ACNS 2023 Field Notes

Gary Hu

June 2023

Contents

1 Introduction

1.1 About ACNS

ACNS is an annual conference focusing on current developments that advance the areas of applied cryptography, cyber security (including network and computer security) and privacy. The goal is to represent both academic research works as well as developments in industrial and technical frontiers. Submissions may focus on the modelling, design, analysis (including security proofs and attacks), development (e.g. implementations), deployment (e.g. system integration), and maintenance (e.g. performance measurements, usability studies) of algorithms, protocols, standards, implementations, technologies devices, systems standing in relation with applied cryptography, cyber security and privacy, while advancing or bringing new insights to the state of the art.

From June 19th to June 22nd, the 21st International Conference on Applied Cryptography and Network Security was held in Kyoto, Japan.

1.2 Field Notes

Unfortunately I could not make the trip out to Kyoto for this conference, but luckily I was able to virtually attend the conference through Zoom.

I wrote down the main ideas of all of the events I attended. They are often copied word for word from the conclusion slide from each presentation and are very similar to the abstracts of the original papers. Sections are ordered in chronological order, talks in each section are ordered by how much I liked them from the most to the least. This document is mostly for my future reference if I want to return to anything from this, but I've publicly posted it in case it benefits anyone else.

The usual disclaimer holds: all intellectual work is due to the original author (and not me), and all errors should be attributed to me.

2 Day 1

Contents

2.1 Web Security

2.1.1 Capturing Antique Browsers in Modern Devices: A Security Analysis of Captive Portal Mini-Browsers

Captive portals are webpages requiring users to authenticate themselves before accessing a public Wi-Fi network, and devices detect these through mini-browsers. What happens if the captive portal is malicious? The two main challenges are that user portal browsers are different from normal browsers and have limited internet access. The authors contributed:

- **Wi-Fi Chameleon**: a tool to analyze attacks on user portal browsers and finds critical vulnerabilities. The two primary attacks are evil-twin attacks and history-stealing attacks. Evil-twin attacks occur through a fake Wi-Fi AP with the same SSID usually done through URI modification, and can be prevented through warning messages/indicator and certificate validation. History-stealing attacks occur when a hacker steals browsing history (cookies, local storage, HSTS records, etc.) and can be prevented through cookies and local storage.
- Defense schemes to secure captive portals, such as a secure browser extension and identity verification. Their secure browser extension has trust on first use plus an HTTPS list, records the correct URI, and prevents URI modification. Their AP identity verification binds SSID with URI, verifies with TLS certificate, and prevents evil-twin attacks.

2.1.2 Those Aren't Your Memories and They're Somebody Else's: Seeding Misinformation in Chat Bot Memories

In 2022, Blender Bot 2.0 was released, an open-source chatbot that builds long-term memory and searches the internet.

The authors revealed that the bots can be easily manipulated to remember misinformation alongside personal knowledge, causing a high likelihood of responding with the misinformation as a fact.

They concluded the following:

1. Personal statements cause the chatbot to remember information.
2. The personal statement content does not impact the rate of memory generation.
3. Personal statements led to memorization.
4. Memorized misinformation can be recalled in responses.
5. When the bot is questioned, there is a 328% increase in misinformation.

2.1.3 Social Honeypot for Humans: Luring People through Self-managed Instagram Pages

Honeypots are profiles on Online Social Networks (OSN) that try to lure spammers and trap them. Information about their activities is used for developing countermeasures to reduce their presence, and are usually deployed on Twitter.

The authors proposed a novel tool for attracting OSN users interested in a target topic, proposing a framework to mimic legit Instagram pages. They generated a topic according to the honeypot purpose, and then automatically generated an image and text using one of four methods: InstaModel, ArtModel, Unsplash-Model, and QuotesModel. They ended up gaining valuable age/location data of spammers, and found out that food accounts were followed the most and sponsored posts performed much better. More importantly, they concluded that topic and engagement plans are more important than generation strategies.

2.1.4 Tiny WFP: Lightweight and Effective Website Fingerprinting Via Wavelet Multi-Resolution Analysis

Tor is widely used and creates opportunities for illicit activities, while website fingerprinting attacks can be used for network supervision. The motivation behind this talk is:

- WF attacks based on deep learning significantly outperform traditional machine learning-based ones.
- Existing traffic analysis mechanisms suffer from enormous storage, communications, and computation overheads.

The authors of this talk created a novel and accurate approach combining multiscale wavelet decomposition and depthwise separable convolution to website fingerprint attacks, which they go over in high detail in this talk.

2.2 Keynote

2.2.1 Challenges and Solutions to Post-Quantum Messaging

A talk about secure messaging protocols (Signal and MLS) and how to fix their vulnerabilities to quantum attacks. These messaging apps are dangerous because of two security threats:

1. They are attached to our smartphones, which are attached to our personal lives and can reveal metadata.
2. We have longer conversations compared to TLS-based communication, which creates a longer period for attacks.

