ETH zürich

Zurich Information Security and Privacy Center (ZISC)

Annual Review 2023

ETH Zurich, Zurich Information Security and Privacy Center (ZISC)

Introduction

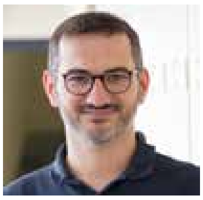

Prof. Srdjan Capkun
Chair of the center

Dr. Kari Kostiainen
Director of the center

The information security landscape is in the middle of a disruption. One reason behind this is the rapid emergence of powerful **new technologies**. To name just one popular example, **generative AI** tools have recently become significantly better at producing content. One obvious threat related to such systems is the easy production of misinformation. Another commonly acknowledged problem is that such models provide no guarantees that the produced text is correct and free of harmful content. However, there are also other, more subtle risks. For example, sophisticated adversaries can influence models and their output by manipulating the training process. Such developments lead to the current worrisome situation where many companies plan to move fast to deploy these technologies, despite the fact that the risks of such systems are still poorly understood.

Another example of new emerging technology is the prospect of **quantum computing**. While it will likely take several years before quantum computers become practical, organizations need to start migrating their IT systems to post-quantum secure implementations today. The scale and complexity of this migration task pose non-trivial technical challenges.

Another reason behind the current disruption is **geopolitical**. The wars in Ukraine and the Middle East have affected the security landscape of Europe irreversibly.

Cyber attacks against critical infrastructure, denial-of-service, and misinformation online are just a few examples of increasingly relevant threats. Also, the growing tension between the US and China has implications beyond these two countries. Export controls and "chip wars" have forced governments and companies to rethink questions related to technology sourcing and **computing sovereignty**. Organizations in Switzerland and elsewhere need to use the latest technologies to remain competitive. However, at the same time, they struggle to find solutions that give them sufficient independence.

Besides such emerging threats, organizations continue to struggle with more **traditional information security problems**: ransomware remains a major headache for companies, phishing campaigns remain prevalent, large data leaks continue to damage the reputation of companies, and the increasing complexity of IT systems makes it difficult to apply best security practices. Although these IT security problems are well understood, the existing tools and technologies are insufficient to eliminate them.

The ZISC Center

The Zurich Information Security and Privacy Center (ZISC) was established in 2003 to bring academia and industry together to address the information security challenges of tomorrow. Today, 20 years later, this mission remains more relevant than ever.

The ZISC center is formed around its eight faculty members whose research groups combined include more than 100 researchers (mainly Ph.D. students and postdocs) working in various aspects of information security, privacy, and cryptography. The center approaches its mission in the following ways:

1. We conduct **PhD-level research projects** that address the information security and privacy challenges in different ways. First, we study emerging technologies to understand their novel security and privacy risks and threats. Second, we develop new security tools and solutions with strong guarantees. Third, we tackle the fundamental open questions of information security and privacy. More information about selected such research projects can be found in the Research Highlights section of this report.

2. **We educate** the next generation of academic researchers, information security experts, and IT professionals. In addition, we support educational programs in Swiss primary schools and gymnasiums. More information about these activities can be found in the Education section of this report.

3. We create an environment where promising young researchers can launch **startup companies** and commercialize their research results. The companies founded by ZISC researchers are listed in the Startup Companies section of this report.

The main research areas of the center include the security and privacy of AI technologies, secure and sovereign computing, foundations of cryptography, future Internet architecture, secure positioning and localization, trusted execution, access control, security protocol verification, and blockchain technology. More information can be found in the Main Research Areas section of this report.

Partnership Model

The typical way to engage with the ZISC center and its researchers is a long-term partnership. ZISC partnership includes the following benefits:

1. **Customized research projects**. The main element of the ZISC partnership is PhD-level research projects that are tailored to meet the needs and interests of our partners. Every year, we conduct several such projects in collaboration with our industry partners. More information about current research collaborations can be found in the Research Projects section of this report.

2. **Expert advice**. We connect our partners to leading research experts (typically professors and postdocs) for technical discussions and practical advice.

3. **Networking**. The ZISC center organizes various events, including a bi-weekly lunch seminar with technical talks.

4. **Continuing education**. We also provide our partners access to cyber-security continuing educational programs. Two programs are available: Certificate of Advanced Studies (https://inf.ethz.ch/continuing-education/cas-cybersecurity.html) and Diploma of Advanced Studies (https://inf.ethz.ch/continuing-education/das-cybersecurity.html).

5. **OpenLab**. Our collaboration working environment called the ZISC OpenLab, allows us to host extended visits to ETH Zurich and various forms of collaboration.

Partners

The research activities of the ZISC center are supported by these partner companies

Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra

armasuisse

NEC

ZURICH

Zürcher
Kantonalbank

SIX

SWISS POST

Associate Partner

**open
systems**

ZISC Faculty Members

The ZISC center includes the following ETH faculty members:

Prof. Dr. David Basin leads the Information Security Group that performs research on methods and tools for the analysis and construction of safe and secure systems.

