Student Seminar: Security Protocols and Applications - Spring 2025

Presentation Topics:

#1: Efficient Anonymous Tokens with Private Metadata Bit

Mentor : Laurane Marco (laurane.marco@epfl.ch)

Abstract: We present a cryptographic construction for anonymous tokens with private metadata bit, called PMBTokens. This primitive enables an issuer to provide a user with a lightweight, single-use anonymous trust token that can embed a single private bit, which is accessible only to the party who holds the secret authority key and is private with respect to anyone else. Our construction generalizes and extends the functionality of Privacy Pass (PETS'18) with this private metadata bit capability. It is based on the DDH and CTDH assumptions in the random oracle model and provides unforgeability, unlinkability, and privacy for the metadata bit. Both Privacy Pass and PMBTokens rely on non-interactive zero-knowledge proofs (NIZKs). We present new techniques to remove the need for NIZKs, while still achieving unlinkability. We implement our constructions and we report their efficiency costs.

Authors: *Ben Kreuter and Tancrede Lepoint and Michele Orrù and Mariana Raykova*

Material:

https://eprint.iacr.org/2020/072

#2: Client-Auditable Verifiable Registries

Mentor : Abdullah Talayhan (abdullah.talayhan@epfl.ch)

Abstract: Verifiable registries allow clients to securely access a key-value mapping maintained by an untrusted server. Applications include distribution of public keys, routing information or software binaries. Existing proposals for verifiable registries rely on global invariants being audited whenever the registry is updated. Clients typically rely on trusted third-party auditors, as large registries become expensive to audit.

We propose several new protocols for client-auditable registries that enable efficient verification of many updates to the registry, removing the need for third-party auditors. Our solutions use incrementally-verifiable computation (IVC) and/or RSA accumulators. Our evaluation shows that our constructions meet practical throughput requirements (60
updates / second), which is 100 xfaster than naive solutions using IVC. Clients save 100-10^4x bandwidth and computation costs over prior solutions requiring auditing every update.

Authors: *Nirvan Tyagi and Ben Fisch and Joseph Bonneau and Stefano Tessaro*

Material:

https://eprint.iacr.org/2021/627

#3: Prio: Private, Robust, and Scalable Computation of Aggregate Statistics

Mentor : Abdullah Talayhan (abdullah.talayhan@epfl.ch)

Abstract: This paper presents Prio, a privacy-preserving system for the collection of aggregate statistics. Each Prio client holds a private data value (e.g., its current location), and a small set of servers compute statistical functions over the values of all clients (e.g., the most popular location). As long as at least one server is honest, the Prio servers learn nearly nothing about the clients' private data, except what they can infer from the aggregate statistics that the system computes. To protect functionality in the face of faulty or malicious clients, Prio uses secret-shared non-interactive proofs (SNIPs), a new cryptographic technique that yields a hundred-fold performance improvement over conventional zero-knowledge approaches. Prio extends classic private aggregation techniques to enable the collection of a large class of useful statistics. For example, Prio can perform a least-squares regression on high-dimensional client-provided data without ever seeing the data in the clear.

Authors: *Henry Corrigan-Gibbs, Dan Boneh*

Material:

https://arxiv.org/abs/1703.06255

#4: A Fast and Simple Partially Oblivious PRF, with Applications

Mentor : Laurane Marco (laurane.marco@epfl.ch)

Abstract: We build the first construction of a partially oblivious pseudorandom function (POPRF) that does not rely on bilinear pairings. Our construction can be viewed as combining elements of the 2HashDH OPRF of Jarecki, Kiayias, and Krawczyk with the Dodis-Yampolskiy PRF. We analyze our POPRF's security in the random oracle model via reduction to a new one-more gap strong Diffie-Hellman inversion assumption. The most significant technical challenge is establishing confidence in the new assumption, which requires new proof techniques that enable us to show that its hardness is implied by the q
q-DL assumption in the algebraic group model.

Our new construction is as fast as the current, standards-track OPRF 2HashDH protocol, yet provides a new degree of flexibility useful in a variety of applications. We show how POPRFs can be used to prevent token hoarding attacks against Privacy Pass, reduce key management complexity in the OPAQUE password authenticated key exchange protocol, and ensure stronger security for password breach alerting services.

Authors: *Nirvan Tyagi and Sofía Celi and Thomas Ristenpart and Nick Sullivan and Stefano Tessaro and Christopher A. Wood*

Material:

https://eprint.iacr.org/2021/864

#5: Efficient ECDSA-based Adaptor Signature for Batched Atomic Swaps

Mentor : Abdullah Talayhan (abdullah.talayhan@epfl.ch)

Abstract: Adaptor signature is a novel cryptographic primitive which ties together the signature and the leakage of a secret value. It has become an important tool for solving the scalability and interoperability problems in the blockchain. Aumayr et al. (Asiacrypt 2021) recently provide the formalization of the adaptor signature and present a provably secure ECDSA-based adaptor signature, which requires zero-knowledge proof in the pre-signing phase to ensure the signer works correctly. However, the number of zero-knowledge proofs is linear with the number of participants. In this paper, we propose efficient ECDSA-based adaptor signature schemes and give security proofs based on ECDSA. In our schemes, the zero-knowledge proofs in the pre-signing phase can be generated in a batch and offline. Meanwhile, the online pre-signing algorithm is similar to the ECDSA signing algorithm and can enjoy the same efficiency as ECDSA. In particular, considering specific verification scenarios, such as (batched) atomic swaps, our schemes can reduce the number of zero-knowledge proofs in the pre-signing phase to one, independent of the number of participants. Last, we conduct an experimental evaluation, demonstrating that the performance of our ECDSA-based adaptor signature reduces online pre-signing time by about 60% compared with the state-of-the-art ECDSA-based adaptor signature.

