one that is S-ZK would be quite modest.

Discussion. Our scheme specification includes the bilinear group, so one might ask whether we have not just shifted the subversion risk from the CRS to the choice of the group. Since the group generation algorithm is deterministic and public, anyone can run the algorithm to re-obtain the group; moreover, different entities can implement it independently if they think that some standardized implementation was subverted, as a check. With the CRS, the situation is different. There is no easy way to check that it was properly generated, at least without compromising security. Perhaps a vocabulary that speaks to this is that the group is *reproducible*, whereas the CRS is not. Someone is trusted to produce it and one cannot easily check that they did it honestly.

Still, one must ask whether the algorithms used allow embedding of backdoors. Here we must look at the specific algorithms. Thus, while one could use a bilinear group in which the discrete-log problem is easy, leading to an insecure scheme, we know it is possible to publicly specify good algorithms. The specifications, given for example in research papers, may be used by anyone to re-produce the results of the algorithms with some faith that there are no backdoors, in the case (as here) that these algorithms are deterministic.

$\underline{\text{Game } \mathrm{KE}_{\mathsf{dGG},\mathsf{M},\mathsf{E}}(\lambda)}$
$(p, \mathbb{G}, \mathbb{G}_T, \mathbf{e}, g) \leftarrow \mathsf{dGG}(1^\lambda)$; $h_0, h_1 \leftarrow\!\!\$\ \mathbb{G}$; $r \leftarrow\!\!\$\ \{0,1\}^{\mathsf{M.rl}(\lambda)}$
$(S_0, S_1, S_2) \leftarrow \mathsf{M}(1^\lambda, h_0, h_1; r)$; $s \leftarrow\!\!\$\ \mathsf{E}(1^\lambda, h_0, h_1, r)$
Return $\big(\mathbf{e}(S_0, S_1) = \mathbf{e}(g, S_2)$ and $g^s \neq S_0$ and $g^s \neq S_1\big)$

$\underline{\text{Game } \mathrm{CDH}_{\mathsf{dGG},\mathsf{A}}(\lambda)}$
$(p, \mathbb{G}, \mathbb{G}_T, \mathbf{e}, g) \leftarrow \mathsf{dGG}(1^\lambda)$; $s, t \leftarrow\!\!\$\ \mathbb{Z}_p$; $C \leftarrow\!\!\$\ \mathsf{A}(1^\lambda, g^s, g^t)$
Return $(C = g^{st})$

$\underline{\text{Game } \mathrm{DLin}_{\mathsf{dGG},\mathsf{A}}(\lambda)}$
$b \leftarrow\!\!\$\ \{0,1\}$; $(p, \mathbb{G}, \mathbb{G}_T, \mathbf{e}, g) \leftarrow \mathsf{dGG}(1^\lambda)$
$u, v, s, t, \xi \leftarrow\!\!\$\ \mathbb{Z}_p$; $b' \leftarrow\!\!\$\ \mathsf{A}(1^\lambda, g^u, g^v, g^{us}, g^{vt}, g^{s+t+b\cdot\xi})$
Return $(b = b')$

Fig. 5. Games defining the knowledge-of-exponent assumption, the CDH assumption and the DLin assumption.

Speaking broadly, we cannot (and do not claim to) prevent all possible subversion. This is not possible. Our goal is to put in defenses that make the most obvious paths harder, one of which is subversion of the CRS.

Bilinear groups. Our construction is based on bilinear groups for which we introduce a new type of knowledge-of-exponent assumption. A bilinear-group generator GGen is a PT algorithm that takes input a security parameter 1^λ and outputs a description of a bilinear group $(p, \mathbb{G}, \mathbb{G}_T, \mathbf{e}, g)$, where p is a prime of length λ, $\mathbb{G}$ and $\mathbb{G}_T$ are groups of order p, g generates $\mathbb{G}$ and $\mathbf{e}\colon \mathbb{G} \times \mathbb{G} \to \mathbb{G}_T$ is a bilinear map that is non-degenerate (i.e. $\langle \mathbf{e}(g,g) \rangle = \mathbb{G}_T$).

While in the cryptographic literature bilinear groups are often assumed to be probabilistically generated, real-world pairing-based schemes are defined for groups that are fixed for every λ. We reflect this by defining the group generator as a deterministic PT algorithm dGG. An advantage of doing so is that every entity in the scheme can compute the group from the security parameter and no party must be trusted with generating the group.

KEA. The knowledge-of-exponent assumption (KEA) [8,25,45] in a group $\mathbb{G}$ states that an algorithm M that is given two random generators g, h of $\mathbb{G}$ and outputs (g^c, h^c) must know c. This is formalized by requiring that there exists an extractor for M which when given M's coins outputs c. Generalizations of KEA were used in the bilinear-group setting in [39]. We introduce a new type of KEA in bilinear groups, which we call DH-KEA, where we assume that if M outputs a Diffie-Hellman (DH) tuple g^s, g^t, g^{st} then it must either know s or t. This should also be the case when M is given two additional random generators h_0, h_1. We note that while an adversary may produce one group element without knowing its discrete logarithm by hashing into the elliptic curve [19,20,55], it seems hard to produce a DH tuple without knowing at least one of the logarithms.

Formally, let $\mathbf{Adv}^{\mathrm{ke}}_{\mathsf{dGG},\mathsf{M},\mathsf{E}}(\lambda) = \Pr[\mathrm{KE}_{\mathsf{dGG},\mathsf{M},\mathsf{E}}(\lambda)]$, where game KE is defined in Fig. 5. The DH-KEA assumption holds for dGG if for every PT M there exists a PT E s.t. $\mathbf{Adv}^{\mathrm{ke}}_{\mathsf{dGG},\mathsf{M},\mathsf{E}}(\cdot)$ is negligible.