Prof. Dr. Srdjan Capkun leads the System Security Group, studying the design and the analysis of security protocols for wired and wireless networks and systems.

Prof. Dr. Dennis Hofheinz leads the Foundations of Cryptography group that designs and analyzes cryptographic building blocks and their use.

Prof. Dr. Ueli Maurer leads the Information Security and Cryptography Group that focuses on information security, theory and application of cryptography and theoretical computer science.

Prof. Dr. Kenny Paterson leads the Applied Cryptography Group whose research focus is on applied cryptography and communication security.

Prof. Dr. Adrian Perrig leads the Network Security Group whose research revolves around building secure and robust network systems – with a particular focus on the design of next-generation Internet architectures.

Prof. Dr. Shweta Shinde leads research in trusted computing and its intersection with system security, program analysis, and formal verification.

Prof. Dr. Florian Tramèr leads the Secure and Private AI Lab whose research currently focuses on understanding and improving the worst-case behavior of machine learning systems.

Research Highlights 2023

Machine learning models memorize more and more data

Prof. Florian Tramèr

Generative machine learning is progressing at a frightening pace!

We now have models that can generate beautiful art, and write text that is often hard to distinguish from human. This progress has been largely enabled by training models on ever larger datasets, primarily scraped from the web. Unfortunately, this sometimes means that these models will memorize some of this data, and randomly regurgitate it.

Generating not-so-fake images

In a long-lasting and fruitful collaboration with researchers at Google Deepmind, we have studied the extent to which various generative models for images and text will spit out memorized training data.
Earlier this year, we were able to demonstrate how state-of-the-art image generation models such as Stable Diffusion or Google's Imagen will sometimes generate images that are virtually identical to real images contained in the model's training data.

While we found that such memorization is very rare, it still fuels the ongoing debate about the possible copyright and privacy implications of training models on large amounts of data scraped from the Web.

Beware of synthetic data

We also caution about the use of generative models to create so-called "synthetic data", which would then hopefully be devoid of any privacy risks. As our work shows, there is a risk that some of the synthetic data created by a generative model would contain clear copies of some of the underlying private training data.

Memorization at scale

We also continued our study of the memorization abilities of large language models, building upon our earlier work that first showcased the ability to extract training data from OpenAI's GPT-2 model in 2019. We now studied much larger models (up to 6B parameters) for which the actual training data has been published (this was not the case for GPT-2: OpenAI has given some information about how the training data was collected, but we don't know exactly what's in it).

Knowledge of the true training data enabled us to do a controlled study where we could measure how model memorization varies with model size, model type, and various properties of the training data distribution.

We found that memorization gets progressively worse as models get larger, and better! So this is a problem that won't just go away with scale.

As a silver lining, we find that properly deduplicating training data is an effective (albeit heuristic) countermeasure against unwarranted memorization.

Overall, our research suggests that modern generative models have a very large capacity for memorization, and that new techniques will be needed to fully prevent models from inadvertently leaking training data.

Further information

"Quantifying Memorization Across Neural Language Models"
Nicholas Carlini, Daphne Ippolito, Matthew Jagielski, Katherine Lee, Florian Tramèr and Chiyuan Zhang
International Conference on Learning Representations, 2023
https://arxiv.org/abs/2202.07646

"Extracting Training Data from Diffusion Models"
USENIX Security Symposium 2023
Nicholas Carlini, Jamie Hayes, Milad Nasr, Matthew Jagielski, Vikash Sehwag, Florian Tramèr, Borja Balle, Daphne Ippolito and Eric Wallace
https://arxiv.org/abs/2301.13188

Researchers

Prof. Florian Tramèr
(Secure and Private AI Lab, ETH Zurich)