Authors: *Binbin Tu, Min Zhang, and Yu Chen*

Material:

https://eprint.iacr.org/2024/140

#6: LadderLeak: Breaking ECDSA With Less Than One Bit Of Nonce Leakage

Mentor : Abdullah Talayhan (abdullah.talayhan@epfl.ch)

Abstract: Although it is one of the most popular signature schemes today, ECDSA presents a number of implementation pitfalls, in particular due to the very sensitive nature of the random value (known as the nonce) generated as part of the signing algorithm. It is known that any small amount of nonce exposure or nonce bias can in principle lead to a full key recovery: the key recovery is then a particular instance of Boneh and Venkatesan's hidden number problem (HNP). That observation has been practically exploited in many attacks in the literature, taking advantage of implementation defects or side-channel vulnerabilities in various concrete ECDSA implementations. However, most of the attacks so far have relied on at least 2 bits of nonce bias (except for the special case of curves at the 80-bit security level, for which attacks against 1-bit biases are known, albeit with a very high number of required signatures). In this paper, we uncover LadderLeak, a novel class of side-channel vulnerabilities in implementations of the Montgomery ladder used in ECDSA scalar multiplication. The vulnerability is in particular present in several recent versions of OpenSSL. However, it leaks less than 1 bit of information about the nonce, in the sense that it reveals the most significant bit of the nonce, but with probability <1. Exploiting such a mild leakage would be intractable using techniques present in the literature so far. However, we present a number of theoretical improvements of the Fourier analysis approach to solving the HNP (an approach originally due to Bleichenbacher), and this lets us practically break LadderLeak-vulnerable ECDSA implementations instantiated over the sect163r1 and NIST P-192 elliptic curves. In so doing, we achieve several significant computational records in practical attacks against the HNP. We submitted a short Abstract summarizing the work, but a long version was accepted to CCS'20 and a full paper is available on ePrint [1]. The work was already presented at the Crypto & Privacy Village at DEFCON [2] and the Workshop on Attacks in Cryptography [3,4] affiliated to CRYPTO 2020. [1] https://eprint.iacr.org/2020/615 [2] https://cryptovillage.org/dc28/ [3] https://www.youtube.com/watch?v=UbjOKMTVMWQ (long) [4] http://www.youtube.com/watch?v=1ddvx2TgPF8&t=22m09s (short)

Authors: *Diego F. Aranha, Felipe R. Novaes, Akira Takahashi, Mehdi Tibouchi, Yuval Yarom*

Material:

https://youtu.be/jeQvDLPQsuw?t=38

- plus those cited above

#7: DECO: Liberating Web Data Using Decentralized Oracles for TLS

Mentor : Betül Durak (betul.durak@microsoft.com)

Abstract: Thanks to the widespread deployment of TLS, users can access private data over channels with end-to-end confidentiality and in- tegrity. What they cannot do, however, is prove to third parties the provenance of such data, i.e., that it genuinely came from a particu- lar website. Existing approaches either introduce undesirable trust assumptions or require server-side modifications.
Users' private data is thus locked up at its point of origin. Users cannot export data in an integrity-protected way to other applica- tions without help and permission from the current data holder. We propose DECO (short for dec
entralized oracle) to address the above problems. DECO allows users to prove that a piece of data accessed via TLS came from a particular website and option- ally prove statements about such data in zero-knowledge, keeping the data itself secret. DECO is the first such system that works without trusted hardware or server-side modifications.
DECO can liberate private data from centralized web-service silos, making it accessible to a rich spectrum of applications. To demonstrate the power of DECO, we implement three applications that are hard to achieve without it: a private financial instrument using smart contracts, converting legacy credentials to anonymous credentials, and verifiable claims against price discrimination.

Authors: *Fan Zhang, Deepak Maram , Harjasleen Malvai, Steven Goldfeder, Ari Juels*

Material:

https://arxiv.org/pdf/1909.00938.pdf

#8: On Valiant's Conjecture: Impossibility of Incrementally Verifiable Computation from Random Oracles

Mentor : Abdullah Talayhan (abdullah.talayhan@epfl.ch)

Abstract: In his landmark paper at TCC 2008 Paul Valiant introduced the notion of ``incrementally verifiable computation'' which enables a prover to incrementally compute a succinct proof of correct execution of a (potentially) long running process. The paper later won the 2019 TCC test of time award. The construction was proven secure in the random oracle model without any further computational assumptions. However, the overall proof was given using a non-standard version of the random-oracle methodology where sometimes the hash function is a random oracle and sometimes it has a short description as a circuit. Valiant clearly noted that this model is non-standard, but conjectured that the standard random oracle methodology would not suffice. This conjecture has been open for 14 years. We prove that under some mild extra assumptions on the proof system the conjecture is true: the standard random-oracle model does not allow incrementally verifiable computation without making computational assumptions. Two extra assumptions under which we can prove the conjecture are 1) the proof system is also zero-knowledge or 2) when the proof system makes a query to its random oracle it can know with non-negligible probability whether the query is fresh or was made by the proof system earlier in the construction of the proof.

Authors: *Mathias Hall-Andersen and Jesper Buus Nielsen*

Material:

https://eprint.iacr.org/2022/542

#9: FROST: Flexible Round-Optimized Schnorr Threshold Signatures

Mentor : Abdullah Talayhan (abdullah.talayhan@epfl.ch)

Abstract: Unlike signatures in a single-party setting, threshold signatures require cooperation among a threshold number of signers each holding a share of a common private key. Consequently, generating signatures in a threshold setting imposes overhead due to network rounds among signers, proving costly when secret shares are stored on network-limited devices or when coordination occurs over unreliable networks. In this work, we present FROST, a Flexible Round-Optimized Schnorr Threshold signature scheme that reduces network overhead during signing operations while employing a novel technique to protect against forgery attacks applicable to similar schemes in the literature. FROST improves upon the state of the art in Schnorr threshold signature protocols, as it can safely perform signing operations in a single round without limiting concurrency of signing operations, yet allows for true threshold signing, as only a threshold t out of n possible partici- pants are required for signing operations, such that t ≤ n. FROST can be used as either a two-round protocol, or optimized to a single-round signing protocol with a pre-processing stage. FROST achieves its efficiency improvements in part by allowing the protocol to abort in the presence of a misbehaving participant (who is then identified and excluded from future operations)—a reasonable model for practical deployment scenarios. We present proofs of security demonstrating that FROST is secure against chosen-message attacks assuming the discrete logarithm problem is hard and the adversary controls fewer participants than the threshold.

Authors: *Chelsea Komlo and Ian Goldberg*

Material:

https://eprint.iacr.org/2020/852

#10: Constant-Size Commitments to Polynomials and Their Applications

Mentor : Abdullah Talayhan (abdullah.talayhan@epfl.ch)

Abstract: We introduce and formally define polynomial commitment schemes, and provide two efficient constructions. A polynomial com- mitment scheme allows a committer to commit to a polynomial with a short string that can be used by a verifier to confirm claimed evaluations of the committed polynomial. Although the homomorphic commitment schemes in the literature can be used to achieve this goal, the sizes of their commitments are linear in the degree of the committed polyno- mial. On the other hand, polynomial commitments in our schemes are of constant size (single elements). The overhead of opening a commit- ment is also constant; even opening multiple evaluations requires only a constant amount of communication overhead. Therefore, our schemes are useful tools to reduce the communication cost in cryptographic pro- tocols. On that front, we apply our polynomial commitment schemes to four problems in cryptography: verifiable secret sharing, zero-knowledge sets, credentials and content extraction signatures.