The talk was presented in 5 steps:

1. **Challenges in PQC:** The need to understand the challenges of PQC is urgent, seen by many (including government agencies). Unfortunately,

we have several problems: 1) PQC causes extremely high communication costs, and 2) there does not exist a PQC counterpart for Diffie-Hellman.

2. **Secure Messaging in Brief:** Secure messaging (SM) is a protocol between a set of users and a server, consisting of how to share private keys and use/maintain them. We want the server to learn nothing about the message.

3. **Post Quantum Signal:** The Signal Protocol is a popular great standard for two-user SM, relying on Diffie-Hellman. It consists of two parts: X3DH (establishes secret key) and Double Rachet (encrypted communication with continuous key updates). Recently, there has been a lot of work in turning both into PQ: A PQ AKE was proposed for a PQ version of X3DH in 2021, and a PQ Double Rachet was proposed in 2019.

4. **Post Quantum MLS:** MLS is another popular scalable solution to SM, build upon a key update abstracted as a continuous group key agreement (CGKA). In 2020, a lattice-based Multi-Recipient PKE was created, leading to an MLS protocol that rivaled Kyber-512 and created better PQ bandwidth costs.

5. **About Metadata Hiding:** SM is not all about E2EE. Metadata reveals a lot of information, and can be split into three layers: messages, static explicit metadata, and dynamic implicit metadata. Recently, a PC Metadata-Hiding MLS was created.

In summary, PQC creates new challenges to SM, which already has many open problems (rising tension in regulation of E2EE, content moderation, and interoperability).

2.3 Lattices and Codes

2.3.1 BIKE Key-Recovery: Combining Power Consumption Analysis and Information-Set Decoding

BIKE is a NIST PQC KEM candidate, which the authors exploit. They start by exploiting the rotation function. Then, they use two implemented versions of this function (written in C and assembly) and combine clustering and information-set decoding to fully exploit BIKE.

2.3.2 Forward Security of Fiat-Shamir Lattice Signatures

Forward security is desired for lattice-based signatures by NIST, but known methods (using short basis) are not compatible for Fiat-Shamir type lattice signatures. The goal of the authors was to design a forward-secure Fiat-Shamir Lattice without any additional short basis while the delegated trapdoor is in a constant dimension. The authors did this using a binary tree structure and a new trapdoor evolution, creating forward-secure Fiat-Shamir lattice signatures with a smaller secret key and efficient key update.

2.3.3 Spherical Gaussian Leftover Hash Lemma via the Rényi Divergence

The discrete Gaussian leftover hash lemma was proved in 2013, stating that the linear transformation of the discrete spherical Gaussian is statistically close to the discrete ellipsoid Gaussian. The authors improved the ellipsoid to spherical, which they called the weak Spherical Gaussian LHL (wSGLHL) theorem. Additionally, the authors applied this to create a sharper LWE self-reduction and mitigated the heuristic for TFHE.

2.3.4 A Gapless Post-Quantum Hash Proof System in the Hamming Metric

The authors created an alternative to Bettaieb's construction for the Hamming metric. Their construction used well-known assumptions (syndrome decoding and variants) and "acknowledged" cryptographic primitives (HQC).

2.3.5 Shorter and Faster Identity-Based Signatures with Tight Security in the (Q)ROM from Lattices

The authors improved identity-based signature (IBS) schemes that offer tight security against adaptive adversaries in both classical and quantum random oracle models. The schemes are applicable to various lattice structures and rely on the SIS/RSIS assumption, achieving shorter signatures, better security bounds, and faster algorithms compared to previous work by simplifying parameters, providing concrete values, and utilizing a non-homogeneous relation to reduce signature size and eliminate a costly trapdoor delegation.

2.4 Symmetric Cryptanalysis

2.4.1 A Novel Automatic Technique Based on MILP to Search for Impossible Differentials

The Mixed Integer Linear Programming (MILP) is a common method to search for impossible differentiables, and the authors aim to design a simpler MILP model (and one that includes key-recovery). The authors proposed a simple modeling technique with two-dimensional binary variables to search for IDs, and then applied their work to Midori-64 (and the current best 11 round ID attack on Midori-64), CRAFT, and SKINNY-64.

2.4.2 TIDAL: Practical Collisions on State-Reduced Keccak Variants

The authors proposed a new collision search strategy TIDAL that produces states after 2 rounds of KECCAK-*p*. TIDAL produced self-symmetric states using squeeze attacks to generate faster collisions in a special target subset and a 1.5 round deterministic trail. The authors also showed that their strategy penetrated higher rounds of Type-I variants and improved the TIDA approach.

2.4.3 Near Collision Attack Against Grain v1

Grain v1 is an 80-bit security lightweight stream cipher that incorrectly claimed to have a near-collision attack. The authors proposed a new near-collision attack on Grain v1. They concluded that Grain v1's low sampling resistance allows for a trade of memory for a number of encryptions, which suggests an algebraic deficiency of Grain v1. The authors also applied this attack to Grain 128 and 128a.

2.4.4 Meet-in-the-Filter and Dynamic Counting with Applications to Speck

The authors proposed a new cryptanalytic tool for differential cryptanalysis for ciphers with a slow or incomplete diffusion layer called meet-in-the-filter (MiF). Then, the authors applied this to create the best-known attacks up to 15 rounds of Speck32.