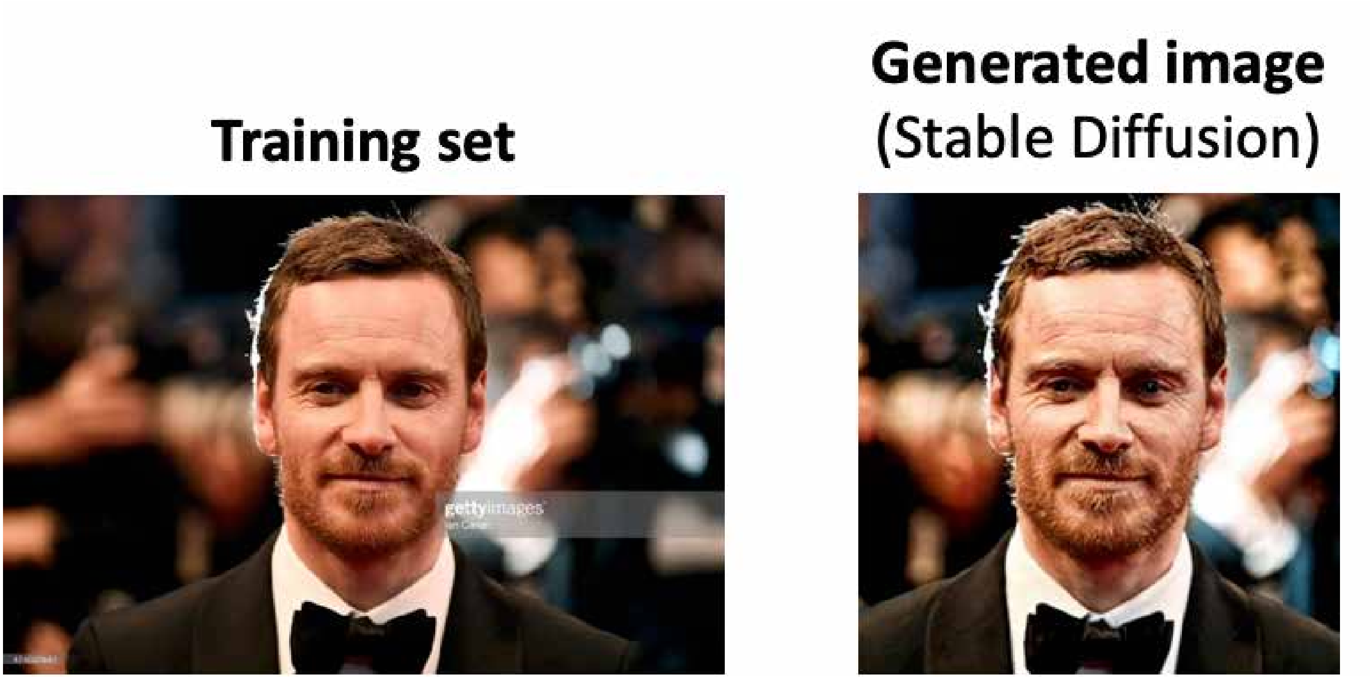

Research Highlights 2023

Fine-grained Internet Path Selection

Prof. Adrian Perrig

The goal of this project is to provide applications communicating over the Internet with **fine-grained transparency and control** regarding on-path forwarding devices. This project is a collaboration between the network security group at ETH Zürich and Armasuisse.

Some applications require communication data not to leave a given jurisdiction, which therefore requires that sensitive traffic is only routed through routers operated in certain countries or regions. Also, applications such as time synchronization services may require packets to be forwarded only over routers with hardware-based time synchronization support such as PTP. Governments or critical infrastructure operators may further desire their applications to avoid sending traffic over network equipment from untrusted manufacturers or to bypass routers running vulnerable software. They may therefore wish to **route traffic over paths consisting of well-known, trusted equipment only**.

To evaluate the usefulness of transparency and control over on-path forwarding devices on the Internet, we conducted a survey among network and security experts. The survey shows that many commercial and private customers desire detailed information about on-path routers and that a large fraction of them are also willing to pay for sending traffic over specific routers.

In its current state, the Internet does not provide end users with the necessary transparency and control regarding on-path forwarding devices to enable these use cases. In particular, the lack of network device information reduces the trustworthiness of the forwarding path and prevents end-user applications requiring specific router capabilities from reaching their full potential. Moreover, the inability to influence the traffic's forwarding path results in applications communicating over undesired routes, while alternative paths with more desirable properties are not usable.

To enable those use cases, we have designed the **flexible attestation-based routing for inter-domain networks (FABRID)**. FABRID runs on top of the SCION next-generation Internet architecture, which already provides path transparency and control at the level of links between autonomous systems and, by design, prevents attacks such as path hijacking and attacks on the integrity of routing messages. FABRID refines SCION's transparency and control by extending it to the level of routers. In FABRID, network operators announce information about the routers deployed in their network via policies, which are disseminated to the user through Internet routing messages.

Applications can learn the available policies and encode the desired policy for each network into their data packets, such that on-path routers know how to forward traffic satisfying the selected policies.

Below, we show an example, where a source end host (S) in network A sends a packet over networks B and C to a target end host (T) in network D. In a first step, networks B and C announce the policies of their internal routes to the source end host. Then, the source end host selects appropriate internal routes according to its preference and embeds the selected routes in the packet header. The packet is then forwarded along the corresponding internal routes (green dotted line) by the networks B and C.

Our design of FABRID includes the specification of an extensible policy language to describe arbitrary and customizable router attributes. Thus, **network operators do not have to reveal sensitive information about their networks**, but can instead decide themselves which router attributes they want to publish. Furthermore, trusted third parties **attest the routers and their properties** in order to prevent network operators from making false claims regarding their network infrastructure.