We note that due to deterministic group generation the assumption does not hold for non-uniform machines M, as their advice for inputs 1^λ could simply be a DH tuple (S_0, S_1, S_2) w.r.t. the group output by $\mathsf{dGG}(1^\lambda)$. However, we follow Goldreich [34] and only consider uniform machines. As a sanity check, we show that DH-KEA holds in the generic-group model. To reflect hashing into elliptic curves, we provide the adversary with an additional generic operation: it can create new group elements without knowing their discrete log. In the full version [6] we show the following.

Theorem 2. *DH-KEA, as defined above, holds in the generic-group model with hashing into the group.*

CDH. The computational Diffie-Hellman assumption in a group $\mathbb{G}$ states that given g^s and g^t for a random s, t, it should be hard to compute g^{st}. Formally, the CDH assumption holds for dGG if $\mathbf{Adv}^{\mathrm{cdh}}_{\mathsf{dGG},\mathsf{A}}(\cdot)$ is negligible for all PT adversaries A, where $\mathbf{Adv}^{\mathrm{cdh}}_{\mathsf{dGG},\mathsf{A}}(\lambda) = \Pr[\mathrm{CDH}_{\mathsf{dGG},\mathsf{A}}(\lambda)]$ and game CDH is specified in Fig. 5.

DLin. The decision linear (DLIN) assumption [18] in a group $\mathbb{G}$ states that given $(g^u, g^v, g^{us}, g^{vt})$ for random u, v, s, t, the element g^{s+t} is indistinguishable from a random group element. Formally, the DLin assumption holds for dGG if $\mathbf{Adv}^{\mathrm{dlin}}_{\mathsf{dGG},\mathsf{A}}(\cdot)$ is negligible for all PT adversaries A, where $\mathbf{Adv}^{\mathrm{dlin}}_{\mathsf{dGG},\mathsf{A}}(\lambda) = 2\Pr[\mathrm{DLin}_{\mathsf{dGG},\mathsf{A}}(\lambda)] - 1$ and game DLin is defined in Fig. 5.

We will make use of the fact that DLin is self-reducible. This means that given a tuple (U, V, S, T, X) one can produce a new tuple (U', V', S', T', X') so that if the original tuple was linear then the new tuple is so too, but with fresh u, v, s and t; and if X is random then (U', V', S', T', X') are all independently

random as well. In particular, consider the following algorithm that takes input a DLin challenge $(U, V, S, T, X) \in \mathbb{G}^5$:

$$\underline{\text{Algorithm } \mathsf{Rnd}(1^\lambda, (U, V, S, T, X))}$$
$$(p, \mathbb{G}, \mathbb{G}_T, \mathbf{e}, g) \leftarrow \mathsf{dGG}(1^\lambda)\ ;\ z, a, b, c, d \leftarrow_{\$} \mathbb{Z}_p$$
$$U' \leftarrow U^c\ ;\ V' \leftarrow V^d\ ;\ S' \leftarrow S^{cz} U^{ca}\ ;\ T' \leftarrow T^{dz} V^{db}\ ;\ X' \leftarrow X^z g^a g^b$$
$$\text{Return } (U', V', S', T', X')$$

Let s, t, ξ be such that $S = U^s, T = V^t, X = g^\xi$. Define $s' := sz + a$ and $t' := tz + b$ and note that they are both uniformly random. We have $S' = (U')^{s'}$, $T' = (V')^{t'}$ and $X' = g^{\xi z + a + b} = g^{(\xi - s - t)z + sz + tz + a + b} = g^{(\xi - s - t)z + s' + t'}$. Thus, if the original challenge was a linear tuple (i.e., $\xi = s + t$) then the new tuple is also linear with new randomness uc, vd, s', t', whereas otherwise (i.e., $\xi - s - t \neq 0$) U', V', S', T' and X' are independently random.

THE SCHEME. Our S-ZK scheme is based on a bilinear-group generator dGG, for which we define *linear commitments* to messages $M \in \mathbb{G}$ as follows:

$$\underline{\mathsf{Ln.C}(M; (\boldsymbol{u}, \boldsymbol{t}))}$$
$$\boldsymbol{C} \leftarrow (g^{u_0}, g^{u_1}, g^{u_0 t_0}, g^{u_1 t_1}, g^{t_0 + t_1} \cdot M)$$
$$\text{Return } \boldsymbol{C}$$

$$\underline{\mathsf{Ln.D}(\boldsymbol{u}, (C_2, C_3, C_4))}$$
$$M \leftarrow C_4 \cdot C_2^{-1/u_0} \cdot C_3^{-1/u_1}$$
$$\text{Return } M$$

Commitments are hiding under DLin. Since (C_2, C_3, C_4) is a linear encryption under public key (C_0, C_1), the logarithms of the latter let one recover the message via $\mathsf{Ln.D}$.

We also use a statistically sound NI system with trivial CRS (also called "non-interactive zap" by GOS [42]) $\mathsf{Z} = (\mathsf{Z.P}, \mathsf{Z.V})$ for the following relation:

$$\underline{\mathsf{R}_Z((x, S_0, S_1, h, \boldsymbol{C}_0, \boldsymbol{C}_1), ((w, (s, \boldsymbol{u}_0, \boldsymbol{u}_1, \boldsymbol{t}_0, \boldsymbol{t}_1)))}$$
$$\text{If } \mathsf{R}(x, w) = \mathsf{true} \text{ then return } \mathsf{true}$$
$$\text{If } (g^s = S_0 \text{ or } g^s = S_1) \text{ and } \boldsymbol{C}_0 = \mathsf{Ln.C}(h^s; (\boldsymbol{u}_0, \boldsymbol{t}_0)) \text{ and } \boldsymbol{C}_1 = \mathsf{Ln.C}(h^s; (\boldsymbol{u}_1, \boldsymbol{t}_1))$$
$$\text{then return } \mathsf{true} \tag{6}$$
$$\text{Return } \mathsf{false}$$

The NI proof system Z can for example be instantiated by the construction from [42], which does not require a CRS, is perfectly sound and WI under the DLin assumption. Our NIZK system $\Pi[\mathsf{R}, \mathsf{dGG}]$ is given in Fig. 6.