2.5 Poster Session

There were a lot of posters, and I didn't have the time to write notes for each of them, so here are the three that I found most intriguing.

2.5.1 A Card-Based Protocol That Lets You Know How Close Two Parties Are in Their Opinions (Agree/Disagree) by Using a Four-Point Likert Scale.

The best poster by far. The abstract is quite perfect, so I'm copying it word for word:

Card-based protocols are commonly used for scenarios where two parties need to perform an AND operation. These protocols are known for providing a non-embarrassing way to confess love, as the other party cannot determine whether the input is 0 or 1.

This paper proposes a new card-based protocol that involves four cards distributed between two parties, and is based on the assumption that all four values are accepted. The protocol asks both parties to indicate how close their opinions are on a four-point scale (0 = not at all disagree, 1 = somewhat disagree, 2 = somewhat agree, and 3 = strongly agree), which is commonly used in surveys. The protocol consists of three patterns: complete agreement, approximate agreement (cases with inputs 0,1 or 2,3), and disagreement. The results are only known to the parties, and the inputs are kept secret from the third party. It should be noted that our Likert scale excludes the neutral response "undecided". Furthermore, the case of our proposal has the algebraic structure as one of the association schemes with 4 points.

2.5.2 AuthZit: Multi-Modal Authentication with Visual Spatial and Text Secrets

People have strong visual-spatial memory and memory via association, so AuthZit encodes authentication secrets as paths through a 3D map of places, and people need to navigate the map in first person and birds-eye perspective to find specific locations and then need to give a textual secret tagged with the location. The authors performed a user study to demonstrate that AuthZit had a high success rate of authentication with little reinforcement. The downside is that registration and log-in will take longer, but the authors believe that the benefits outweigh the costs.

2.5.3 Integrating Quantum Key Distribution into Hybrid Quantum-Classical Networks

The poster can be split into three areas:

- Conceptual architecture of hybrid quantum-classical networks built with a bottom-up approach by gradual incorporation of QKD links into the existing communication infrastructure. The architecture is built with a bottom-up approach by gradually incorporating QKD links into the existing communication infrastructure, emphasizing the need to convey the benefits of QKD-generated keys to users not directly connected with quantum links.

- A novel Butterfly protocol for delivering QKD-generated keys as a service. Concrete proposals exist for protocols providing quantum keys as a service, addressing reliance on a single geographically closest QKD node. The authors introduce the Butterfly protocol to overcome the vulnerability of relying on a single node by assuming the compromise of two independent TLS communication links, connecting users independently to two Key Distribution Centres (KDC) linked with paired QKD devices, establishing a QKD-keyed TLS session.

- Specification of requirements for key synchronization and QKD device authentication in hybrid networks, and discussion of solving these problems with the available technologies. Two important aspects unaddressed by quantum cryptography: Key synchronization among multiple nodes and QKD device authentication over the network. Secure solutions for these problems exist but face efficiency limitations, posing a challenge for developing scalable protocols. Known solutions include the BBN Key relay protocol, OTP, universal hash function, and PQC-based authentication schemes.

2.6 General Thoughts

Here is a ranking from best to worst in terms of section.

- **Web Security.** This section was absolutely awesome in my opinion and I had a really tough time ranking which talk I liked the most because they were all amazing. At the end I just used a RNG to rank them. Ironically I thought this section was going to be the most boring of the day and almost skipped it, so I'm very glad that I didn't skip it.
- **Keynote.** The keynote was exceptionally good. The topic was quite interesting but new to me, so bonus points to Shuichi Katsumata for presenting this at the perfect level for me to understand. Slides were amazing, pace was amazing, I have no criticisms for this talk.
- **Poster Session.** I felt like this section was hit or miss. Most of the posters were not that interesting to me at all, but then there were those three posters that were super, super cool. The card poster is probably the coolest thing I've seen all year.
- **Symmetric Cryptanalysis.** I was most excited for this section and it didn't disappoint, but I didn't think any of the talks were extremely mindblowing either. It was basically what I had expected.
- **Lattices and Codes.** I thought that many of the authors went way too fast paced in this section, but I was lost halfway through almost all of the talks. Maybe I just don't have the proper prerequisites.

Can't wait for day 2!

3 Day 2

Contents

3.1 Machine Learning

3.1.1 Fast and Efficient Malware Detection with Joint Static and Dynamic Features Through Transfer Learning

Malware detection is a critical but very challenging task. The authors developed an accurate and fast malware detection model that improve the detection of malware through the 1D-CNN extractor to learn about the patterns in static and dynamic features of malware. They also developed a knowledge distillation to transfer rich knowledge from a large Teacher model (static and dynamic) to a smaller student model (only static), allowing them to benefit from dynamic analysis without the long delays previously necessary.