Authors: *Aniket Kate, Gregory M. Zaverucha, and Ian Goldberg*

Material:

https://www.iacr.org/archive/asiacrypt2010/6477178/6477178.pdf

#11: CanDID: Can-Do Decentralized Identity with Legacy Compatibility, Sybil-Resistance, and Accountability

Mentor : Betül Durak (betul.durak@microsoft.com)

Abstract: We present CanDID, a platform for practical, user-friendly realization of decentralized identity, the idea of empowering end users with management of their own credentials.

While decentralized identity promises to give users greater control over their private data, it burdens users with management of private keys, creating a significant risk of key loss. Existing and proposed approaches also presume the spontaneous availability of a credential-issuance ecosystem, creating a bootstrapping problem. They also omit essential functionality, like resistance to Sybil attacks and the ability to detect misbehaving or sanctioned users while preserving user privacy.

CanDID addresses these challenges by issuing credentials in a user-friendly way that draws securely and privately on data from existing, unmodified web service providers. Such legacy compatibility similarly enables CanDID users to leverage their existing online accounts for recovery of lost keys. Using a decentralized committee of nodes, CanDID provides strong confidentiality for user's keys, real-world identities, and data, yet prevents users from spawning multiple identities and allows identification (and blacklisting) of sanctioned users.

We present the CanDID architecture and its technical innovations and report on experiments demonstrating its practical performance.

Authors: *Deepak Maram and Harjasleen Malvai and Fan Zhang and Nerla Jean-Louis and Alexander Frolov and Tyler Kell and Tyrone Lobban and Christine Moy and Ari Juels and Andrew Miller*

Material:

https://eprint.iacr.org/2020/934

#12: Token Binding over HTTP

Mentor : Betül Durak (betul.durak@microsoft.com)

Abstract: This document describes a collection of mechanisms that allow HTTP servers to cryptographically bind security tokens (such as cookies and OAuth tokens) to TLS connections.

We describe both first-party and federated scenarios. In a first-party scenario, an HTTP server is able to cryptographically bind the security tokens that it issues to a client -- and that the client subsequently returns to the server -- to the TLS connection between the client and the server. Such bound security tokens are protected from misuse, since the server can generally detect if they are replayed inappropriately, e.g., over other TLS connections.

Federated Token Bindings, on the other hand, allow servers to cryptographically bind security tokens to a TLS connection that the client has with a different server than the one issuing the token.

This document is a companion document to "The Token Binding Protocol Version 1.0" (RFC 8471).

Authors: *Various Authors (IETF)*

Material:

https://datatracker.ietf.org/doc/rfc8473/

#13: Oblivious Pseudorandom Functions (OPRFs) using Prime-Order Groups

Mentor : Laurane Marco (laurane.marco@epfl.ch)

Abstract: An Oblivious Pseudorandom Function (OPRF) is a two-party protocol between client and server for computing the output of a Pseudorandom Function (PRF). The server provides the PRF secret key, and the client provides the PRF input. At the end of the protocol, the client learns the PRF output without learning anything about the PRF secret key, and the server learns neither the PRF input nor output. An OPRF can also satisfy a notion of 'verifiability', called a VOPRF. A VOPRF ensures clients can verify that the server used a specific private key during the execution of the protocol. A VOPRF can also be partially-oblivious, called a POPRF. A POPRF allows clients and servers to provide public input to the PRF computation. This document specifies an OPRF, VOPRF, and POPRF instantiated within standard prime-order groups, including elliptic curves.

Authors: *Various Authors (IETF)*

Material:

https://datatracker.ietf.org/doc/draft-irtf-cfrg-voprf/

#14: Picnic: Post Quantum Signatures

Mentor : Abdullah Talayhan (abdullah.talayhan@epfl.ch)

Abstract: We propose a new class of post-quantum digital signature schemes that: (a) derive their security entirely from the security of symmetric-key primitives, believed to be quantum-secure, and (b) have extremely small keypairs, and, (c) are highly parameterizable.

In our signature constructions, the public key is an image y = f(x) of a one-way function f and secret key x. A signature is a non-interactive zero-knowledge proof of x, that incorporates a message to be signed. For this proof, we leverage recent progress of Giacomelli et al. (USENIX'16) in constructing an efficient Sigma-protocol for statements over general circuits. We improve this Sigma-protocol to reduce proof sizes by a factor of two, at no additional computational cost. While this is of independent interest as it yields more compact proofs for any circuit, it also decreases our signature sizes.

We consider two possibilities to make the proof non-interactive: the Fiat-Shamir transform and Unruh's transform (EUROCRYPT'12, '15,'16). The former has smaller signatures, while the latter has a security analysis in the quantum-accessible random oracle model. By customizing Unruh's transform to our application, the overhead is reduced to 1.6x when compared to the Fiat-Shamir transform, which does not have a rigorous post-quantum security analysis.

We implement and benchmark both approaches and explore the possible choice of f, taking advantage of the recent trend to strive for practical symmetric ciphers with a particularly low number of multiplications and end up using LowMC (EUROCRYPT'15).

Authors: *Melissa Chase and David Derler and Steven Goldfeder and Claudio Orlandi and Sebastian Ramacher and Christian Rechberger and Daniel Slamanig and Greg Zaverucha*

Material:

https://microsoft.github.io/Picnic/

#15:Cryptographic Smooth Neighbors

Mentor: Laurane Marco (laurane.marco@epfl.ch)

Abstract: We revisit the problem of finding two consecutive *B*-smooth integers by giving an optimised implementation of the Conrey-Holmstrom-McLaughlin ``smooth neighbors'' algorithm. While this algorithm is not guaranteed to return the complete set of *B*-smooth neighbors, in practice it returns a very close approximation to the complete set, but does so in a tiny fraction of the time of its exhaustive counterparts. We exploit this algorithm to find record-sized solutions to the pure twin smooth problem. Though these solutions are still not large enough to be cryptographic parameters themselves, we feed them as input into known methods of searching for twins to yield cryptographic parameters that are much smoother than those given in prior works. Our methods seem especially well-suited to finding parameters for the SQISign signature scheme, particularly those that are geared towards high-security levels.