In the data plane, we increase goodput by minimizing the size of the packet header describing the desired policies, and apply efficient measures to ensure authenticity and secrecy of the embedded policies. To demonstrate FABRID's viability, we implement and evaluate a high-speed variant in DPDK, achieving **line rate forwarding at 160 Gbps**, and deploy a Go-based version in the global SCIONLab research testbed.

Publications

Krähenbühl, C., Wyss, M., Basin, D., Lenders, V., Perrig, A., & Strohmeier, M. (2023). FABRID: flexible attestation-based routing for inter-domain networks. In Proceedings of the USENIX Security Symposium (pp. 5755–5772).

Researchers

Prof. Adrian Perrig (ETH)
Prof. David Basin (ETH)
Cyrill Krähenbühl (ETH)
Marc Wyss (ETH)
Vincent Lenders (armasuisse)
Martin Strohmeier (armasuisse)

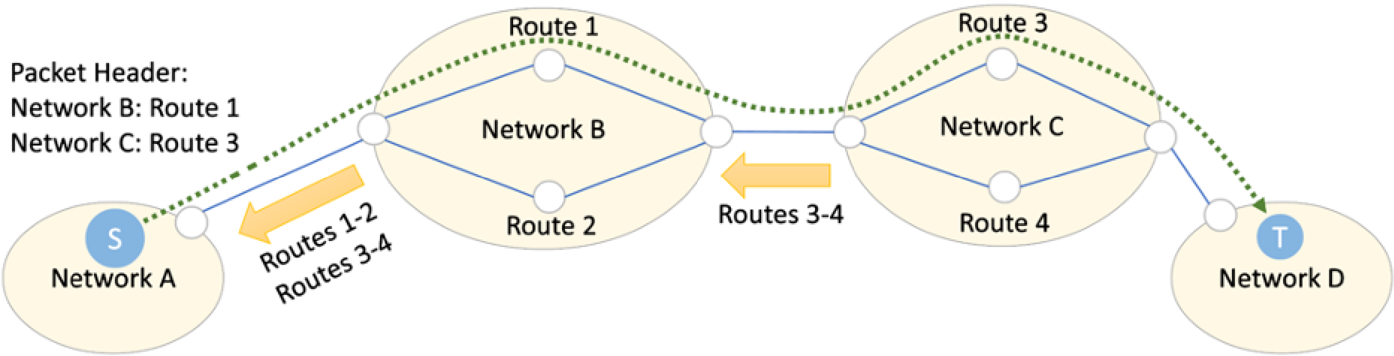

Research Highlights 2023

Dynamic trade-offs in Cryptographic Protocols

Prof. Ueli Maurer

Cryptographic protocols solve a wide variety of problems with far reaching applications, like secure e-voting, privacy-preserving machine learning, and distributed financial systems.
Given a task to carry out, one must provide a protocol (a set of instructions) guaranteeing to individual parties that, provided they follow their instructions, they will achieve the desired goal, despite some of the parties not following the instructions correctly and maybe even voluntarily cheating. A simple example of a task could be agreeing on a common random value. The guarantee in this case could be that 1) everybody gets the same value, and 2) the value is sampled from the wanted distribution, despite 3) up to half of the parties not following the instructions.

Protocol are designed with respect to a specific assumption. Formal security proofs then ensure that, whenever the assumption hold, the protocol provides a certain guarantee. Typically, stronger assumptions allow to design a protocol that achieves stronger security guarantees.

Examples of assumptions on the communications channels are that messages are delivered within some known time, or do not contain more than a certain number of errors, while examples of assumptions on the adversary are the extent to which they can force parties to deviate from the protocol, or their amount of computational power.

As soon as the assumption on which a certain protocol relies is voided, however, the protocol fails to provide the guarantee completely. For example, if the privacy of a secure computation protocol assumes that all messages are delivered within one minute, even a one second delay on a single message (maybe due to an unexpected network overload) causes the privacy guarantee to completely break down.

In real-world applications, one faces the dilemma of whether to choose a protocol providing a very strong guarantee, but relying on very strong assumption, or a protocol relying on a weak assumption and providing a similarly weaker guarantee, despite believing the stronger assumption might be satisfied most of the time!

To avoid this dilemma, we promote a different approach to protocol design: that is providing a single protocol that if some stronger assumption is satisfied, provides a stronger guarantee, but if only a weaker assumption is satisfied, still provides some (weaker) guarantee.

Protocols with fallback guarantees have been investigated in different settings. In [C89], Chaum initiated this field of research providing a multi-party computation protocol achieving unconditional security assuming an honest majority of parties and cryptographic security otherwise.

In [DHL21], [DL22] we present a multi-party computation protocol and a secure message-transmission protocol that, if the underlying network is reliable (messages are delivered within some known time) are very efficient and tolerate a high corruption threshold, but even when the network is less reliable and messages are arbitrarily delayed, tolerate some (lower) corruption threshold.

