A Note on Zero-Knowledge Simulator of the CROSS Identification Protocol

Shai Levin

University of Auckland, New Zealand

Abstract. We point out flaw in zero-knowledge of the CROSS identification protocol, CROSS-ID [BBB+24], which allows a distinguisher to distinguish real and simulated transcripts given access to the witness. Moreover, we show that the real and simulated transcripts are not statistically indistinguishable, and therefore the protocol can only satisfy *weak computational* (rather than strong, statistical or perfect) Honest Verifier Zero-knowledge. This issue is still present in version 2.0 updated on January 31, 2025, which resolves the security losses attained via the attacks of [BLP+25]

Keywords: NIST-signatures · Cryptanalysis · Zero-Knowledge · Code-based cryptography

CROSS [BBB+24] is a candidate submitted to the NIST additional digital signatures competition. The CROSS protocol has been selected as a round 2 candidate and is thus under more scrutiny by the cryptographic community. This is particularly the case for CROSS, which has recently been attacked in [BLP+25], leading the authors to update their parameters to account for the loss in bit-security. In this note, we show that the identification protocol which the CROSS signature is based on does not satisfy their included definition of honest verifier zero-knowledge. In particular, we show that (a) the distributions of real and simulated transcripts are not statistically close, and (b) given access to the witness, it is possible to distinguish between the real and simulated transcripts with overwhelming advantage.

It is unclear what relevance this has to the security of the resulting signature scheme, when the Fiat-Shamir transform has been applied. However, this is particularly relevant to the security of fully-anonymous ring signatures constructed via generic transformation of identification protocols, as was seen in the case of SQISign [BLL24].

Note: We assume the reader of this note is accompanied with the CROSS security document [BBB+25]. The relevant content is presented in Section 4.2 and Figure 4. Hence, except for the definitions of zero-knowledge, we will follow all notation and definitions that are present in the document.

Identifying Variants of Honest Verifier Zero Knowledge We consider the various notions of honest verifier zero-knowledge in the context of 5-round indentification protocols, which are extended from the classic definitions of 3-round sigma protocols. Below is the definition present in the v2.0 security document of CROSS, retrieved from [BBB+25]:

Definition 1 (Honest Verifier Zero-Knowledge)**.** Let $\Pi = (\mathcal{P}, \mathcal{V})$ be interactive proof system for a hard relation $\mathcal{R} \subseteq X \times Y$. We say that Π is honest-verifier zero-knowledge if there exists a probabilistic polynomial time algorithm $\mathcal{S}$, called the simulator, such that the following two distribution ensembles are indistinguishable:

$$\{(x, w, \langle \mathcal{P}(x, w), \mathcal{V}(x) \rangle) \mid (x, w) \leftarrow\$ \mathcal{R}\} \text{ and } \{(x, w, \mathcal{S}(x)) \mid (x, w) \leftarrow\$ \mathcal{R}\}$$

E-mail: shai.levin@auckland.ac.nz (Shai Levin)

Date of this document: 2025-02-26.

where $\langle \mathcal{P}(x,w), \mathcal{V}(x) \rangle$ denotes the transcript of an honest interaction between a prover and the verifier with their respective inputs.

There is some ambiguity in the definition. It is unclear what the authors mean by indistinguishability. In the security proof for zero-knowledge, they claim that certain values in the simulated transcript follow the same statistical distribution as those in real transcripts, but typically the witness is only included in the distribution given to the distinguisher in the setting of computational zero-knowledge. Unfortunately, we show that the CROSS protocol can only satisfy the variant below (in the ROM):