Authors: *Giacomo Bruno, Maria Corte-Real Santos, Craig Costello, Jonathan Komada Eriksen, Michael Naehrig, Michael Meyer, Bruno Sterner*

Material:

https://eprint.iacr.org/2022/1439

#16 : Threshold BBS+ Signatures for Distributed Anonymous Credential Issuance

Mentor : Abdullah Talayhan (abdullah.talayhan@epfl.ch)

Abstract: We propose a secure multiparty signing protocol for the BBS+ signature scheme; in other words, an anonymous credential scheme with threshold issuance. We prove that due to the structure of the BBS+ signature, simply verifying the signature produced by an otherwise semi-honest protocol is sufficient to achieve composable security against a malicious adversary. Consequently, our protocol is extremely simple and efficient: it involves a single request from the client (who requires a signature) to the signing parties, two exchanges of messages among the signing parties, and finally a response to the client; in some deployment scenarios the concrete cost bottleneck may be the client's local verification of the signature that it receives. Furthermore, our protocol can be extended to support the strongest form of blind signing and to serve as a distributed evaluation protocol for the Dodis-Yampolskiy Oblivious VRF. We validate our efficiency claims by implementing and benchmarking our protocol.

Authors: *Jack Doerner, Yashvanth Kondi, Eysa Lee, abhi shelat, LaKyah Tyner*

Material: https://eprint.iacr.org/2023/602

#17 : Separate Your Domains: NIST PQC KEMs, Oracle Cloning and Read-Only Indifferentiability

Mentor : Abdullah Talayhan (abdullah.talayhan@epfl.ch)

Abstract: It is convenient and common for schemes in the random oracle model to assume access to multiple random oracles (ROs), leaving to implementations the task (we call it oracle cloning) of constructing them from a single RO. The first part of the paper is a case study of oracle cloning in KEM submissions to the NIST Post-Quantum Cryptography standardization process. We give key-recovery attacks on some submissions arising from mistakes in oracle cloning, and find other submissions using oracle cloning methods whose validity is unclear. Motivated by this, the second part of the paper gives a theoretical treatment of oracle cloning. We give a definition of what is an "oracle cloning method" and what it means for such a method to "work," in a framework we call read-only indifferentiability, a simple variant of classical indifferentiability that yields security not only for usage in single-stage games but also in multi-stage ones. We formalize domain separation, and specify and study many oracle cloning methods, including common domain-separating ones, giving some general results to justify (prove read-only indifferentiability of) certain classes of methods. We are not only able to validate the oracle cloning methods used in many of the unbroken NIST PQC KEMs, but also able to specify and validate oracle cloning methods that may be useful beyond that.

Authors: *Mihir Bellare, Hannah Davis, and Felix Günther*

Material:

https://eprint.iacr.org/2020/241

#18 : Passive SSH Key Compromise via Lattices

Mentor : Abdullah Talayhan (abdullah.talayhan@epfl.ch)

Abstract: We demonstrate that a passive network attacker can opportunistically obtain private RSA host keys from an SSH server that experiences a naturally arising fault during signature computation. In prior work, this was not believed to be possible for the SSH protocol because the signature included information like the shared Diffie-Hellman secret that would not be available to a passive network observer. We show that for the signature parameters commonly in use for SSH, there is an efficient lattice attack to recover the private key in case of a signature fault. We provide a security analysis of the SSH, IKEv1, and IKEv2 protocols in this scenario, and use our attack to discover hundreds of compromised keys in the wild from several independently vulnerable implementations.

Authors: *Keegan Ryan, Kaiwen He, George Arnold Sullivan, Nadia Heninger*

Material:

https://eprint.iacr.org/2023/1711

#19 : Lightweight Techniques for Private Heavy Filters

Mentor : Abdullah Talayhan (abdullah.talayhan@epfl.ch)

Abstract: This paper presents Poplar, a new system for solving the private heavy-hitters problem. In this problem, there are many clients and a small set of data-collection servers. Each client holds a private bitstring. The servers want to recover the set of all popular strings, without learning anything else about any client's string. A web-browser vendor, for instance, can use Poplar to figure out which homepages are popular, without learning any user's homepage. We also consider the simpler private subset-histogram problem, in which the servers want to count how many clients hold strings in a particular set without revealing this set to the clients. Poplar uses two data-collection servers and, in a protocol run, each client sends only a single message to the servers. Poplar protects client privacy against arbitrary misbehavior by one of the servers and our approach requires no public-key cryptography (except for secure channels), nor general-purpose multiparty computation. Instead, we rely on incremental distributed point functions, a new cryptographic tool that allows a client to succinctly secret-share the labels on the nodes of an exponentially large binary tree, provided that the tree has a single non-zero path. Along the way, we develop new general tools for providing malicious security in applications of distributed point functions. A limitation of Poplar is that it reveals to the servers slightly more information than the set of popular strings itself. We precisely define and quantify this leakage and explain how to ameliorate its effects. In an experimental evaluation with two servers on opposite sides of the U.S., the servers can find the 200 most popular strings among a set of 400,000 client-held 256-bit strings in 54 minutes. Our protocols are highly parallelizable. We estimate that with 20 physical machines per logical server, Poplar could compute heavy hitters over ten million clients in just over one hour of computation.

Authors: *Dan Boneh, Elette Boyle, Henry Corrigan-Gibbs, Niv Gilboa, Yuval Ishai*

Material:

https://eprint.iacr.org/2021/017

#20: Quantum Lightning Never Strikes the Same State Twice

Mentor : Abdullah Talayhan (abdullah.talayhan@epfl.ch)

Abstract: Public key quantum money can be seen as a version of the quantum no-cloning theorem that holds even when the quantum states can be verified by the adversary. In this work, investigate quantum lightning, a formalization of "collision-free quantum money" defined by Lutomirski et al. [ICS'10], where no-cloning holds even when the adversary herself generates the quantum state to be cloned. We then study quantum money and quantum lightning, showing the following results:
- We demonstrate the usefulness of quantum lightning by showing several potential applications, such as generating random strings with a proof of entropy, to completely decentralized cryptocurrency without a block-chain, where transactions is instant and local.
- We give win-win results for quantum money/lightning, showing that either signatures/hash functions/commitment schemes meet very strong recently proposed notions of security, or they yield quantum money or lightning.
- We construct quantum lightning under the assumed multi-collision resistance of random degree-2 systems of polynomials.
- We show that instantiating the quantum money scheme of Aaronson and Christiano [STOC'12] with indistinguishability obfuscation that is secure against quantum computers yields a secure quantum money scheme

Authors: *Mark Zhandry*

Material:

https://arxiv.org/abs/1711.02276

#21: Anonymous Tokens with Hidden Metadata Bit from Algebraic MACs

Mentor : Betül Durak (betul.durak@microsoft.com)

Abstract: On the one hand, the web needs to be secured from malicious activities such as bots or DoS attacks; on the other hand, such needs ideally should not justify services tracking people's activities on the web. Anonymous tokens provide a nice tradeoff between allowing an issuer to ensure that a user has been vetted and protecting the users' privacy. However, in some cases, whether or not a token is issued reveals a lot of information to an adversary about the strategies used to distinguish honest users from bots or attackers.