Theorem 3. *Let* R *be an* **NP** *relation and let* dGG *be a bilinear-group generator. Then* $\Pi[\mathsf{R}, \mathsf{dGG}]$*, defined in Fig. 6, satisfies (1) soundness under DH-KEA and CDH; and (2) subversion zero knowledge under DH-KEA and DLin.*

Below we give some intuition. A proof can be found in the full version [6].

$\underline{\Pi.\mathsf{Pg}(1^\lambda)}$
$(p, \mathbb{G}, \mathbb{G}_T, \mathbf{e}, g) \leftarrow \mathsf{dGG}(1^\lambda)$; $t, s_0, s_1 \leftarrow\!\!\$\ \mathbb{Z}_p$; $h \leftarrow g^t$
$S_0 \leftarrow g^{s_0}$; $S_1 \leftarrow g^{s_1}$; $S_2 \leftarrow g^{s_0 s_1}$; $crs \leftarrow (S_0, S_1, S_2, h)$; Return crs

$\underline{\Pi.\mathsf{V}\big(1^\lambda, (S_0, S_1, S_2, h), x, \pi\big)}$
$(p, \mathbb{G}, \mathbb{G}_T, \mathbf{e}, g) \leftarrow \mathsf{dGG}(1^\lambda)$; $(\boldsymbol{C}_0, \boldsymbol{C}_1, \boldsymbol{D}_0, \boldsymbol{D}_1, \zeta) \leftarrow \pi$
If $\mathbf{e}(S_0, S_1) \neq \mathbf{e}(g, S_2)$ then return false
For $i, j = 0, 1$ do
  If $\mathbf{e}(C_{0,i}, C_{1,j}) \neq \mathbf{e}(g, D_{i,j})$ then return false
Return $\mathsf{Z.V}((x, S_0, S_1, h, \boldsymbol{C}_0, \boldsymbol{C}_1), \zeta)$

$\underline{\Pi.\mathsf{P}(1^\lambda, (S_0, S_1, S_2, h), x, w)}$
If $\mathsf{R}(x, w) = \mathsf{false}$ then return $\perp$
$(p, \mathbb{G}, \mathbb{G}_T, \mathbf{e}, g) \leftarrow \mathsf{dGG}(1^\lambda)$
If $\mathbf{e}(S_0, S_1) \neq \mathbf{e}(g, S_2)$ then return $\perp$
$C_{0,0}, \ldots, C_{0,4}, C_{1,2}, C_{1,3}, C_{1,4} \leftarrow\!\!\$\ \mathbb{G}$; $u_0, u_1 \leftarrow\!\!\$\ \mathbb{Z}_p$; $C_{1,0} \leftarrow g^{u_0}$; $C_{1,1} \leftarrow g^{u_1}$
For $i, j = 0, 1$ do $D_{i,j} \leftarrow C_{0,i}^{u_j}$
$\zeta \leftarrow\!\!\$\ \mathsf{Z.P}((x, S_0, S_1, h, \boldsymbol{C}_0, \boldsymbol{C}_1), (w, \perp))$; $\pi \leftarrow (\boldsymbol{C}_0, \boldsymbol{C}_1, \boldsymbol{D}_0, \boldsymbol{D}_1, \zeta)$
Return π

Fig. 6. NIZK scheme $\Pi[\mathsf{R}, \mathsf{dGG}]$ satisfying SND and S-ZK

Soundness. Assume an adversary A outputs a proof $\pi = (\boldsymbol{C}_0, \boldsymbol{C}_1, \boldsymbol{D}_0, \boldsymbol{D}_1, \zeta)$ for a false statement. Since there does not exist a witness w, by statistical soundness of the proof ζ, R_Z must return 1 in the second line in Eq. (6), meaning $\boldsymbol{C}_0$ and $\boldsymbol{C}_1$ are commitments to either $h^{\log S_0}$ or $h^{\log S_1}$; intuitively, the adversary has thus broken the CDH assumption either for challenge (S_0, h) or (S_1, h).

To make this formal, we construct an algorithm B that on input (g^s, h) outputs h^s with probability close to $\frac{1}{2}$. We first construct four machines $\mathsf{M}_{i,j}$, $0 \leq i, j \leq 1$ that are given given (S, h), set $S_b \leftarrow S$ for a random b, complete this to a CRS, on which they run A; when A returns π, $\mathsf{M}_{i,j}$ outputs $(C_{0,i}, C_{1,j}, D_{i,j})$. By DH-KEA there exist four extractors $\mathsf{E}_{i,j}$ which on input (S, h) and $\mathsf{M}_{i,j}$'s coins (which include A's coins) return either $u_{0,i} = \log C_{0,i}$ or $u_{1,j} = \log C_{1,j}$.

Using $\mathsf{M}_{0,0}, \mathsf{M}_{0,1}, \mathsf{M}_{1,0}, \mathsf{M}_{1,1}$, we define B: given a CDH challenge (S, h), it picks coins $\bar{r}$ and uses $\bar{r}$ to pick $b \leftarrow\!\!\$\ \{0, 1\}$, $s' \leftarrow\!\!\$\ \mathbb{Z}_p$ and coins r for A; it sets $S_b \leftarrow S$, $S_{1-b} \leftarrow g^{s'}$ and $S_2 \leftarrow S^{s'}$ and runs A on input (S_0, S_1, S_2, h) and coins r to get π containing $(\boldsymbol{C}_0, \boldsymbol{C}_1, \boldsymbol{D}_0, \boldsymbol{D}_1)$; it then runs all $\mathsf{E}_{i,j}$ on input $(S, h, \bar{r})$, which each returns either $u_{0,i} = \log C_{0,i}$ or $u_{1,j} = \log C_{1,j}$. This implies that for some i, B obtains both $u_{i,0}$ and $u_{i,1}$. Using this, B recovers $T \leftarrow \mathsf{Ln.D}((u_{i,0}, u_{i,1}), (C_{i,2}, C_{i,3}, C_{i,4}))$, which it outputs. By soundness of ζ, we have either $T = h^{\log S_0}$ or $T = h^{\log S_1}$. Since A has no information on where the challenge S was embedded, B solves CDH with probability $\frac{1}{2}$.