Definition 2 (Weak Computational Honest Verifier Zero-knowledge)**.** Given a random oracle O, let $\Pi = (\mathcal{P}^O, \mathcal{V}^O)$ be interactive proof system for a hard relation $\mathcal{R} \subseteq X \times Y$. We say that Π is weak computational honest-verifier zero-knowledge if there exists a probabilistic polynomial time algorithm $\mathcal{S}^O$, called the simulator, such that for any $\lambda \in \mathbb{N}$, and any PPT distinguisher D that makes polynomially many queries to the random oracle O, the following quantity is negligible in λ:

$$\Pr[D^O(x, \langle \mathcal{P}^O(x,w), \mathcal{V}^O(x) \rangle) = 1 \mid (x,w) \leftarrow\!\$\, \mathcal{R}] - \Pr[D^O(x, \mathcal{S}^O(x)) = 1 \mid (x,w) \leftarrow\!\$\, \mathcal{R}]$$

We denote the above quantity as the advantage of D.

In particular, we first provide a distinguisher for $\Pi_{\mathsf{CROSS\text{-}ID}}$ which shows that it does not satisfy Definition 1 (i.e. it is not *strong* honest verifier zero-knowledge). Then, we show that $\Pi_{\mathsf{CROSS\text{-}ID}}$ is not statistically zero-knowledge, since for a fixed instance-witness pair, the distributions of real and simulated transcripts are not statistically close.

An efficient distinguisher given access to the witness We define a distinguisher D for Definition 1 of $\Pi_{\mathsf{CROSS\text{-}ID}}$ as follows. On input:

- instance $(G, \mathbf{H}, \mathbf{s})$ where $G \subseteq \mathbb{E}^n$, $\mathbf{H} \in \mathbb{F}_p^{(n-k)\times n}$, $\mathbf{s} \in \mathbb{F}_p^{n-k}$,
- witness $\mathbf{e} \in G$ such that $\mathbf{s} = \mathbf{e}\mathbf{H}^\top$,
- and transcript $T = (\mathtt{cmt}_0, \mathtt{cmt}_1, \mathtt{chall}_1, \mathtt{digest}_y, \mathtt{chall}_2, \mathtt{resp})$

The distinguisher, given oracle access to $O = \mathsf{Hash}(\cdot)$, performs the following steps:

1. If $\mathtt{chall}_2 = 0$, parse $\mathtt{resp}$ as Seed and perform the following:
   (a) Compute $(\mathbf{e}', \mathbf{u}') \leftarrow \mathsf{CSPRNG}(\mathsf{Seed})$ and hence with knowledge of $\mathbf{e}$, obtain $\mathbf{v} = \mathbf{e} \star (\mathbf{e}')^{-1}$, and $\mathbf{u} = \mathbf{v} \star \mathbf{u}'$. Lastly, compute $\mathbf{s}' = \mathbf{u}\mathbf{H}^\top$.
   (b) If $\mathtt{cmt}_0 = \mathsf{Hash}(\mathbf{s}'|\mathbf{v})$, output 1. Otherwise, output 0.
2. Otherwise, $\mathtt{chall}_2 = 1$, parse $\mathtt{resp}$ as $(\mathbf{y}, \mathbf{v})$ and perform the following:
   (a) With knowledge of $\mathbf{e}$, compute $\mathbf{e}' = \mathbf{e} \star \mathbf{v}^{-1}$ and solve for $\mathbf{u}' = \mathbf{y} - \mathtt{chall}_1 \star \mathbf{e}'$.
   (b) If $\mathtt{cmt}_1 = \mathsf{Hash}(\mathbf{u}'|\mathbf{v})$, output 1. Otherwise, output 0.

The only time the distinguisher will output 1 for a simulated transcript is when the simulator chooses a random bit string that coincides with the commitment for an honest execution of the protocol with the same challenges and responses. This occurs with probability $2^{-2\lambda}$. Hence the distinguisher has overwhelming advantage $1 - 2^{-2\lambda}$ in distinguishing the distributions in Definition 1.

Statistical distance between real and simulated transcripts Recall the statistical distance (or total variable distance) between two distributions (or random variables) X and Y is defined as:

$$\Delta(X,Y) = \frac{1}{2}\sum_{x} |\Pr[X=x] - \Pr[Y=x]|$$

The distributions are said to be statistically indistinguishable for a parameter $\lambda \in \mathbb{N}$ if $\Delta(X,Y) \leq \mathsf{negl}(n)(\lambda)$.