Other examples include [FHH+03], in which the authors present broadcast protocols providing different validity and consistency guarantees depending on the number of corruptions. In HLM+11] even more fine grained dynamic trade-offs for multi-party computation are provided. We will further investigate new settings in which protocols providing dynamic trade-offs are to be preferred, from a practical perspective, to protocols following a traditional design, provide new protocols and explore novel protocol design techniques.

Further information

[DL22]
Synchronous Perfectly Secure Message Transmission with Optimal Asynchronous Fallback Guarantees
Giovanni Deligios, Chen-Da Liu-Zhang
https://eprint.iacr.org/2022/1397.pdf

[DHL21]
Round-Efficient Byzantine Agreement and Multi-Party Computation with Asynchronous Fallback
Giovanni Deligios, Martin Hirt, and Chen-Da Liu-Zhang
https://eprint.iacr.org/2021/1141.pdf

[C89]
The Spymasters Double-Agent Problem: Multiparty Computations Secure Unconditionally from Minorities and Cryptographically from Majorities.
David Chaum
https://dblp.org/search?q=david+chaum+spymaster

[FHH+03]
Two-Threshold Broadcast and Detectable Multi-Party Computation
Matthias Fitzi, Martin Hirt, Thomas Holenstein, and Jürg Wullschleger
https://crypto.ethz.ch/pubs/FHHW03

[HLM+11]
Graceful Degradation in Multi-Party Computation
Martin Hirt, Christoph Lucas, Ueli Maurer, and Dominik Raub
https://crypto.ethz.ch/pubs/HLMR11

Researchers

Ueli Maurer (ETH)
Giovanni Deligios (ETH)

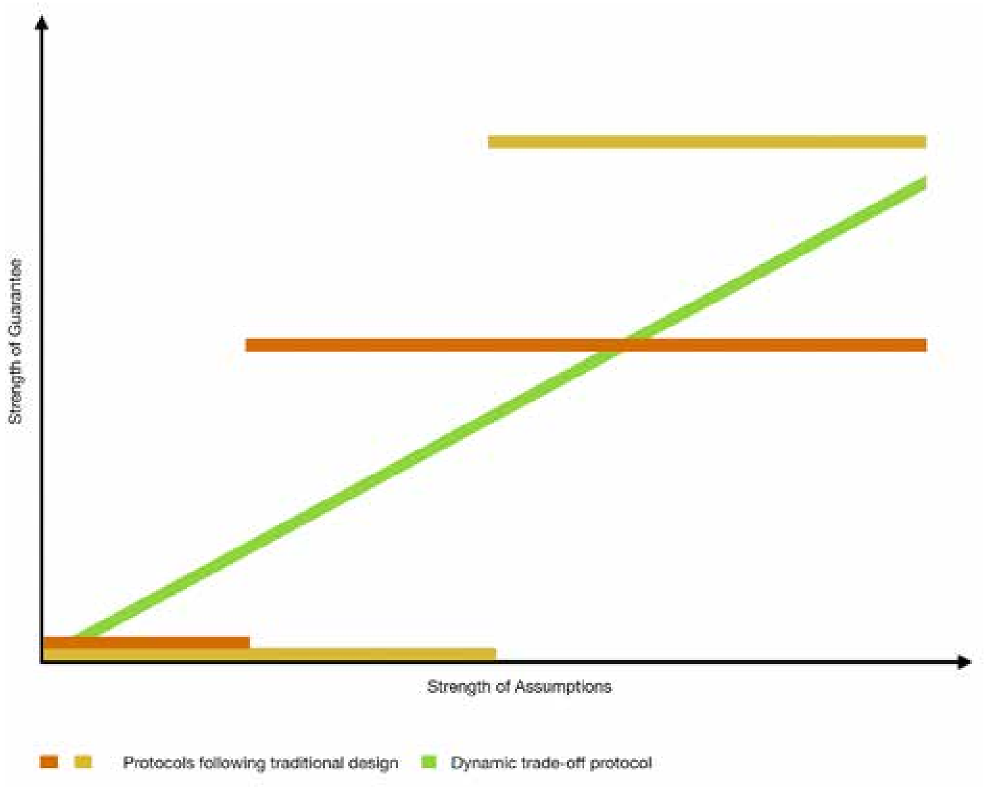

Research Highlights 2023

Privacy Analysis of Web3

Prof. Shweta Shinde

Increasing Popularity of Web3

In light of the recent buzz surrounding the Metaverse and NFTs, Web3 is steadily gaining popularity. Web3's primary objective is to decentralize the internet through the use of decentralized applications (DApps), with wallets serving as vital intermediaries between these applications and users. Popular wallets like MetaMask are now in use by millions of users. Regrettably, Web3 is frequently promoted as a more secure and private alternative to Web 2.0. However, it is important to note that decentralized applications and wallets are built on conventional technologies that do not prioritize user privacy.