$\underline{\Pi.\mathsf{Pg}(1^\lambda)}$
$\sigma \leftarrow\!\!\$\ \{0,1\}^{2\lambda}$; $m_1 \leftarrow\!\!\$\ \mathsf{Z.V}(1^\lambda)$; Return $crs \leftarrow (\sigma, m_1)$

$\underline{\Pi.\mathsf{P}(1^\lambda, (\sigma, m_1), x, w)}$
$m_2 \leftarrow\!\!\$\ \mathsf{Z.P}(1^\lambda, (\sigma, x), (\bot, w), m_1)$; $\pi \leftarrow m_2$; Return π

$\underline{\Pi.\mathsf{V}(1^\lambda, (\sigma, m_1), x, \pi)}$
Return $\mathsf{Z.D}(1^\lambda, (\sigma, x), m_1, \pi)$

Fig. 7. NIZK scheme $\Pi[\mathsf{R}, \mathsf{dGG}]$ satisfying SND and S-ZK

Subversion zero knowledge. By DH-KEA, for every X that outputs a CRS of the form $(g^{s_0}, g^{s_1}, g^{s_0 s_1}, h)$ there exists an algorithm E that extracts either s_0 or s_1. To show S-ZK we first construct a simulator S. Its first part $\mathsf{S.crs}$ picks r, runs $crs \leftarrow \mathsf{X}(1^\lambda, r)$ and sets $s \leftarrow\!\!\$\ \mathsf{E}(1^\lambda, r)$ if crs is correctly formed and $s \leftarrow \bot$ otherwise, and outputs crs, r and the trapdoor $std \leftarrow s$. It is immediate that crs_1 output by X on coins r_1 is indistinguishable from crs_0, r_0 output by $\mathsf{S.crs}$.

We next construct a proof simulator $\mathsf{S.pf}$ for statements x under $crs = (S_0, S_1, S_2, h)$ using trapdoor s. Like $\Pi.\mathsf{P}$ it returns $\bot$ if crs is malformed. Else, it chooses $\boldsymbol{u}_0, \boldsymbol{t}_0, \boldsymbol{u}_1, \boldsymbol{t}_1$ and defines $\boldsymbol{C}_0$ and $\boldsymbol{C}_1$ as commitments to h^s and computes the corresponding elements $D_{i,j} \leftarrow g^{u_{0,i} u_{1,j}}$. Since either $g^s = S_0$ or $g^s = S_1$, $\mathsf{S.pf}$ has thus a witness for the statement $(x, S_0, S_1, h, \boldsymbol{C}_0, \boldsymbol{C}_1) \in \mathsf{R}_Z$, which it uses to compute a proof ζ. The simulated proof is $\pi \leftarrow (\boldsymbol{C}_0, \boldsymbol{C}_1, \boldsymbol{D}_0, \boldsymbol{D}_1, \zeta)$, which we now argue is indistinguishable from a real proof output by $\Pi.\mathsf{P}$ under DLin by a series of game hops.

We first note that when constructing ζ, instead of witness $(s, \boldsymbol{u}_0, \boldsymbol{u}_1, \boldsymbol{t}_0, \boldsymbol{t}_1)$ we could use w; this is indistinguishable under WI, which for the GOS system follows from DLin. In the next game hop, we replace $\boldsymbol{C}_0$ by a random quintuple and construct the $D_{i,j}$'s as in $\Pi.\mathsf{P}$; this is indistinguishable under DLin. In the final game hop we replace $\boldsymbol{C}_1$ by a random quintuple. This is also reduced to DLin using the fact that we can compute the $D_{i,j}$'s using the logarithms of $\boldsymbol{C}_0$. The result is a proof π that is distributed like one output by $\Pi.\mathsf{P}$.

6.2 Subversion SND and Subversion WI

In this section we prove result **P2**: there exists an NI system that is simultaneously SND, WI, S-SND and S-WI. We call Π an NI system with *trivial* CRS if $crs = \varepsilon$ and $\Pi.\mathsf{P}$ and $\Pi.\mathsf{V}$ ignore input crs. In Lemma 4 we observe that if such a Π is SND and WI then it is also S-SND and S-WI. (Intuitively, if the CRS is ignored then there's no harm in subverting it.) In Theorem 5 we then notice that an NI system with trivial CRS exists [42] which is SND and WI under the DLin assumption in bilinear groups (defined on p. 19). As in this instantiation the group is chosen by the prover (rather than fixed as for **P1**), it needs to be *verifiable* [42] (that is, one can efficiently check that it is a bilinear group).

Lemma 4. *Let* R *be an* **NP** *relation. Let* Π *be an NI system with trivial CRS for* R*. If* Π *is SND and WI then it is also S-SND and S-WI.*

Proof. Let A be an S-SND adversary. Define B against SND: on input $(1^\lambda, \varepsilon)$, run $(crs, x, \pi) \leftarrow\!\!\$\ \mathsf{A}(1^\lambda)$ and return (x, π). Since $\Pi.\mathsf{V}(1^\lambda, \varepsilon, x, \pi) = \Pi.\mathsf{V}(1^\lambda, crs, x, \pi)$, we have $\Pr[\mathrm{SND}_{\Pi,\mathsf{R},\mathsf{B}}(\lambda)] = \Pr[\text{S-SND}_{\Pi,\mathsf{R},\mathsf{A}}(\lambda)]$. Thus, if Π is SND, it is S-SND.