We rely on the following classical results from [Sho05, Thm 8.32, Thm 8.31].

Lemma 1. *If S and T are finite sets, and X and Y are random variables taking values in S, and $f : S \to T$ is a function, then $\Delta(X,Y) \geq \Delta(f(X), f(Y))$.*

Lemma 2. *Let X and Y be random variables taking the values in a set S. For every $S' \subseteq S$, we have $\Delta(X,Y) \geq |\Pr[X \in S'] - \Pr[Y \in S']|$.*

Given $(P^O, V^O) = \Pi_{\mathsf{CROSS\text{-}ID}}$, let us consider the distribution of transcripts for a fixed $(x,w) \in \mathcal{R}_{\mathsf{CROSS\text{-}ID}}$ and instantiation of a random oracle O. Let

$$\mathcal{T}_{\text{real}} = \{\langle P^O(x,w), V^O(x)\rangle\} \text{ and } \mathcal{T}_{\text{sim}} = \{\mathcal{S}^O(x)\}$$

be the random variables associated to real and simulated transcripts respectively over the set of valid transcripts Ω. From now on, parse elements of the set $T \in \Omega$ as $(\mathtt{cmt}_0, \mathtt{cmt}_1, \mathtt{chall}_1, \mathtt{digest}_y, \mathtt{chall}_2, \mathtt{resp})$.

First, we observe that since the real transcripts use a λ-bit seed, which determines the resulting values for $\mathtt{cmt}_0$ and $\mathtt{cmt}_1$. However, the digests for the random oracle are length 2λ. Hence there are at least $2^{2\lambda} - 2^{\lambda}$ possible values for $\mathtt{cmt}_0$ and $\mathtt{cmt}_1 \in \{0,1\}^{2\lambda}$ which are never used in the real transcripts, but are in the unopened commitment of the simulated transcript. Consider the function $f : \Omega \to \{0,1\} \times \{0,1\}^{2\lambda}$ which sends:

$$(\mathtt{cmt}_0, \mathtt{cmt}_1, \mathtt{chall}_1, \mathtt{digest}_y, \mathtt{chall}_2, \mathtt{resp}) \mapsto (\mathtt{chall}_2, \mathtt{cmt}_{1-\mathtt{chall}_2})$$

That is, the distribution of $\mathtt{chall}_2$ and the $(1-\mathtt{chall}_2)$-th commitment, which remains unopened. Let $\mathcal{T}'_{\text{real}} := f(\mathcal{T}_{\text{real}})$ and $\mathcal{T}'_{\text{sim}} := f(\mathcal{T}_{\text{sim}})$, and observe that the latter distribution is uniformly distributed.

Define the function $g_{0,\mathbf{e}} : \{0,1\}^{\lambda} \to \{0,1\}^{2\lambda}$ that takes as input a λ-bit seed, and outputs the resulting commitment $\mathtt{cmt}_0$ in the real transcript for a given witness $\mathbf{e}$. Similarly, define $g_{1,\mathbf{e}}$ to output the commitment $\mathtt{cmt}_1$ in the real transcript for a given seed and $\mathbf{e}$. Now, let

$$\begin{aligned} S_0 &= \{0\} \times \{0,1\}^{2\lambda} \setminus \{0\} \times g_{1,\mathbf{e}}(\{0,1\}^{\lambda}), \\ \text{and } S_1 &= \{1\} \times \{0,1\}^{2\lambda} \setminus \{1\} \times g_{0,\mathbf{e}}(\{0,1\}^{\lambda}) \end{aligned}$$

Which is precisely the set of values which arise in $\mathcal{T}'_{\text{sim}}$ and not in $\mathcal{T}'_{\text{real}}$. We note that

$$|S_0| \geq 2^{2\lambda} - 2^{\lambda} \text{ and } |S_1| \geq 2^{2\lambda} - 2^{\lambda}. \tag{1}$$