Privacy Concerns

Web tracking is pervasive on the internet, involving both explicit (e.g., cookies) and implicit (e.g., browser fingerprinting) methods. While third-party cookies were once common for tracking, most modern browsers now block them by default. An alternative method gaining popularity is browser fingerprinting, which uniquely identifies users based on their browser's configuration. Unlike cookies, fingerprinting is stateless and challenging to mitigate without disabling JavaScript. DApps often include third-party tracking scripts. Once a user connects their wallet to a DApp, these scripts can access sensitive information such as the user's wallet address without prior consent, potentially compromising anonymity. Third-party scripts can also access a user's IP address, potentially linking multiple wallet addresses to their IP, thereby raising privacy concerns.

Wallet Address Leakage

We measured wallet address leakage across 616 popular DApps from DappRadar.com and 100 popular wallets from Google's Chrome Web Store. Our framework intercepted HTTP and WebSocket traffic including cookies for all of these DApps and wallets and verified whether a user's wallet address is included. Our framework identified 211 unique DApp websites (35% of the DApps) which leak the user's wallet address to third-parties. We studied the privacy policies of the top 20 third-parties and observe that 95% state that they collect the user's IP address. Hence, these third-parties can potentially link a user's wallet address to its IP address.

Web3-Based Browser Fingerprinting

We measured browser fingerprinting by visiting the top 1 Million Tranco websites as of November 8th, 2022. We found 1,325 websites calling wallet APIs out of which 1, 099 leverage these APIs to perform browser fingerprinting. We used the categorization service provided by SafeDNS to categorize these websites. We find that most websites belong to the category "Pornography & Sexuality". The most popular website performing web3-based browser fingerprinting was xhamster.com. However, we also found evidence in news websites such as The New York times but also on social media websites such as TikTok.

Effectiveness of Blocklists

Since half of the calls to wallet APIs originate from third-party sources, we evaluated the effectiveness of blocklists in blocking the detected third-party scripts. We downloaded the most up to date blocklists from Disconnect, DuckDuckGo, EasyList, EasyPrivacy, and Whotracks.me and counted how many of the encountered third-parties would have been blocked by each of these blocklists. Whotracks.me demonstrated the highest level of protection by blocking 43% of the identified third-parties, while Disconnect provided the least protection by blocking only 12% of the third-parties. We also explored the impact of combining all five blocklists simultaneously to enhance protection. Our results reveal that the combination of these blocklists resulted in blocking 56% of the third-parties, representing a 13% improvement compared to relying solely on Whotracks.me's blocklist.

Rethinking Web3 Privacy

Achieving privacy on the web is non-trivial. The information that we either expose explicitly (e.g., email address at login) or implicitly (e.g., IP address when visiting a website) to a party on the web can be used to track our habits online and result in negative consequences on our daily lives offline. Web3 leverages many design decisions from Web 2.0. Wallets could try to achieve more privacy by introducing a more fine-grained permission system. However, malicious parties could still link these different wallet addresses together by levering other information such as a user's IP address. Web3-based fingerprinting could be contained by redesigning the way decentralized applications interact with wallets. However, this would require standardization as well as DApp developers changing their current implementations.

Further information

Website: https://github.com/christoftorres/Web3-Privacy

Publication:
Is Your Wallet Snitching On You? An Analysis on the Privacy Implications of Web3. Christof Ferreira Torres, Fiona Willi, Shweta Shinde. In Proceedings of the 32nd USENIX Security Symposium (USENIX Security 2023). https://www.usenix.org/conference/usenixsecurity23/presentation/torres

Researchers

Dr. Christof Ferreira Torres (Secure & Trustworthy Systems Group, ETH)
Prof. Shweta Shinde (Secure & Trustworthy Systems Group, ETH)