Let A be a WI adversary. Define B against S-WI: on input $(1^\lambda, \varepsilon)$, run $(crs, st) \leftarrow\!\!\$\ \mathsf{A}(1^\lambda)$; $b' \leftarrow\!\!\$\ \mathsf{A}^{\text{PROVE}}(1^\lambda, crs, st)$ and return b'; forward A's queries to own oracle (this simulates A's oracle since $\Pi.\mathsf{P}(1^\lambda, \varepsilon, x, w_b) = \Pi.\mathsf{P}(1^\lambda, crs, x, w_b)$). We have $\Pr[\mathrm{WI}_{\Pi,\mathsf{R},\mathsf{B}}(\lambda)] = \Pr[\text{S-WI}_{\Pi,\mathsf{R},\mathsf{A}}(\lambda)]$. Thus, if Π is WI, it is S-WI. □

Theorem 5. *Let* R *be an* **NP** *relation. If the decision-linear assumption holds for a verifiable bilinear group then there exists an NI system* Π *for* R *that is S-SND and S-WI.*

Proof. Let Π be the NI system presented in [42]. Π is an NI system with trivial CRS satisfying SND and WI under the DLin assumption. By Lemma 4 it follows that Π is also S-SND and S-WI. □

6.3 Soundness, ZK and Subversion WI

We prove result **P3** by presenting an NI system that is SND, ZK, and S-WI.

ZAPS. A zap [30] for a relation R is a 2-move protocol (cf. Sect. 4.4), where the first move is *public-coin* and is generated *independently of the statement* to be proved. Zaps retain soundness and witness-indistinguishability even if the statements are chosen adaptively after the first move m_1 is fixed. Consequently, the same m_1 can be reused for many proofs. We denote zaps by

$$m_1 \leftarrow\!\!\$\ \mathsf{Z.V}(1^\lambda)\ ;\ m_2 \leftarrow\!\!\$\ \mathsf{Z.P}(1^\lambda, x, w, m_1)\ ;\ b \leftarrow \mathsf{Z.D}(x, m_1, m_2)\ .$$

Dwork and Naor [30] show that zaps can be constructed from any NIZK in the shared random string model. Concretely, zaps can be based on any family of doubly-enhanced trapdoor permutations, when the underlying NIZK is instantiated with the system of FLS [32].

THE SCHEME. The CRS of our scheme consists of a random bit string σ of length 2λ and the first move m_1 of a zap. A proof consists of the second move of the zap for statement (x, σ), proving that either $x \in L$ or s is the pre-image of σ under a PRG G. The formal description of Π follows.

Let $\mathsf{G}\colon \{0,1\}^\lambda \to \{0,1\}^{2\lambda}$ be a pseudorandom generator and let Z be a zap for the following relation R_Z:

$\underline{\mathsf{R}_Z((\sigma, x), (s, w))}$
If $\sigma = \mathsf{G}(s)$ then return true
Return $\mathsf{R}(x, w)$

Then NI system $\Pi[\mathsf{G}, \mathsf{Z}]$ is given in Fig. 7.

Theorem 6. *Let* R *be an* **NP** *relation. Let* G *be a length-doubling function and* Z *be a zap for relation* R_Z*. If* G *is pseudorandom and* Z *is sound and witness-indistinguishable then* $\Pi[\mathsf{G},\mathsf{Z}]$ *is SND, ZK and S-WI.*

Proof. Soundness of Π follows from the soundness of the zap and the fact that the probability that a randomly sampled string σ is in the range of the PRG G is negligible. ZK follows as in [32]: The ZK simulator picks $s \leftarrow\!\!\$\ \{0,1\}^\lambda$, sets the CRS to be $\sigma \leftarrow \mathsf{G}(s)$ and $m_1 \leftarrow\!\!\$\ \mathsf{Z.V}(1^\lambda)$. When the simulator is challenged to prove a theorem x, it has a witness for $(\sigma, x) \in \mathsf{R}_Z$ and can therefore compute $\pi \leftarrow\!\!\$\ \mathsf{Z.P}(1^\lambda, (\sigma, x), (s, \bot), m_1)$. Indistinguishability of the simulated CRS and proofs follows from the pseudorandomness of G and zap-WI (defined below).

To show S-WI, we prove that from an adversary A winning game $\text{S-WI}_{\Pi,\mathsf{R},\mathsf{X},\mathsf{A}}$ we can construct an adversary B winning the WI game of the underlying zap for relation R_Z. We denote this game by $\text{Z-WI}_{\mathsf{Z},\mathsf{R}_Z,\mathsf{B}}$ and define it in Fig. 8. Note that it reflects the stronger notion of WI where the verifier can obtain several proofs, for theorems of her choice, computed using the same first move m_1.

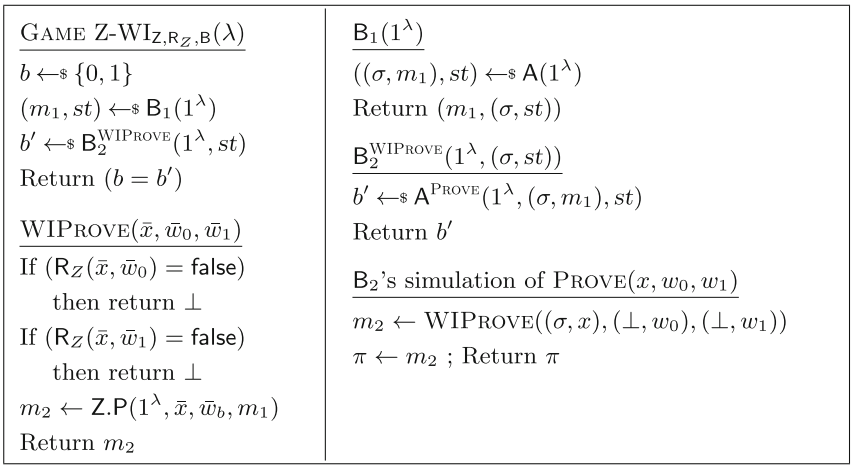

Fig. 8. Game defining WI for zaps (left) and adversary in proof of S-WI of Π

In its first stage B runs A to obtain a CRS consisting of σ and the first message m_1 and returns m_1. B then simulates oracle Prove(x, w_0, w_1) for A by accessing its own oracle WIProve. Figure 8 specifies adversary B. Plugging its description into game $\text{Z-WI}_{\mathsf{Z},\mathsf{R}_Z,\mathsf{B}}$, we obtain