3.1.2 Efficient Network Representation for GNN-based Intrusion Detection

The authors introduced a Graph Neural Network (GNN) based framework to identify malicious communication flows. It has three parts: A Graph Structure-Agnostic (responsible for embedding the node attributes to extract the most relevant information), an Attention-based Feature extractor (exploits the embedded features generated by the GSA layer to aggregate neighbors' data while assigning an importance score to each one of them), and a Spatial Feature Extractor (extracts the spatial information from the graph using a Convolutional Graph Network GCN). Then, the authors show that their framework is better than classical ones using many different comparison tools.

3.1.3 EVADE: Efficient Moving Target Defense for Autonomous Network Topology Shuffling Using Deep Reinforcement Learning

EVADE is a defense that periodically changes a network topology to thwart potential attackers for protecting a given network. The authors improved this defense through a fractal-based environment (FSS) that can significantly reduce the training complexity of our DRL algorithms, a vulnerability-aware ranking algorithm (VREN) to strategically adapt edges for efficient and effective network configurations, and a density optimization (DO)-based greedy algorithm to further reduce the search space for DRL algorithms.

3.1.4 Steal from Collaboration: Spy Attack by a Dishonest Party in Vertical Federated Learning

The authors explore and formulate a new privacy leakage path in VFL, proposing two novel effective methods of spy attacks for the cases where the adversary is the active party and the passive party. The authors also discuss four possible defenses against point out their weakness which highlights the need for designing advanced defense strategies.

3.2 Side-Channel and Fault Attacks

3.2.1 HS-based Error Correction Algorithm for Noisy Binary GCD Side-Channel Sequences

The authors proposed an error correction algorithm:

1. Proposed Z-encoding to efficiently enumerate candidates using the same Expand as the Heninger-Shacham's CRYPTO2009 method.
2. Introduced two types of loss functions: likelihood-based and norm-based

Then, they evaluated and confirmed the high performance of their algorithm.

3.2.2 Formal Verification of Arithmetic Masking in Hardware and Software

Masking protects implementations against physical side-channel attacks, an area growing in demand to the rise of PQC that's caused modern research to be centered around efficient Boolean masking schemes for well-known symmetric cryptographic algorithms. The authors created the first formal verification approach for arithmetic and Boolean masking, and find many flaws in classical implementations that would fail in PQC.

3.2.3 Divide and Rule: DiFA-Division Property Based Fault Attacks on PRESENT and GIFT

The authors applied a bit based division property with fault attacks, introduced the first bit invariance property, used the bit invariance property to reduce the key space size, and discussed about the even-nibble property of GIFT-128.

3.2.4 Layered Binary Templating

The authors present a new and efficient attack technique called layered binary templating attacks (LBTA), leading them to discover first-come-first-serve data placement and data deduplication during compilation. Most notably, the compiler can introduce spacial granularity between sensitive data. Finally, the authors exploited this, concluding that chromium-based applications were the most susceptible.

3.3 Embedded Security

3.3.1 QuantumCharge: Post-Quantum Cryptography for Electric Vehicle Charging

With the rise of quantum computers, there is a growing threat to EV charging security. The authors introduced QuantumCharge, which is an extension to ISO 15118, a transition to PQC, Crypto-agility, EV keys in HSM, and Signature-based. Lastly, the authors produced a proof-of-concept implementation and used Tamarin security verification.

3.3.2 A Forkcipher-based Pseudo-Random Number Generator

The authors took a well-tested PRNG (CTR-DRBG), rebuilt it with a new primitive (Forkcipher), and improved it: optimized internal components, increased it's sped by 33

3.3.3 DMA'n'Play: Practical Remote Attestation Based on Direct Memory Access

The authors created DMA'n'Play, which allows remote attestation using DMA and direct monitoring of memory contents, removing several practical limitations on attestation schemes. They also produced an external device DMA'n'Play ToGo to relay measurements to external entity, and integrated their idea into two real-world devices (a drone and a syringe).

3.3.4 Recommendation for a holistic secure embedded ISA extension

After studying glitching defenses and control-flow/memory integrity, the authors create a design to deter against the threats they studied. Additionally, the authors implemented their recommendations using the gem5 simulator system, and evaluated it using the MiBench2 benchmarks.

3.4 Elliptic Curves and Pairings

3.4.1 Pairings in Rank-1 Constraint Systems

The author is the first to efficiently implement pairings as a SNARK arithmetic circuit through the R1CS model. With many optimizations, the author implements his ideas in gnark for the curves BLS12-377 and BLS24-315.

3.4.2 Binary Kummer Line

The author increases the security and efficiency of scalar multiplications using the binary Kummer Line through a concrete binary Kummer line: BKL251. BKL251 provides 128-bit security, and is faster than both BEd251 and CURVE2251.

3.4.3 Generalized Asynchronous Remote Key Generation for Pairing-based Cryptosystems

Asynchronous Remote Key Generation (ARKG) allows for a party to create public keys with corresponding private keys that can only be later computed by an intended party only. The author proposes Asymmetric Key Generation (AKG) and an extension ϕ-AKG, a novel generic approach for building ARKG schemes. The author also constructs ARKG schemes for many popular primitives, and then benchmarks them.