In this work, we focus on designing an anonymous token protocol between a client and an issuer (also a verifier) that enables the issuer to support its fraud detection mechanisms while preserving users' privacy. This is done by allowing the issuer to embed a hidden (from the client) metadata bit into the tokens. We first study an existing protocol from CRYPTO 2020 which is an extension of Privacy Pass from PoPETs 2018; that protocol aimed to provide support for a hidden metadata bit, but provided a somewhat restricted security notion. We demonstrate a new attack, showing that this is a weakness of the protocol, not just the definition. In particular, the metadata bit hiding is weak in the setting where the attacker can redeem some tokens and get feedback on whether the bit extraction succeeded.

We then revisit the formalism of anonymous tokens with private metadata bit, consider the more natural notion, and design a scheme which achieves it. In order to design this new secure protocol, we base our construction on algebraic MACs instead of PRFs. Our security definitions capture a realistic threat model where adversaries could, through direct feedback or side channels, learn the embedded bit when the token is redeemed. Finally, we compare our protocol with one of the CRYPTO 2020 protocols. We obtain 20% more efficient performance.

Authors: *Melissa Chase, F. Betül Durak, Serge Vaudenay*

Material:

https://eprint.iacr.org/2022/1622

#22: Fair exchange on smart contracts

Mentor : Serge Vaudenay (serge.vaudenay@epfl.ch)

Abstract: We introduce a blockchain Fair Data Exchange (FDE) protocol, enabling a storage server to transfer a data file to a client atomically: the client receives the file if and only if the server receives an agreed-upon payment. We put forth a new definition for a cryptographic scheme that we name verifiable encryption under committed key (VECK), and we propose two instantiations for this scheme. Our protocol relies on a blockchain to enforce the atomicity of the exchange and uses VECK to ensure that the client receives the correct data (matching an agreed-upon commitment) before releasing the payment for the decrypting key. Our protocol is trust-minimized and requires only constant-sized on-chain communication, concretely 3 signatures, 1 verification key, and 1 secret key, with most of the data stored and communicated off-chain. It also supports exchanging only a subset of the data, can amortize the server's work across multiple clients, and offers a general framework to design alternative FDE protocols using different commitment schemes. A prominent application of our protocol is the Danksharding data availability scheme on Ethereum, which commits to data via KZG polynomial commitments. We also provide an open-source implementation for our protocol with both instantiations for VECK, demonstrating our protocol's efficiency and practicality on Ethereum.

Authors: *Ertem Nusret Tas, István András Seres, Yinuo Zhang, Márk Melczer, Mahimna Kelkar, Joseph Bonneau, Valeria Nikolaenko*

Material:

https://eprint.iacr.org/2024/418

https://dl.acm.org/doi/10.1145/3658644.3690248

#23: zkLogin: Privacy-Preserving Blockchain Authentication with Existing Credentials

Mentor : Serge Vaudenay (serge.vaudenay@epfl.ch)

Abstract: For many users, a private key based wallet serves as the primary entry point to blockchains. Commonly recommended wallet authentication methods, such as mnemonics or hardware wallets, can be cumbersome. This difficulty in user onboarding has significantly hindered the adoption of blockchain-based applications.
We develop zkLogin, a novel technique that leverages identity tokens issued by popular platforms (any OpenID Connect enabled platform e.g., Google, Facebook, etc.) to authenticate transactions. At the heart of zkLogin lies a signature scheme allowing the signer to sign using their existing OpenID accounts and nothing else. This improves the user experience significantly as users do not need to remember a new secret and can reuse their existing accounts.
zkLogin provides strong security and privacy guarantees. Unlike prior works, zkLogin's security relies solely on the underlying platform's authentication mechanism without the need for any additional trusted parties (e.g., trusted hardware or oracles). As the name suggests, zkLogin leverages zero-knowledge proofs (ZKP) to ensure that the sensitive link between a user's off-chain and on-chain identities is hidden, even from the platform itself.
zkLogin enables a number of important applications outside blockchains. It allows billions of users to produce \textit{verifiable digital content leveraging their existing digital identities}, e.g., email address. For example, a journalist can use zkLogin to sign a news article with their email address, allowing verification of the article's authorship by any party.
We have implemented and deployed zkLogin on the Sui blockchain as an additional alternative to traditional digital signature-based addresses.

Authors: *Foteini Baldimtsi, Konstantinos Kryptos Chalkias, Yan Ji, Jonas Lindstrøm, Deepak Maram, Ben Riva, Arnab Roy, Mahdi Sedaghat, Joy Wang*

Material:

https://arxiv.org/abs/2401.11735

https://dl.acm.org/doi/10.1145/3658644.3690356

#24: Complete Knowledge: Preventing Encumbrance of Cryptographic Secrets

Mentor : Serge Vaudenay (serge.vaudenay@epfl.ch)

Abstract: Most cryptographic protocols model a player's knowledge of secrets in a simple way. Informally, the player knows a secret in the sense that she can directly furnish it as a (private) input to a protocol, e.g., to digitally sign a message. The growing availability of Trusted Execution Environments (TEEs) and secure multiparty computation, however, undermines this model of knowledge. Such tools can encumber a secret sk and permit a chosen player to access sk conditionally, without actually knowing sk. By permitting selective access to sk by an adversary, encumbrance of secrets can enable vote-selling in cryptographic voting schemes, illegal sale of credentials for online services, and erosion of deniability in anonymous messaging systems. Unfortunately, existing proof-of-knowledge protocols fail to demonstrate that a secret is unencumbered. We therefore introduce and formalize a new notion called complete knowledge (CK). A proof (or argument) of CK shows that a prover does not just know a secret, but also has fully unencumbered knowledge, i.e., unrestricted ability to use the secret. We introduce two practical CK schemes that use special-purpose hardware, specifically TEEs and off-the-shelf mining ASICs. We prove the security of these schemes and explore their practical deployment with a complete, end-to-end prototype with smart-contract verification that supports both. We show how CK can address encumbrance attacks identified in previous work. Finally, we introduce two new applications enabled by CK that involve proving ownership of blockchain assets.