Game Z-WI$_{\mathsf{Z},\mathsf{R}_Z,\mathsf{B}}(\lambda)$
$b \leftarrow\!\!\$\ \{0,1\}$
$((\sigma, m_1), st) \leftarrow\!\!\$\ \mathsf{A}(1^\lambda)$
$b' \leftarrow\!\!\$\ \mathsf{A}^{\text{Prove}}(1^\lambda, (\sigma, m_1), st)$
Return $(b = b')$

Prove(x, w_0, w_1)
If $\mathsf{R}_Z((\sigma, x), (\bot, w_0)) = \mathsf{false}$ then return $\bot$
If $\mathsf{R}_Z((\sigma, x), (\bot, w_1)) = \mathsf{false}$ then return $\bot$
$m_2 \leftarrow \mathsf{Z.P}(1^\lambda, (\sigma, x), (\bot, w_b), m_1)$
Return m_2

As this is precisely the description of game $\text{S-WI}_{\Pi,\mathsf{R},\mathsf{A}}$, we have

$$\Pr[\text{Z-WI}_{\mathsf{Z},\mathsf{R}_Z,\mathsf{B}}(\lambda)] = \Pr[\text{S-WI}_{\Pi,\mathsf{R},\mathsf{A}}(\lambda)] \ . \tag{7}$$

Since Z is zap-WI, $2\Pr[\text{Z-WI}_{\mathsf{Z},\mathsf{R}_Z,\mathsf{B}}(\cdot)] - 1$ is negligible and thus by Eq. (7) $\mathbf{Adv}^{\text{s-wi}}_{\Pi,\mathsf{R},\mathsf{A}}(\cdot)$ is negligible, which proves the theorem. □

Acknowledgments. Bellare was supported in part by NSF grants CNS-1228890 and CNS-1526801, ERC Project ERCC FP7/615074 and a gift from Microsoft corporation. Fuchsbauer was supported in part by the European Research Council under the European Communitys Seventh Framework Programme (FP7/2007-2013 Grant Agreement no. 339563 CryptoCloud). This work was done in part while Bellare and Scafuro were visiting the Simons Institute for the Theory of Computing, supported by the Simons Foundation and by the DIMACS/Simons Collaboration in Cryptography through NSF grant CNS-1523467. We thank Yuval Ishai for helpful discussions and information.

References

1. Abe, M., Fuchsbauer, G., Groth, J., Haralambiev, K., Ohkubo, M.: Structure-preserving signatures and commitments to group elements. In: Rabin, T. (ed.) CRYPTO 2010. LNCS, vol. 6223, pp. 209–236. Springer, Heidelberg (2010). doi:10.1007/978-3-642-14623-7_12
2. Abe, M., Groth, J., Ohkubo, M., Tibouchi, M.: Structure-preserving signatures from type II pairings. In: Garay, J.A., Gennaro, R. (eds.) CRYPTO 2014. LNCS, vol. 8616, pp. 390–407. Springer, Heidelberg (2014). doi:10.1007/978-3-662-44371-2_22
3. Ateniese, G., Magri, B., Venturi, D.: Subversion-resilient signature schemes. In: Ray, I., Li, N., Kruegel, C. (eds.) ACM CCS 15, pp. 364–375. ACM Press, October 2015
4. Barak, B., Lindell, Y., Vadhan, S.P.: Lower bounds for non-black-box zero knowledge. In: 44th FOCS, pp. 384–393. IEEE Computer Society Press, October 2003
5. Barak, B., Pass, R.: On the possibility of one-message weak zero-knowledge. In: Naor, M. (ed.) TCC 2004. LNCS, vol. 2951, pp. 121–132. Springer, Heidelberg (2004). doi:10.1007/978-3-540-24638-1_7
6. Bellare, M., Fuchsbauer, G., Scafuro, A.: NIZKs with an untrusted CRS: Security in the face of parameter subversion. Cryptology ePrint Archive, Report 2016/372 (2016). http://eprint.iacr.org/2016/372
7. Bellare, M., Goldwasser, S.: New paradigms for digital signatures and message authentication based on non-interative zero knowledge proofs. In: Brassard, G. (ed.) CRYPTO 1990. LNCS, vol. 435, pp. 194–211. Springer, Heidelberg (1990)
8. Bellare, M., Palacio, A.: The knowledge-of-exponent assumptions and 3-round zero-knowledge protocols. In: Franklin, M. (ed.) CRYPTO 2004. LNCS, vol. 3152, pp. 273–289. Springer, Heidelberg (2004). doi:10.1007/978-3-540-28628-8_17
9. Bellare, M., Paterson, K.G., Rogaway, P.: Security of symmetric encryption against mass surveillance. In: Garay, J.A., Gennaro, R. (eds.) CRYPTO 2014. LNCS, vol. 8616, pp. 1–19. Springer, Heidelberg (2014). doi:10.1007/978-3-662-44371-2_1
10. Bellare, M., Rogaway, P.: The security of triple encryption and a framework for code-based game-playing proofs. In: Vaudenay, S. (ed.) EUROCRYPT 2006. LNCS, vol. 4004, pp. 409–426. Springer, Heidelberg (2006). doi:10.1007/11761679_25

