Axolotl

Robust Asynchronous Threshold Schnorr

Hila Dahari-Garbian [1] Ariel Nof [2] Luke Parker[3]

[1]Reichman University @ Herzliya, Israel
hiladahari41@gmail.com

[2]Bar-Ilan University @ Ramat-Gan, Israel
ariel.nof@biu.ac.il

[3]Serai DEX @ Florida, United States of America
lukeparker@serai.exchange

Presented on 2026-10-07 @ NIST TCPT3

exponent-Verifiable Random Function (eVRF)

- A pseudo-random function is a deterministic function whose output is indistinguishable from random
- A exponent-Verifiable Random Function[1] has participants publish the output in exponent (committed to) with an efficient proof to verify the committed-to value is the correct output of the pseudo-random function
- Instantiations from the Decisional Diffie-Hellman problem over an elliptic curve and the Paillier cryptosystem were provided

[1] ia.cr/2024/397, 10.1007/978-3-031-91098-2_8

Axolotl

- A 'multi-party eVRF' gadget
- Allows an arbitrary threshold to verifiably evaluate a pseudo-random function, revealing functions of the output
- One-round, robust protocol, with the output independent to the subset of participants
- Distinct from summing local evaluations of eVRFs in that:
  - Avoids the Asynchronous Common Subset problem
  - Any subset can reconstruct without additional sharing
- Distinct from replicated-secret-sharing in that the secrets are Shamir secret shares, not a factorial amount of entropy
- Each participant multicasts a message of size $\mathcal{O}(\lambda)$
- Low-overhead prover with complexity $\mathcal{O}(\lambda + t * \log_2(t))$

Background

- We work with the group $\hat{G} : \hat{G}_q \times \hat{H}$
- $\hat{G}_q$ is an unknown-order group
- $\hat{H}$ is a q-order group with an easy discrete-log problem
- For $G \in \hat{G}$, we consider $G + i \cdot H \; \forall \; 0 \leq i < q$ members of the same "coset"
- We have the surjection $\psi : \hat{G} \to \hat{G}_q$ where for $G \in \hat{G}$ and any member C of the same coset, we have $\psi(G) = \psi(C)$
- We have the injection $\phi : \hat{G}_q \to \hat{G}$ where $\phi(\psi(G))$ is in the same coset as G
- The two together form a "coset labeling function"[2] where $\phi(\psi(G))$ chooses a deterministic element from the coset

[2] ia.cr/2022/363, 10.1007/978-3-031-15985-5_15

Pseudo-Random Function

- We denote the evaluation key $k \in \mathbb{Z}$
- Sample $G = \mathcal{H}_{\hat{G}}(\text{context})$
- The partial evaluation is the Diffie-Hellman $K = k \cdot G$
- Output the evaluation $o = \text{dlog}_H(K - \phi(\psi(K)))$
- o is indistinguishable from random in $\mathbb{F}_q$ under the Decisional Diffie-Hellman assumption in $\hat{G}$
- We are able to efficiently prove the correctness of K, deferring outputting the evaluation to be 'in the clear', avoiding having to prove correct application of the surjection, injection

Functional Evaluation

- Participants sample a mask $m \in \mathbb{F}_q$
- Participants reconstruct not K but $K + m \cdot H$
- $\text{dlog}_H(K + m \cdot H - \phi(\psi(K + m \cdot H))) = o + m$
- For a q-order generator E where the discrete-log problem is hard, participants also publish $m \cdot E$
- We then have $o \cdot E = ((m + o) \cdot E) - m \cdot E$
- Similar techniques allow evaluating other functions

Threshold Setup

- We require Shamir-secret-sharing the evaluation key $k \in \mathbb{Z}$
- There is prior art for such a setup[3]
- Because the secret shares do not undergo a modular reduction, the resulting shares are of size $\mathcal{O}(\lambda + t * \log_2(t))$
- Axolotl's prover complexity is equal to a share's size
- In a non-threshold setting, a non-interactive setup with secret shares of size $\mathcal{O}(\lambda)$ is plausible (but not our focus)

[3] ia.cr/2022/1437, 10.1007/978-3-031-38557-5_20

Threshold Schnorr

- Invoke the Axolotl gadget (as above) to learn the nonce commitment $R = r \cdot E$
- Invoke the Axolotl gadget again to output the function $r + c * x$
- Two-round protocol, asynchronous, over peer-to-peer channels
- Supports the optimal corruption threshold (no honest-majority assumption)
- Guarantees delivery of the signature so long as the liveness threshold is satisfied
- Approximately a couple kibibytes multicast by each participant for each round at the 128-bit security level
- Approximately a few seconds with our constant-time prover

Other Notes

- We plan to publish our preprint shortly
- Our concrete proposal uses a class group as seen in the Castagnos-Laguillaumie framework[4] and assumes the hardness of the Decisional Diffie-Hellman and Adaptive Root problems
- We are planning detailed specification of all algorithms, including descriptions, pseudo-code, and time/space complexities
- We wonder if there's an adaptively-secure variant of Axolotl, premised on the same techniques used for threshold BLS signatures[5]

[4] ia.cr/2015/047, 10.1007/978-3-319-16715-2_26
[5] ia.cr/2023/1553, 10.1007/978-3-031-68394-7_9