Authors: *Mahimna Kelkar, Kushal Babel, Philip Daian, James Austgen, Vitalik Buterin, Ari Juels,*

Material:

https://eprint.iacr.org/2023/044

https://doi.org/10.1145/3658644.3690273

https://doi.org/10.1007/978-3-642-04138-9_29

#25: Multivariate Blind Signatures Revisited

Mentor : Max Duparc (max.duparc@epfl.ch)

Abstract: In 2017, Petzoldt, Szepieniec, and Mohamed proposed a blind signature scheme, based on multivariate cryptography. This construction has been expanded on by several other works. This short paper shows that their construction is susceptible to an efficient polynomial-time attack. The problem is that the authors implicitly assumed that for a random multivariate quadratic map and a collision-resistant hash function , the function is a binding commitment, which is not the case. There is a "folklore" algorithm that can be used to, given any pair of messages, efficiently produce a commitment that opens to both of them. We hope that by pointing out that multivariate quadratic maps are not binding, similar problems can be avoided in the future.

Authors: *Ward Beullens*

Material:

https://eprint.iacr.org/2024/720

#26: Rapidash: Foundations of Side-Contract-Resilient Fair Exchange

Mentor : Serge Vaudenay (serge.vaudenay@epfl.ch)

Abstract: Fair exchange is a fundamental primitive enabled by blockchains, and is widely adopted in applications such as atomic swaps, payment channels, and DeFi. Most existing designs of blockchain-based fair exchange protocols consider only the participating users as strategic players, and assume the miners are honest and passive. However, recent works revealed that the fairness of commonly deployed fair exchange protocols can be broken entirely in the presence of user-miner collusion. In particular, a user can bribe the miners to help it cheat — a phenomenon also referred to as Miner Extractable Value (MEV). In this work, we provide the first formal treatment of side-contract-resilient fair exchange where users and miners may enter into arbitrary contracts on the side. We propose a new fair exchange protocol called Rapidash, and prove that the protocol is incentive compatible in the presence of user-miner collusion. In particular, we show that Rapidash satisfies a coalition-resistant Nash equilibrium absent external incentives. Further, even when there exist arbitrary but bounded external incentives, Rapidash still protects honest players and ensures that they cannot be harmed. Last but not least, our game-theoretic formulations also lay the theoretical groundwork for studying side-contract-resilient fair exchange protocols. Finally, to showcase the instantiability of Rapidash with a wide range of blockchain systems, we present instantiations of Rapidash that are compatible with Bitcoin and Ethereum while incurring only a minimal overhead in terms of costs for the users.

Authors: *Hao Chung, Elisaweta Masserova, Elaine Shi, Sri AravindaKrishnan Thyagarajan*

Material:

https://eprint.iacr.org/2022/1063

#27: MPC in the head using the subfield bilinear collision problem

Mentor : Lewis Glabush (lewis.glabush@epfl.ch)

Abstract: In this paper, we introduce the subfield bilinear collision problem and use it to construct an identification protocol and a signature scheme. This construction is based on the MPC-in-the-head paradigm and uses the Fiat-Shamir transformation to obtain a signature.

Authors: *Janik Huth, Antoine Joux*

Material:

https://eprint.iacr.org/2023/1685 +
https://www.springerprofessional.de/en/mpc-in-the-head-using-the-subfield-bilinear-collision-problem/27475950

#28: How to Prove False Statements: Practical Attacks on Fiat-Shamir

Mentor : Betül Durak **(**betul.durak@microsoft.com **)**

Abstract: The Fiat-Shamir (FS) transform is a prolific and powerful technique for compiling public-coin interactive protocols into non-interactive ones. Roughly speaking, the idea is to replace the random coins of the verifier with the evaluations of a complex hash function. The FS transform is known to be sound in the random oracle model (i.e., when the hash function is modeled as a totally random function). However, when instantiating the random oracle using a concrete hash function, there are examples of protocols in which the transformation is not sound. So far all of these examples have been contrived protocols that were specifically designed to fail. In this work we show such an attack for a standard and popular interactive succinct argument, based on the GKR protocol, for verifying the correctness of a non-determinstic bounded-depth computation. For every choice of FS hash function, we show that a corresponding instantiation of this protocol, which was been widely studied in the literature and used also in practice, is not (adaptively) sound when compiled with the FS transform. Specifically, we construct an explicit circuit for which we can generate an accepting proof for a false statement. We further extend our attack and show that for every circuit and desired output , we can construct a functionally equivalent circuit , for which we can produce an accepting proof that outputs (regardless of whether this statement is true). This demonstrates that any security guarantee (if such exists) would have to depend on the specific implementation of the circuit , rather than just its functionality. Lastly, we also demonstrate versions of the attack that violate non-adaptive soundness of the protocol -- that is, we generate an attacking circuit that is independent of the underlying cryptographic objects. However, these versions are either less practical (as the attacking circuit has very large depth) or make some additional (reasonable) assumptions on the underlying cryptographic primitives.

Authors: *Dmitry Khovratovich, Ron D. Rothblum, Lev Soukhanov*

Material:

https://eprint.iacr.org/2025/118

#29: DiStefano: Decentralized Infrastructure for Sharing Trusted Encrypted Facts and Nothing More

Mentor : Betül Durak **(**betul.durak@microsoft.com **)**

Abstract: We design DiStefano: an efficient, maliciously-secure framework for generating private commitments over TLS-encrypted web traffic, for verification by a designated third-party. DiStefano provides many improvements over previous TLS commitment systems, including: a modular protocol specific to TLS 1.3, support for arbitrary verifiable claims over encrypted data, client browsing history privacy amongst pre-approved TLS servers, and various optimisations to ensure fast online performance of the TLS 1.3 session. We build a permissive open-source implementation of DiStefano integrated into the BoringSSL cryptographic library (used by Chromium-based Internet browsers). We show that DiStefano is practical in both LAN and WAN settings for committing to facts in arbitrary TLS traffic, requiring 1 s and 80 KiB to execute the complete online phase of the protocol.