11. Ben-Sasson, E., Chiesa, A., Green, M., Tromer, E., Virza, M.: Secure sampling of public parameters for succinct zero knowledge proofs. In: 2015 IEEE Symposium on Security and Privacy (SP), pp. 287–304. IEEE (2015)
12. Ben-Sasson, E., Chiesa, A., Tromer, E., Virza, M.: Scalable zero knowledge via cycles of elliptic curves. In: Garay, J.A., Gennaro, R. (eds.) CRYPTO 2014. LNCS, vol. 8617, pp. 276–294. Springer, Heidelberg (2014). doi:10.1007/978-3-662-44381-1_16
13. Ben-Sasson, E., Chiesa, A., Tromer, E., Virza, M.: Succinct non-interactive zero knowledge for a Von Neumann architecture. In: 23rd USENIX Security Symposium (USENIX Security 14), pp. 781–796 (2014)
14. Bernstein, D.J., Lange, T., Niederhagen, R., Dual, E.C.: A standardized back door. Cryptology ePrint Archive, Report 2015/767 (2015). http://eprint.iacr.org/2015/767
15. Bitansky, N., Canetti, R., Paneth, O., Rosen, A.: On the existence of extractable one-way functions. In: Shmoys, D.B. (ed.) 46th ACM STOC, pp. 505–514. ACM Press, May/June 2014
16. Blum, M., De Santis, A., Micali, S., Persiano, G.: Noninteractive zero-knowledge. SIAM J. Comput. **20**(6), 1084–1118 (1991)
17. Blum, M., Feldman, P., Micali, S.: Non-interactive zero-knowledge and its applications (extended abstract). In: 20th ACM STOC, pp. 103–112. ACM Press, May 1988
18. Boneh, D., Boyen, X., Shacham, H.: Short group signatures. In: Franklin, M. (ed.) CRYPTO 2004. LNCS, vol. 3152, pp. 41–55. Springer, Heidelberg (2004). doi:10.1007/978-3-540-28628-8_3
19. Boneh, D., Franklin, M.: Identity-based encryption from the weil pairing. In: Kilian, J. (ed.) CRYPTO 2001. LNCS, vol. 2139, pp. 213–229. Springer, Heidelberg (2001). doi:10.1007/3-540-44647-8_13
20. Brier, E., Coron, J.-S., Icart, T., Madore, D., Randriam, H., Tibouchi, M.: Efficient indifferentiable hashing into ordinary elliptic curves. In: Rabin, T. (ed.) CRYPTO 2010. LNCS, vol. 6223, pp. 237–254. Springer, Heidelberg (2010). doi:10.1007/978-3-642-14623-7_13
21. Canetti, R., Goldreich, O., Goldwasser, S., Micali, S.: Resettable zero-knowledge (extended abstract). In: 32nd ACM STOC, pp. 235–244. ACM Press, May 2000
22. Canetti, R., Pass, R., Shelat, A.: Cryptography from sunspots: how to use an imperfect reference string. In: 48th FOCS, pp. 249–259. IEEE Computer Society Press, October 2007
23. Chase, M., Lysyanskaya, A.: On signatures of knowledge. In: Dwork, C. (ed.) CRYPTO 2006. LNCS, vol. 4117, pp. 78–96. Springer, Heidelberg (2006). doi:10.1007/11818175_5
24. Checkoway, S., Fredrikson, M., Niederhagen, R., Everspaugh, A., Green, M., Lange, T., Ristenpart, T., Bernstein, D.J., Maskiewicz, J., Shacham, H.: On the practical exploitability of Dual EC in TLS implementations. In: USENIX Security (2014)
25. Damgård, I.: Towards practical public key systems secure against chosen ciphertext attacks. In: Feigenbaum, J. (ed.) CRYPTO 1991. LNCS, vol. 576, pp. 445–456. Springer, Heidelberg (1992). doi:10.1007/3-540-46766-1_36
26. Santis, A., Crescenzo, G., Ostrovsky, R., Persiano, G., Sahai, A.: Robust non-interactive zero knowledge. In: Kilian, J. (ed.) CRYPTO 2001. LNCS, vol. 2139, pp. 566–598. Springer, Heidelberg (2001). doi:10.1007/3-540-44647-8_33