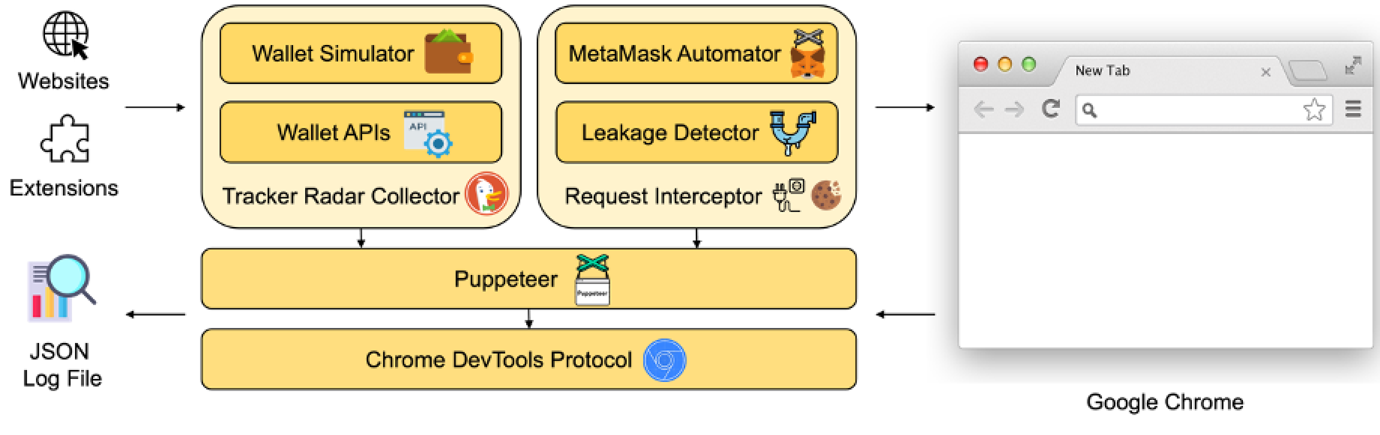

Research Highlights 2023

Security Analysis of MongoDB Queryable Encryption

Dr. Zichen Gui

Dr. Tianxin Tang

Encrypted Databases with Queries

Databases are fundamental to a wide array of services. Typically, databases are deployed on local servers or hosted through cloud services. Data breaches are a prevalent concern and pose a significant threat to many services.

To address this issue, most major database vendors now offer encryption-at-rest as a feature. This ensures that the database remains encrypted when it is not in use. However, during query operations, the server can still access the plaintext database in memory. This means that if the server is compromised, attacks such as unauthorized memory dumps become possible, presenting considerable risks when the data is in use.

Is it even possible to preserve data privacy from the server when data is in use? There is an active research area focused on encrypted database solutions that support efficient queries without resorting to plaintexts. This area is called searchable symmetric encryption (SSE) or structured encryption (STE). Encrypted database solutions built from SSE/STE can be deployed on local servers and remote cloud services. In both settings, these solutions are resilient against attacks on data in use.

Queryable Encryption

MongoDB is a popular NoSQL database system. In June 2023, MongoDB announced Queryable Encryption (QE) as a new feature in the 6.0 release. QE is the first SSE/STE scheme introduced by a major database vendor. At the time, QE was only in preview and supported only equality queries. QE with equality queries has become generally available (GA) in MongoDB 7.0. MongoDB has also promised other search functionalities in future releases.

Unlike encrypted database solutions built based on encryption-at-rest, MongoDB claims that QE achieves much stronger security. Specifically, they stated that data encrypted with QE "remains secure in-transit, at-rest, in memory, in logs, and in backups" when the research took place.[1]

The strong security claims of QE make it tempting for businesses to adopt QE as their encrypted database solution. On the other hand, the construction and the security properties of QE are also of great interest to the research community. QE offers a unique opportunity for us to study the practical aspects of encrypted database schemes. In the sections below, we give an overview of our findings on QE.

MongoDB and QE

QE was implemented as an extension to the MongoDB system. Its integration leverages MongoDB's native libraries and requires only minimal modifications to the existing system. This approach offers significant advantages, notably in saving engineering effort and reducing expenses, making the deployment of QE seamless.

MongoDB and QE: : Vulnerabilities

The approach taken by MongoDB to integrate QE also has drawbacks. In particular, it means that QE will interact with other MongoDB components when in use. These interactions have not been studied in the literature before (as all schemes were standalone constructions) and they may leak more information than what they intended. Indeed, we identified vulnerabilities arising from the interplay between QE and MongoDB's logging system.

To provide some context, MongoDB's logging system produces a querylog, storing transaction details for debugging purposes, and an oplog, tracking write operations to maintain database consistency. Our analysis reveals that these logs contain statistical information about the queries and data. Alarmingly, they can be exploited in real-world scenarios! We demonstrated the vulnerability experimentally by launching an inference attack against US census microdata encrypted with QE.

[1] https://web.archive.org/web/20220608105422/https://www.mongodb.com/products/queryable-encryption; MongoDB no longer claims security "in memory, in logs, and in backups" on the new version of the website: https://www.mongodb.com/products/capabilities/security/encryption.

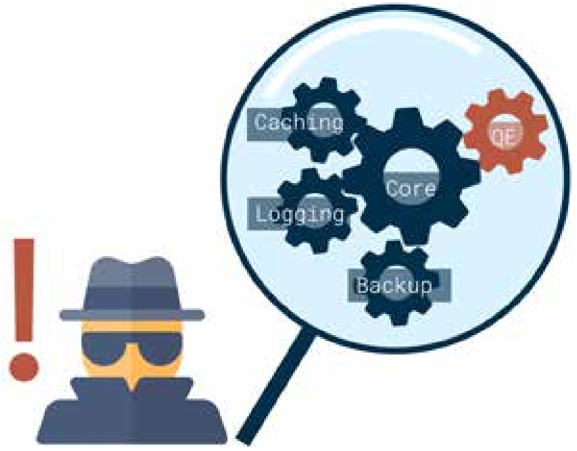

With the knowledge of auxiliary statistical information, we are able to recover the encrypted data with high accuracy. This shows that QE's security claim, data encrypted by QE "stays secure in transit, in memory, in logs, and in backups", simply does not hold up. It highlights the importance of understanding interactions between system components when designing a secure encrypted data system.