Equality is the best case, where $g_{i,\mathbf{e}}$ is injective for both $i \in \{0,1\}$. Hence, we have:

$$\begin{aligned} \Delta(\mathcal{T}_{\text{real}}, \mathcal{T}_{\text{sim}}) &\geq \Delta(\mathcal{T}'_{\text{real}}, \mathcal{T}'_{\text{sim}}) && \text{(By Lemma 1)} \\ &\geq |\Pr[\mathcal{T}'_{\text{real}} \in S_0 \cup S_1] - \Pr[\mathcal{T}'_{\text{sim}} \in S_0 \cup S_1]| && \text{(By Lemma 2)} \\ &= |0 - \Pr[\mathcal{T}'_{\text{sim}} \in S_0 \cup S_1]| \\ &= \Pr[\mathcal{T}'_{\text{sim}} \in S_0 \cup S_1] \\ &= \frac{|S_0 \cup S_1|}{|\{0,1\} \times \{0,1\}^{2\lambda}|} \end{aligned}$$

$$\geq \frac{2^{2\lambda} - 2^{\lambda} + 2^{2\lambda} - 2^{\lambda}}{2 \cdot 2^{2\lambda}} \qquad \text{(By Equation (1))}$$
$$= 1 - 2^{-2\lambda-1}$$

Hence we have that real and protocol transcripts for a fixed instance $(x, w) \in \mathcal{R}_{\mathsf{CROSS\text{-}ID}}$ are far from statistically indistinguishable, with statistical distance at least $1 - 2^{-2\lambda-1} = 1 - \mathsf{negl}(n)(\lambda)$.

Conclusion In short, the problem present in CROSS-ID is that the commitments sent to the verifier are not hiding. Generally, a hash based commitment needs fresh, independent randomness in order to be hiding. We believe this flaw is due to the extreme optimisations of CROSS, in an attempt to minimise their signature sizes.

Acknowledgements

This work is supported by funding from the Ministry for Business, Innovation and Employment, New Zealand. The author would like to thank Prof. Steven Galbraith and his reading group at University of Auckland for their intuition and guidance on this work.

References

[BBB+24] Marco Baldi, Alessandro Barenghi, Sebastian Bitzer, Patrick Karl, Felice Manganiello, Alessio Pavoni, Gerardo Pelosi, Paolo Santini, Jonas Schupp, Freeman Slaughter, Antonia Wachter-Zeh, and Violetta Weger. CROSS — Codes and Restricted Objects Signature Scheme. Technical report, National Institute of Standards and Technology, 2024. available at https://csrc.nist.gov/Projects/pqc-dig-sig/round-2-additional-signatures.

[BBB+25] Marco Baldi, Alessandro Barenghi, Sebastian Bitzer, Patrick Karl, Felice Manganiello, Alessio Pavoni, Gerardo Pelosi, Paolo Santini, Jonas Schupp, Freeman Slaughter, Antonia Wachter-Zeh, and Violetta Weger. CROSS security details document v2.0, 2025. URL: https://www.cross-crypto.com/nist-submission.html.

[BLL24] Giacomo Borin, Yi-Fu Lai, and Antonin Leroux. Erebor and durian: Full anonymous ring signatures from quaternions and isogenies. Cryptology ePrint Archive, Report 2024/1185, 2024. URL: https://eprint.iacr.org/2024/1185.

[BLP+25] Michele Battagliola, Riccardo Longo, Federico Pintore, Edoardo Signorini, and Giovanni Tognolini. A revision of CROSS security: Proofs and attacks for multi-round fiat-shamir signatures. Cryptology ePrint Archive, Paper 2025/127, 2025. URL: https://eprint.iacr.org/2025/127.

[Sho05] Victor Shoup. *A Computational Introduction to Number Theory and Algebra*. Cambridge University Press, 2005.