3.5 Isogeny-Based Cryptography

3.5.1 Efficient Isogeny Proofs Using Generic Techniques

The authors showed that generic zk-SNARKs are a viable tool for primitives with little structure. The main idea was ℓ-isogeny path finding relation $\rightarrow$ polynomial system of equations $\rightarrow$ R1CS instance $\rightarrow$ generic proof system $\rightarrow$ isogeny NIZK-PoK. So [ID protocol, VDF, Signature(?), $\ldots$] = [OWF] + [zk-SNARK]. Additionally zk-SNARKs are practical but are hard to assess without accurate benchmarking tools, and there is still missing literature on the non-malleability via Fiat-Sharmir/BCS transform.

3.5.2 Practical Robust DKG Protocols for CSIDH

The talk started with Hard Homogenous Spaces (HHS) and cryptography/DKG on it. Then, the authors introduced two DKG protocols that worked with Shamir secret sharing, and are shown to have both lower isogeny computations and communication cost compared to eEtended/Structured CSI-RAShi.

3.5.3 Low Memory Attacks on Small Key CSIDH

CSIDH is a secure PQ KEP with many appealing properties, but for it to be efficient one needs small keys. The authors introduce Restricted Effective Group Actions (REGA) to connect CSIDH and Diffie-Hellman, a REGA-based Diffie-Hellman protocol, and adapt several techniques to the REGA-DLOG$_m$ setting to create more efficient time-memory tradeoffs.

3.6 General Thoughts

Here is a ranking from best to worst in terms of section.

- Elliptic Curves and Pairings. This is where my interests were, and I'm really satisfied and happy with all three talks. I had read Pairings in Rank-1 Constraints a few months ago and got really stuck, so it was nice to come full circle after a few months of forgetting about that paper and realizing that I now had the mathematical ability to understand it. The other two talks were amazing, too. Favorite section of the day by far.
- Embedded Security. I really like the electrical vehicle talk, and I had better prerequisites for this section so I could understand the content a lot better than the other ones. I disliked the recommendation talk, but the other 3 were quite good.
- Isogeny Based Cryptography. I really like the efficient isogeny proofs talk, and the other two weren't terrible either. Only downside was that it was the last session of the day and I was feeling a bit tired.
- Machine learning. I'm not a huge fan of machine learning, but I don't hate it either. This section was what I expected, so I wasn't a huge fan

nor did I hate it.

- Side-Channel and Fault Attacks. I lacked prerequisites for this section and was frequently lost.

Can't wait for day 3!

4 Day 3

Contents

4.1 Privacy Preserving Protocols

4.1.1 Constant-Round Multiparty Private Function Evaluation With (Quasi-)Linear Complexities

PFE is a special case of MPC but current PFE protocols are unpractical due to large round complexity. The authors proposed the first time time constant-round multiparty PFE protocols secure against any number of corrupted parties under the semi-honest security model. The authors have two constructions: the first from oblivious evaluation of switching network (OSN) protocol and the second based on a single homomorphic encryption. Additionally, the authors optimized their constructions with half-gate technology.

4.1.2 A Framework for UC Secure Privacy Preserving Biometric Authentication using Efficient Functional Encryption

The authors modeled the first privacy preserving biometric based 2FA as an ideal functionality in universal composability and proposed a general protocol that uses functional encryption and prove that it UC-realizes our ideal functionality. Additionally, they showed how to instantiate their framework with efficient, state of the art inner-product functional encryption (allowing the computation of the Euclidean distance/Hamming distance/cosine similarity between encrypted biometric templates), implemented it, and evaluated its performance.

4.1.3 Predicate Private Set Intersection With Linear Complexity

Private Set Intersection (PSI) enables two parties to learn the intersection of their input sets without exposing other items outside the intersection. The authors created a new primitive called Predicate Private Set Intersection (PPSI) that considers the setting where each item in the input set has an associated payload, and the desired output is a subset of the intersection obtained by evaluating certain conditions over the payload. Additionally, the authors implemented an efficient PPSI protocol and evaluated its performance.

4.1.4 Private Information Retrieval with Result Verification for More Servers

In an multi-server information-theoretic PIR with result verification (PIR-RV) model where the client can detect the existence of malicious servers even if only one server is honest, the authors constructed a PIR-RV protocol and showed discussed its security.

4.2 Homomorphic Cryptography

4.2.1 PIE: p-adic Encoding for High-Precision Arithmetic in Homomorphic Encryption

The authors applied p-adic number theory techniques to construct a generic rational encoder which is compatible with HE. This included a new coding scheme PIE that can be combined with both AGCD-based and RLWE-based HE to perform high precision arithmetic. Additionally, the authors demonstrated how to attach PIE to the AGCD-based IDGHV scheme and the RLWE-based (modified) Fan-Vercauteren scheme, and compared their work to previous work.

4.2.2 Analysis and Prevention of Averaging Attacks against Obfuscation Protocols

The authors concluded that averaging attacks perform exceptionally well if obfuscation is based on random values sampled independently for each query, and showed an attack analysis able to be generalized to all protocols that employ probabilistic output obfuscation. Additionally, the authors proposed the paradigm of Data-Dependent Deterministic Obfuscation (D3O) that effectively prevents averaging attacks, proposed a construction, demonstrated its practicality, and evaluated its performance.