Authors: *Dmitry Khovratovich, Ron D. Rothblum, Lev Soukhanov*

Material:

https://eprint.iacr.org/2023/1063

#30: Side-channel Assisted Existential Forgery Attack on Dilithium - A NIST PQC candidate

Mentor : Keng-Yu Chen (keng-yu.chen@epfl.ch)

Abstract: The recent lattice-based signature scheme Dilithium, submitted as part of the CRYSTALS (Cryptographic Suite for Algebraic Lattices) package, is one of a number of strong candidates submitted for the NIST standardisation process of post-quantum cryptography. The Dilithium signature scheme is based on the Fiat-Shamir paradigm and can be seen as a variant of the Bai-Galbraith scheme (BG) combined with several improvements from previous ancestor lattice-based schemes like GLP and BLISS signature schemes. One of the main features of Dilithium is the compressed public-key, which is a rounded version of the LWE instance. This implies that Dilithium is not breakable with the knowledge of only the secret or the error of the LWE instance, unlike its ancestor lattice-based signature schemes. In this paper, we investigate the security of Dilithium against a combination of side-channel and classical attacks. Side-channel attacks on schoolbook and optimised polynomial multiplication algorithms in the signing procedure are shown to extract the secret component of the LWE instance, which is just one among the multiple components of the secret-key of Dilithium. We then propose an alternative signing procedure, through which it is possible to forge signatures with only the extracted portion of the secret-key, without requiring the knowledge of all its elements. Thus showing that Dilithium too breaks on just knowing the secret portion of the LWE instance, similar to previous lattice-based schemes.

Authors: *Prasanna Ravi, Mahabir Prasad Jhanwar, James Howe, Anupam Chattopadhyay, and Shivam Bhasin*

Material:

https://eprint.iacr.org/2018/821

#31: The Hidden Parallelepiped Is Back Again: Power Analysis Attacks on Falcon

Mentor : Keng-Yu Chen (keng-yu.chen@epfl.ch)

Abstract: Falcon is a very efficient and compact lattice-based signature finalist of the NIST's Post-Quantum standardization campaign. This work assesses Falcon's side-channel resistance by analyzing two vulnerabilities, namely the pre-image computation and the trapdoor sampling. The first attack is an improvement of Karabulut and Aysu (DAC 2021). It overcomes several difficulties inherent to the structure of the stored key like the Fourier representation and directly recovers the key with a limited number of traces and a reduced complexity. The main part of this paper is dedicated to our second attack: we show that a simple power analysis during the signature execution could provide the exact value of the output of a subroutine called the base sampler. This intermediate value does not directly lead to the secret and we had to adapt the so-called hidden parallelepiped attack initially introduced by Nguyen and Regev in Eurocrypt 2006 and reused by Ducas and Nguyen in Asiacrypt 2012. We extensively quantify the resources for our attacks and experimentally demonstrate them with Falcon's reference implementation on the ELMO simulator (McCann, Oswald and Whitnall USENIX 2017) and on a Chipwhisperer Lite with STM32F3 target (ARM Cortex M4).
These new attacks highlight the need for side-channel protection for one of the three finalists of NIST's standardization campaign by pointing out the vulnerable parts and quantifying the resources of the attacks.

Authors: *Morgane Guerreau, Ange Martinelli, Thomas Ricosset, Thales,Mélissa Rossi*

Material:

https://eprint.iacr.org/2022/057

#32: How (not) to Use Welch's T-test in Side-Channel Security Evaluations

Mentor : Keng-Yu Chen (keng-yu.chen@epfl.ch)

Abstract: The Test Vector Leakage Assessment (TVLA) methodology is a qualitative tool relying on Welch's T-test to assess the security of cryptographic implementations against side-channel attacks. Despite known limitations (e.g., risks of false negatives and positives), it is sometimes considered as a pass-fail test to determine whether such implementations are "safe" or not (without clear definition of what is "safe"). In this note, we clarify the limited quantitative meaning of this test when used as a standalone tool. For this purpose, we first show that the straightforward application of this approach to assess the security of a masked implementation is not sufficient. More precisely, we show that even in a simple (more precisely, univariate) case study that seems best suited for the TVLA methodology, detection (or lack thereof) with Welch's T-test can be totally disconnected from the actual security level of an implementation. For this purpose, we put forward the case of a realistic masking scheme that looks very safe from the TVLA point-of-view and is nevertheless easy to break. We then discuss this result in more general terms and argue that this limitation is shared by all "moment-based" security evaluations. We conclude the note positively, by describing how to use moment-based analyzes as a useful ingredient of side-channel security evaluations, to determine a "security order".

Authors: *François-Xavier Standaert*

Material:

https://eprint.iacr.org/2017/138

#33: Random-Oracle Uninstantiability from Indistinguishability Obfuscation

Mentor : Keng-Yu Chen (keng-yu.chen@epfl.ch)

Abstract: Assuming the existence of indistinguishability obfuscation (iO), we show that a number of prominent transformations in the random-oracle model are uninstantiable in the standard model. We start by showing that the Encrypt-with-Hash transform of Bellare, Boldyreva and O'Neill (CRYPTO 2007) for converting randomized public-key encryption schemes to deterministic ones is not instantiable in the standard model. To this end, we build on the recent work of Brzuska, Farshim and Mittelbach (CRYPTO 2014) and rely on the existence of iO for circuits or iO for Turing machines to derive uninstantiability for hash functions of a priori bounded polynomial size and arbitrary polynomial size, respectively. The techniques that we use to establish this result are flexible and lend themselves to a number of other transformations such as the classical Fujisaki--Okamoto transform (CRYPTO 1998) and transformations akin to those by Bellare and Keelveedhi (CRYPTO 2011) and Douceur et al. (ICDCS 2002) for obtaining KDM-secure encryption and de-duplication schemes respectively. Our results call for a re-assessment of scheme design in the random-oracle model and highlight the need for new transforms that do not suffer from iO-based attacks.

Authors: *Chris Brzuska, Pooya Farshim, and Arno Mittelbach*

Material:

https://eprint.iacr.org/2014/867

#34: Relational Hash

Mentor : Keng-Yu Chen (keng-yu.chen@epfl.ch)

Abstract: A Traditional cryptographic hash functions allow one to easily check whether the original plaintexts are equal or not, given a pair of hash values. Probabilistic hash functions extend this concept where given a probabilistic hash of a value and the value itself, one can efficiently check whether the hash corresponds to the given value. However, given distinct probabilistic hashes of the same value it is not possible to check whether they correspond to the same value. In this work we introduce a new cryptographic primitive called \emph{Relational Hash} using which, given a pair of (relational) hash values, one can determine whether the original plaintexts were related or not. We formalize various natural security notions for the Relational Hash primitive - one-wayness, twin one-wayness, unforgeability and oracle simulatibility. We develop a Relational Hash scheme for discovering linear relations among bit-vectors (elements of) and -vectors. Using the linear Relational Hash schemes we develop Relational Hashes for detecting proximity in terms of hamming distance. The proximity Relational Hashing schemes can be adapted to a privacy preserving biometric identification scheme, as well as a privacy preserving biometric authentication scheme secure against passive adversaries

Authors: *Avradip Mandal and Arnab Roy*

Material:

https://eprint.iacr.org/2014/394

#35: Breaking SIDH in polynomial time

Mentor : Max Duparc (max.duparc@epfl.ch)

Abstract: We show that we can break SIDH in classical polynomial time, even with a random starting curve.