27. Dodis, Y., Ganesh, C., Golovnev, A., Juels, A., Ristenpart, T.: A formal treatment of backdoored pseudorandom generators. In: Oswald, E., Fischlin, M. (eds.) EUROCRYPT 2015. LNCS, vol. 9056, pp. 101–126. Springer, Heidelberg (2015). doi:10.1007/978-3-662-46800-5_5
28. Dodis, Y., Haralambiev, K., López-Alt, A., Wichs, D.: Efficient public-key cryptography in the presence of key leakage. In: Abe, M. (ed.) ASIACRYPT 2010. LNCS, vol. 6477, pp. 613–631. Springer, Heidelberg (2010). doi:10.1007/978-3-642-17373-8_35
29. Dolev, D., Dwork, C., Naor, M.: Nonmalleable cryptography. SIAM J. Comput. **30**(2), 391–437 (2000)
30. Dwork, C., Naor, M.: Zaps and their applications. In: 41st FOCS, pp. 283–293. IEEE Computer Society Press, November 2000
31. Escala, A., Groth, J.: Fine-tuning groth-sahai proofs. In: Krawczyk, H. (ed.) PKC 2014. LNCS, vol. 8383, pp. 630–649. Springer, Heidelberg (2014). doi:10.1007/978-3-642-54631-0_36
32. Feige, U., Lapidot, D., Shamir, A.: Multiple non-interactive zero knowledge proofs based on a single random string (extended abstract). In: 31st FOCS, pp. 308–317. IEEE Computer Society Press, October 1990
33. Garg, S., Goyal, V., Jain, A., Sahai, A.: Bringing people of different beliefs together to do UC. In: Ishai, Y. (ed.) TCC 2011. LNCS, vol. 6597, pp. 311–328. Springer, Heidelberg (2011). doi:10.1007/978-3-642-19571-6_19
34. Goldreich, O.: A uniform-complexity treatment of encryption and zero-knowledge. J. Cryptology **6**(1), 21–53 (1993)
35. Goldreich, O., Micali, S., Wigderson, A.: Proofs that yield nothing but their validity or all languages in NP have zero-knowledge proof systems. J. ACM **38**(3), 691–729 (1991)
36. Goldreich, O., Oren, Y.: Definitions and properties of zero-knowledge proof systems. J. Cryptology **7**(1), 1–32 (1994)
37. Goldwasser, S., Micali, S., Rackoff, C.: The knowledge complexity of interactive proof systems. SIAM J. Comput. **18**(1), 186–208 (1989)
38. Groth, J.: Simulation-sound NIZK proofs for a practical language and constant size group signatures. In: Lai, X., Chen, K. (eds.) ASIACRYPT 2006. LNCS, vol. 4284, pp. 444–459. Springer, Heidelberg (2006). doi:10.1007/11935230_29
39. Groth, J.: Short pairing-based non-interactive zero-knowledge arguments. In: Abe, M. (ed.) ASIACRYPT 2010. LNCS, vol. 6477, pp. 321–340. Springer, Heidelberg (2010). doi:10.1007/978-3-642-17373-8_19
40. Groth, J.: Efficient fully structure-preserving signatures for large messages. In: Iwata, T., Cheon, J.H. (eds.) ASIACRYPT 2015. LNCS, vol. 9452, pp. 239–259. Springer, Heidelberg (2015). doi:10.1007/978-3-662-48797-6_11
41. Groth, J., Ostrovsky, R.: Cryptography in the multi-string model. In: Menezes, A. (ed.) CRYPTO 2007. LNCS, vol. 4622, pp. 323–341. Springer, Heidelberg (2007). doi:10.1007/978-3-540-74143-5_18
42. Groth, J., Ostrovsky, R., Sahai, A.: Non-interactive zaps and new techniques for NIZK. In: Dwork, C. (ed.) CRYPTO 2006. LNCS, vol. 4117, pp. 97–111. Springer, Heidelberg (2006). doi:10.1007/11818175_6
43. Groth, J., Ostrovsky, R., Sahai, A.: Perfect non-interactive zero knowledge for NP. In: Vaudenay, S. (ed.) EUROCRYPT 2006. LNCS, vol. 4004, pp. 339–358. Springer, Heidelberg (2006). doi:10.1007/11761679_21
44. Groth, J., Sahai, A.: Efficient non-interactive proof systems for bilinear groups. In: Smart, N. (ed.) EUROCRYPT 2008. LNCS, vol. 4965, pp. 415–432. Springer, Heidelberg (2008). doi:10.1007/978-3-540-78967-3_24

45. Hada, S., Tanaka, T.: On the existence of 3-round zero-knowledge protocols. In: Krawczyk, H. (ed.) CRYPTO 1998. LNCS, vol. 1462, pp. 408–423. Springer, Heidelberg (1998). doi:10.1007/BFb0055744
46. Katz, J., Kiayias, A., Zhou, H.-S., Zikas, V.: Distributing the setup in universally composable multi-party computation. In: Halldórsson, M.M., Dolev, S. (eds.) 33rd ACM PODC, pp. 20–29. ACM, July 2014
47. Kiayias, A., Zacharias, T., Zhang, B.: DEMOS-2: scalable E2E verifiable elections without random oracles. In: Ray, I., Li, N., Kruegel, C. (eds.) ACM CCS 15, pp. 352–363. ACM Press, October 2015
48. Micali, S.: CS proofs (extended abstracts). In: 35th FOCS, pp. 436–453. IEEE Computer Society Press, November 1994
49. Naor, M., Yung, M.: Public-key cryptosystems provably secure against chosen ciphertext attacks. In: 22nd ACM STOC, pp. 427–437. ACM Press, May 1990
50. Pass, R.: On deniability in the common reference string and random oracle model. In: Boneh, D. (ed.) CRYPTO 2003. LNCS, vol. 2729, pp. 316–337. Springer, Heidelberg (2003). doi:10.1007/978-3-540-45146-4_19
51. Patarin, J., Goubin, L.: Asymmetric cryptography with S-Boxes Is it easier than expected to design efficient asymmetric cryptosystems? In: Han, Y., Okamoto, T., Qing, S. (eds.) ICICS 1997. LNCS, vol. 1334, pp. 369–380. Springer, Heidelberg (1997). doi:10.1007/BFb0028492
52. Paterson, K.G.: Imprimitive permutation groups and trapdoors in iterated block ciphers. In: Knudsen, L. (ed.) FSE 1999. LNCS, vol. 1636, pp. 201–214. Springer, Heidelberg (1999). doi:10.1007/3-540-48519-8_15
53. Rijmen, V., Preneel, B.: A family of trapdoor ciphers. In: Biham, E. (ed.) FSE 1997. LNCS, vol. 1267, pp. 139–148. Springer, Heidelberg (1997). doi:10.1007/BFb0052342
54. Russell, A., Tang, Q., Yung, M., Zhou, H.-S.: Cliptography: clipping the power of kleptographic attacks. Cryptology ePrint Archive, Report 2015/695 (2015). http://eprint.iacr.org/2015/695
55. Shallue, A., Woestijne, C.E.: Construction of rational points on elliptic curves over finite fields. In: Hess, F., Pauli, S., Pohst, M. (eds.) ANTS 2006. LNCS, vol. 4076, pp. 510–524. Springer, Heidelberg (2006). doi:10.1007/11792086_36
56. Wee, H.: Lower bounds for non-interactive zero-knowledge. In: Vadhan, S.P. (ed.) TCC 2007. LNCS, vol. 4392, pp. 103–117. Springer, Heidelberg (2007). doi:10.1007/978-3-540-70936-7_6
57. Young, A., Yung, M.: The dark side of "Black-Box" cryptography or: should we trust capstone? In: Koblitz, N. (ed.) CRYPTO 1996. LNCS, vol. 1109, pp. 89–103. Springer, Heidelberg (1996). doi:10.1007/3-540-68697-5_8
58. Young, A., Yung, M.: Kleptography: using cryptography against cryptography. In: Fumy, W. (ed.) EUROCRYPT 1997. LNCS, vol. 1233, pp. 62–74. Springer, Heidelberg (1997). doi:10.1007/3-540-69053-0_6