4.2.3 FLSwitch: Towards Secure and Fast Model Aggregation for Federated Deep Learning with a Learning State-Aware Switch

The authors proposed FLSwitch, a secure and fast FL solution consisting of three novel components, a new secure aggregation protocol based on the Pailliar HE and a residue number coding system outperforming the state-of-the-art HE-based solutions, a fast FL aggregation protocol with an extremely light overhead of learning on ciphertexts, and a learning state-aware decision model to switch between two protocols during an FL task. Finally, the authors demonstrated the applicability of FLSwitch.

4.3 Keynote II

4.3.1 Language-enforced Data Confidentiality against Memory Disclosure and Transient Execution Attacks

Memory disclosure vulnerabilities (which can quickly turn into remotely exploitable vulnerabilities) and transient execution attacks are two increasingly important threats.

- Memory disclosure vulnerabilities, which leaks sensitive process data, has become an increasingly important threat because previous methods are becoming harder. While control flow hijacking has become more difficult as exploit mitigations technologies become more advanced, the new concern is how computer memory is handled. Even with exploit mitigations, hackers can still find ways to access and steal valuable data from computer processes through memory disclosure vulnerabilities.
- Transient execution attacks leak otherwise inaccessible process data through residual micro-architectural side channel attacks. Current technology is insufficient in defending against these attacks and its variants, and in-process and cross-process leakage of sensitive application data continues.

However, there are several solutions:

- Language-enforced data confidentiality is practical. We defend against software vulnerability exploitation by continuously finding and fixing bugs, and we can defend against software and hardware vulnerability exploitation by retrofitting memory safety to C/C++ and rewriting critical components in Rust/Go. Both of these are very doable.
- In-memory encryption offers future proof protection against memory disclosure and transient attacks. There are many open source prototypes for this. One example of this is DynPTA, a selective data protection approach that combines static analysis with scoped dynamic data flow tracking (DFT) to keep a subset of manually annotated sensitive data always encrypted in memory. Soon, another open source prototype, LeakLess will be released.

To summarize, memory disclosure vulnerabilities and transient execution attacks can be defended against, but there is still a lot of work to be done.

4.4 General Thoughts

Here is a ranking from best to worst in terms of section.

- Homomorphic Cryptography. This was the section I was the most excited for the day and it delivered. I liked the p-adic talk a lot: p-adics were an area of math that I've always found cool, and never knew that you could apply that to homomorphic cryptography until today. The other two talks weren't as great but they were still pretty good.
- Keynote II. The talk was presented very well and the presenter is obviously very intelligent, but it's not an area that I'm particularly interested in. I'm not that interested in software/hardware, and more about the math. All of this was knew to me and it was presented in an interesting manner, but it's just not an area that I'm super interested in.
- Privacy Preserving Protocols. I didn't really like this section, as I kept getting lost and none of the talks really caught my eye.

Can't wait for day 4!

5 Day 4

Contents

5.1 Encryption

5.1.1 On the Complete Non-Malleability of the Fujisaki-Okamoto Transform

Malleability is a concern for Fujisaki-Okamoto Transforms, which turn weakly secure PKEs into strongly secure KEMs in the ROM. The authors introduced three game-based security notions NM-CPA*, NM-CCA1*, and NM-CCA2* for KEM schemes, attacks against 2 of 4 FO transforms, and displayed a relationship between NM-ATK* PKE and NM-ATK* KEM.

5.1.2 Anonymous (Hierarchical) Identity-Based Encryption from Broader Assumptions

Identity-Based Encryption (IBE) allows a sender to encrypt messages to a receiver without knowing the receiver-specific public key, but only using the receiver's identity and a master public key that is small, i.e., polynomial in the security parameter. The authors contributed a more efficient anonymous IBE from (relaxed) blind Chameleon encryption, the first anonymous hierarchical IBE from Computational Diffie-Hellman (CDH) and first anonymous (hierarchical) IBE from ϕ-hiding, an anonymous (hierarchical) IBE from lattices, and new tools for non-blackbox anonymous (hierarchical) IBE (Blind CE, Blind Hash Encryption, Blind One-Time Signature with Encryption).

5.1.3 Optimal Security Notion for Decentralized Multi-Client Functional Encryption

Constructions for (Decentralized) Multi-Client Functional Encryption ((D)MCFE) is an active area of research. In short, the authors criticize current (D)MCFE schemes for the class of inner products and propose a stronger security notion. Additionally, the authors showed an optimal proof for specific functionalities and a feasibility result for IP-DMCFE.

5.1.4 Publicly Auditable Functional Encryption

Suppose we have a client who does not want to disclose his data, a server that needs to perform functional evaluation on uploaded data, and an external 3rd party that wants to verify that the data complies with policies. The authors solved part of this problem by introducing publicly auditable FE definitions, as well as constructions from FE, commitments, and NIWIs (and UD Multi Input FE construction).