Authors: *Damien Robert*

Material:

https://eprint.iacr.org/2022/1038

#36: Constant time lattice reduction in dimension 4 with application to SQIsign

Mentor : Max Duparc (max.duparc@epfl.ch)

Abstract: In this paper we propose a constant time lattice reduction algorithm for integral dimension-4 lattices. Motivated by its application in the SQIsign post-quantum signature scheme, we provide for the first time a constant time LLL-like algorithm with guarantees on the length of the shortest output vector. We implemented our algorithm and ensured through various tools that it indeed operates in constant time. Our experiments suggest that in practice our implementation outputs a Minkowski reduced basis and thus can replace a non constant time lattice reduction subroutine in SQIsign.

Authors: *Otto Hanyecz, Alexander Karenin, Elena Kirshanova, Péter Kutas, Sina Schaeffler*

Material:

https://eprint.iacr.org/2025/027

#37: More Practical Single-Trace Attacks on the Number Theoretic Transform

Mentor : Max Duparc (max.duparc@epfl.ch)

Abstract: Single-trace side-channel attacks are a considerable threat to implementations of classic public-key schemes. For lattice-based cryptography, however, this class of attacks is much less understood, and only a small number of previous works show attacks. Primas et al., for instance, present a single-trace attack on the Number Theoretic Transform (NTT), which is at the heart of many efficient lattice-based schemes.

They, however, attack a variable-time implementation and also require a rather powerful side-channel adversary capable of creating close to a million multivariate templates. Thus, it was an open question if such an attack can be made practical while also targeting state-of-the-art constant-time implementations.

In this paper, we answer this question positively. First, we introduce several improvements to the usage of belief propagation, which underlies the attack. And second, we change the target to encryption instead of decryption; this limits attacks to the recovery of the transmitted symmetric key, but in turn, increases attack performance. All this then allows successful attacks even when switching to univariate Hamming-weight templates. We evaluate the performance and noise resistance of our attack using simulations, but also target a real device. Concretely, we successfully attack an assembly-optimized constant-time Kyber implementation running on an ARM Cortex M4 microcontroller while requiring the construction of only 213 templates.

Authors: *Peter Pessl and Robert Primas*

Material:

https://eprint.iacr.org/2019/795

#38: Efficient verifiable delay functions

Mentor : Max Duparc (max.duparc@epfl.ch)

Abstract: We construct a verifiable delay function (VDF). A VDF is a function whose evaluation requires running a given number of sequential steps, yet the result can be efficiently verified. They have applications in decentralised systems, such as the generation of trustworthy public randomness in a trustless environment, or resource-efficient blockchains. To construct our VDF, we actually build a trapdoor VDF. A trapdoor VDF is essentially a VDF which can be evaluated efficiently by parties who know a secret (the trapdoor). By setting up this scheme in a way that the trapdoor is unknown (not even by the party running the setup, so that there is no need for a trusted setup environment), we obtain a simple VDF. Our construction is based on groups of unknown order such as an RSA group, or the class group of an imaginary quadratic field. The output of our construction is very short (the result and the proof of correctness are each a single element of the group), and the verification of correctness is very efficient.

Authors: *Benjamin Wesolowski*

Material:

https://eprint.iacr.org/2018/623

#39: SQISignHD: New Dimensions in Cryptography

Mentor : Max Duparc (max.duparc@epfl.ch)

Abstract: We introduce SQIsignHD, a new post-quantum digital signature scheme inspired by SQIsign. SQIsignHD exploits the recent algorithmic breakthrough underlying the attack on SIDH, which allows to efficiently represent isogenies of arbitrary degrees as components of a higher dimensional isogeny. SQIsignHD overcomes the main drawbacks of SQIsign. First, it scales well to high security levels, since the public parameters for SQIsignHD are easy to generate: the characteristic of the underlying field needs only be of the form . Second, the signing procedure is simpler and more efficient. Our signing procedure implemented in C runs in 28 ms, which is a significant improvement compared to SQISign. Third, the scheme is easier to analyse, allowing for a much more compelling security reduction. Finally, the signature sizes are even more compact than (the already record-breaking) SQIsign, with compressed signatures as small as 109 bytes for the post-quantum NIST-1 level of security. These advantages may come at the expense of the verification, which now requires the computation of an isogeny in dimension , a task whose optimised cost is still uncertain, as it has been the focus of very little attention. Our experimental sagemath implementation of the verification runs in around 600 ms, indicating the potential cryptographic interest of dimension isogenies after optimisations and low level implementation.

Authors: *Pierrick Dartois, Antonin Leroux, Damien Robert, Benjamin Wesolowski*

Material:

https://eprint.iacr.org/2023/436

#40: On Protecting SPHINCS+ Against Fault Attacks

Mentor : Max Duparc (max.duparc@epfl.ch)

Abstract: SPHINCS+ is a hash-based digital signature scheme that was selected by NIST in their post-quantum cryptography standardization process. The establishment of a universal forgery on the seminal scheme SPHINCS was shown to be feasible in practice by injecting a fault when the signing device constructs any non-top subtree. Ever since the attack has been made public, little effort was spent to protect the SPHINCS family against attacks by faults. This paper works in this direction in the context of SPHINCS+ and analyzes the current algorithms that aim to prevent fault-based forgeries. First, the paper adapts the original attack to SPHINCS+ reinforced with randomized signing and extends the applicability of the attack to any combination of faulty and valid signatures. Considering the adaptation, the paper then presents a thorough analysis of the attack. In particular, the analysis shows that, with high probability, the security guarantees of SPHINCS+ significantly drop when a single random bit flip occurs anywhere in the signing procedure and that the resulting faulty signature cannot be detected with the verification procedure. The paper shows both in theory and experimentally that the countermeasures based on caching the intermediate W-OTS+s offer a marginally greater protection against unintentional faults, and that such countermeasures are circumvented with a tolerable number of queries in an active attack. Based on these results, the paper recommends real-world deployments of SPHINCS+ to implement redundancy checks.

Authors: *Aymeric Genêt*

Material:

https://eprint.iacr.org/2023/042