5.2 Advanced Primitives

5.2.1 Robustly Reusable Fuzzy Extractors in a Post-Quantum World

A Fuzzy Extractor is a pair of algorithms used to generate and recover keys from fixed, noisy, entropic sources. We care about this because of authentication: PUFs and biometrics. In this talk, the authors showed that the Hash Fuzzy Extractor is secure in the QROM with somewhat reasonable parameters for appropriate choices of underlying Secure Sketch, such as Syndrome Code-offset. Additionally, appealing to the QROM rather than using a Standard Model Robustly Reusable Fuzzy Extractor significantly improves parameters required for 128-bit security.

5.2.2 GeT a CAKE: Generic Transformations from Key Encapsulation Mechanisms to Password Authenticated Key Exchanges

Password Authenticated Key Exchange (PAKE) are important and there is now a desire to turn this into standardized post-quantum encapsulation mechanisms (KEM). The authors proposed two new generic and natural constructions in the UC model to turn a KEM into a PAKE. In short, the authors concluded: CAKE/OCAKE use implicit/explicit authentication to create a Pake from a well-defined KEM; in the UC model, CAKE uses an adaptive corruption and erasure model, while OCAKE is static corruption. Additionally, the authors applied this to Kyber.

5.2.3 Subversion-Resilient Authenticated Encryption without Random Oracles

The question the authors aimed to answer was: Is it possible to achieve authenticated encryption without random oracles in the offline watchdog model? They proved that subversion-resilient AE is possible through canonical verification/decryption built from wPRFs, relying on symmetric random keys for recomputation.

5.2.4 Scored Anonymous Credentials

Anonymous credentials are useful for obvious reasons, but this is a problem when users misbehave. The authors overcame the unspoken global halting issue of prior works, proposing scored anonymous credentials, a new system that design storing a number of active sessions with volatile scores downgradable before finalized with proper session judging and efficiency/flexibility using verifiable shuffle.

5.3 Multiparty Computation

5.3.1 Game-Theoretically Secure Protocols for the Ordinal Random Assignment Problem

The Classical Ordinal Random Assignment Problem is stated as follow: There are n players and n items, and each player has a total preference order for the items. Devise a (random) mechanism to give a one-to-one matching satisfying certain good properties. Maximin security requires that when an adversary is acting to affect the allocation for an honest player in the protocol, the honest could always expect to get their preferred item with a "higher" chance than when the adversary is absent. The authors contribute two results:

- For $n \geq 4$ players, any mechanism that achieves both strong equal treatment and ordinal efficiency cannot be realized by a maximin secure protocol (against a fail-stop adversary) that terminates with a bounded number of rounds.

- There exists a mechanism that achieves both strong equal treatment and stability (when all players are honest) and can be realized by a maximin secure protocol against a fail-stop adversary controlling up to $n-1$ corrupted players.

5.3.2 Explicit and Nearly Tight Lower Bound for 2-party Perfectly Secure FSS

Function Secret Sharing (FSS) is an MPC technique that aims to share a function or program among multiple parties. The authors found a lower bound for the Information Theoretic FSS for an arbitrary function class $\mathscr{F}$ and give a construction very close to this bound.

5.3.3 A New Approach to Garbled Circuits

Garbled circuits are cryptographic primitives with a wide variety of applications. The authors introduced a new garbling scheme without an encrypted truth-table and more efficient complexities per gate, showing that their scheme satisfies PPT adversaries in the RO model and how it can be extended to support free XOR.

5.3.4 Multi-Theorem Fiat-Shamir Transform from Correlation-Intractable Hash Functions

The title explains what this talk is really well: the authors proposed the first Fiat-Shamir paradigm to obtain a multi-theorem adaptive NIZK relying on correlation intractable hash functions.

5.4 Blockchain

5.4.1 Mt. Random: Multi-Tiered Randomness Beacons

The authors proposed Mt. Random, a multi-tiered randomness beacon that combines PVSS and (T)VRF techniques in order to provide an optimal efficiency/quality trade-off without sacrificing security guarantees. Their tier includes:

- PVSS: large batches of uniform randomness but has quadratic complexity. Our Results: gradual release
- VRFs: cheap randomness but has (bounded) bias.
- Threshold VRFs: cheap uniform randomness but needs complex setup and reseeding. Our Results: better distributed key generation
- VDFs/TLPs: cheap uniform randomness but needs complex setup and concrete parameters are unknown.

5.4.2 Revisiting Transaction Ledger Robustness in the Miner Extractable Value Era

What happens when the miners of a transaction ledger are dictatorial? The authors introduced a new property called content preference robustness (CPR), which ensures rational liveness and provides rational transaction order preservation. The authors showed how to achieve these properties, and then showed a construction of a generic CPR ledger compiler leveraging TLPs.

5.4.3 An Empirical Analysis of Security and Privacy Risks in Android Cryptocurrency Wallet Apps

The authors analyzed the source code of wallet apps to find potential security issues:

We systematically analyse wallet apps' source codes and find potential security issues: requesting dangerous permission, embedding third-party libraries for advertising and tracking purposes, presence of malware code, and using anti-analysis techniques. They found out that

- 1.7% apps using sensitive permission such as android.permission.DOWNLOAD WITHOUT NOTIFICATION in their code which, once requested, enables the app to download any file or malware executables without user consent,
- 5.4% wallet apps embed malware code in their source code according to VirusTotal,
- 2.8% apps embed CrossLibrary Data Harvesting (XLDH) library which "illegally" extracts user information from legitimate libraries such as Facebook, Google, Twitter, and Dropbox.

5.5 General Thoughts

Here is a ranking from best to worst in terms of section.

- Multiparty Computation. Favorite section of the entire conference. Basically every talk was cool, and ordinal random assignment problem one was amazing!! I found it really hard to rank the rest in terms of preference because they were all amazing. I was pretty excited for this section and it was awesome.
- Blockchain. This section was also amazing! I really enjoyed the mt. random and the mev talk, the last one wasn't as interesting to me so I will rank this below multiparty computation. I was the most excited for this section for the day and the first talk definitely exceeded my expectations.
- Advanced Primitives. I really liked the fuzzy extractors and CAKE one, both were new ideas to me that seemed to make sense intuitively and the motivation behind both was explained really well. The other two talks were still quite interesting, but these two stood out to me.

- Encryption. Of all the talks, I only understood more than half of the FO transform talk, which was cool. The other talks were unfortunately a bit too technical for me, but I'm sure they were great.

6 Conclusion

6.1 Favorite Sections

- Web Security.
- Keynote I.
- Elliptic Curves and Pairings.
- Homomorphic Cryptography.
- Multiparty Computation.
- Blockchain.
- Advanced Primitives.

6.2 Favorite Talks

- Capturing Antique Browsers in Modern Devices: A Security Analysis of Captive Portal Mini-Browsers
- Those Aren't Your Memories and They're Somebody Else's: Seeding Misinformation in Chat Bot Memories
- Challenges and Solutions to Post-Quantum Messaging
- BIKE Key-Recovery: Combining Power Consumption Analysis and Information-Set Decoding
- A Card-Based Protocol That Lets You Know How Close Two Parties Are in Their Opinions (Agree/Disagree) by Using a Four-Point Likert Scale.
- AuthZit: Multi-Modal Authentication with Visual Spacial and Text Secrets
- QuantumCharge: Post-Quantum Cryptography for Electric Vehicle Charging
- Pairings in Rank-1 Constraint Systems
- Efficient Isogeny Proofs Using Generic Techniques
- PIE: p-adic Encoding for High-Precision Arithmetic in Homomorphic Encryption
- On the Complete Non-Malleability of the Fujisaki-Okamoto Transform
- Robustly Reusable Fuzzy Extractors in a Post-Quantum World
- GeT a CAKE: Generic Transformations from Key Encaspulation Mechanisms to Password Authenticated Key Exchanges
- Game-Theoretically Secure Protocols for the Ordinal Random Assignment Problem

- Mt. Random: Multi-Tiered Randomness Beacons

6.3 General Thoughts

Overall, this was probably my favorite cryptography conference I've attended virtually, and the reverse sleep schedule was well worth it for a few days. 3am talks are such a vibe...

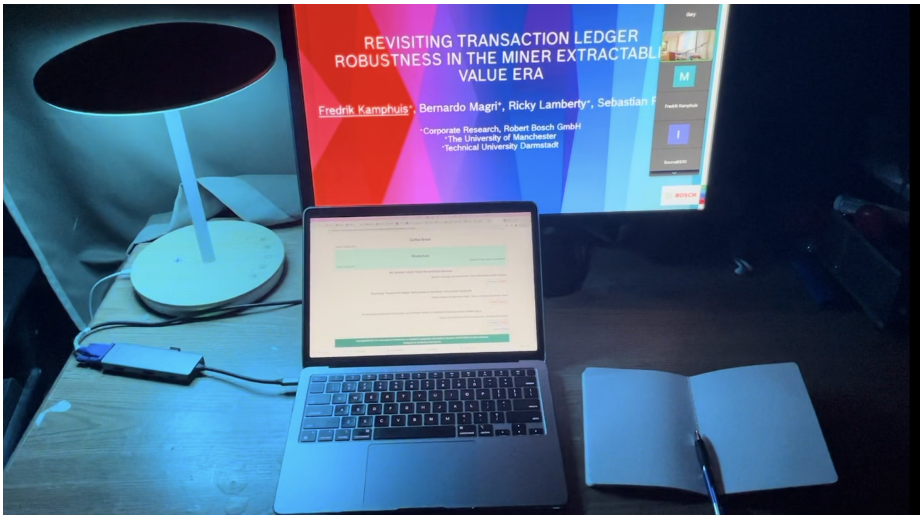
REVISITING TRANSACTION LEDGER ROBUSTNESS IN THE MINER EXTRACTABL
VALUE ERA
Fredrik Kamphuis*, Bernardo Magri+, Ricky Lamberty*, Sebastian F
*Corporate Research, Robert Bosch GmbH
+The University of Manchester
*Technical University Darmstadt
Gary
M
Fredrik Kamphuis
I
BOSCH
Coffee Break
Blockchain