Countermeasures

It is tempting to believe that a quick fix to QE's issues is on the horizon. However, simple countermeasures to address all vulnerabilities seem elusive. A straightforward countermeasure, like disabling the logs, is only effective for the querylog we exploited. Indeed, MongoDB deactivated this specific log for QE in their 7.0 release. However, this solution is not feasible for the oplog, essential for maintaining database consistency. An alternative is encrypting the oplog, but this approach introduces another layer of complexity. Since the server requires access to the oplog to synchronize databases, it must be able to access the encrypted log. And if we grant the server access to the decryption key, we end up back to square one --- the server can recover and access the plaintext database again. These two approaches seem to be the only solutions for handling the logs. Effectively addressing these log-related and other integration vulnerabilities should be a focal point of future research on encrypted databases.

Lessons for System Engineers

From our analysis, we learned that MongoDB overlooked the interaction between QE and its system components when it first introduced QE. The consequence is that QE is vulnerable to attacks through these system interactions. This observation extends beyond this specific scenario, offering a broader lesson for engineers tasked with designing and maintaining a system.

When considering adding a new security component, complications that potentially introduce security risks emerge, especially concerning its interactions with the existing system. A significant concern is that if a security component is found to be insecure and needs to be updated, there is a possibility that architectural changes are required. This was evident with MongoDB 7.0, which introduced breaking changes in QE, rendering it incompatible with version 6.0. This implies that when system engineers decide to incorporate a new security component that has not been properly scrutinized, they may find themselves in the position of dealing with the overheads of updating the architecture of the existing system if breaking changes have been introduced to the security component later.

In addition, the way the new security component interacts with the existing system can be intricate and easy to overlook. Our security analysis of MongoDB's QE highlights this point: vulnerabilities often reside in the interactions. Therefore, when integrating a security module, system engineers must be educated to meticulously attend to these interaction details, be aware of potential risks, and actively consult with security specialists. Simultaneously, during security audits, it is vital to scrutinize the implementation, particularly the interplay between the new and existing components.

Lessons for Business Customers

For business customers, our analysis of QE underscores a critical concern: even though emerging security products aim to enhance security, potential flaws in their design or implementation can unintentionally compromise it. When deploying a security product that claims to address specific concerns, it is essential to understand that regardless of how promising its security claims may appear, intrinsic risks still apply. Our findings further highlight the importance of security audits. These audits are instrumental in identifying vulnerabilities and security risks. Deploying a fresh security product without such vetting carries inherent risks.

Further information

Publication: "Security Analysis of MongoDB Queryable Encryption". In Proceedings of the USENIX Security Symposium (USENIX Security, 2023). https://www.usenix.org/system/files/usenixsecurity23-gui_1.pdf

Artifact website: https://gitlab.com/mongodbqe/mongo

Researchers

Dr. Zichen Gui (Applied Cryptography Group, ETH Zurich)
Prof. Kenneth G. Paterson (Applied Cryptography Group, ETH Zurich)
Dr. Tianxin Tang (Applied Cryptography Group, ETH Zurich)

Research Highlights 2023

Time for Change: How Clocks Break UWB Secure Ranging

Prof. Srdjan Capkun

Many security-relevant applications require secure and reliable means to bound the physical distance between two devices, even in the presence of an external attacker. Some examples are mobile payments, Passive Keyless Entry-and-Start Systems (PKES) of cars, and other access control use cases. In these applications, an attacker is usually interested in shortening the distance measured by the devices, either to bypass a security check or gain access. For example, if an attacker convinces a PKES-enabled car that the keys of its owner are close, the car will unlock.

Ultra-Wide Band (UWB) is a secure-ranging technology widely deployed in consumer devices. UWB devices measure their distance with sub-decimeter accuracy from the time-of-flight of the radio signals they exchange. Chips following the IEEE 802.15.4z standard are currently available in many smartphones (e.g., iPhone, Samsung Galaxy S21+) and gadgets (e.g., AirTag, HomePod, Apple Watch, and Samsung SmartTag+). A new standard, IEEE 802.15.4ab, is under development and promises to achieve better ranging performance in challenging environments, which is going to increase UWBs reliability and enable new applications.

However, designing secure ranging systems is challenging, and several attacks are possible at the physical layer. Although existing research has analyzed the security of UWB ranging extensively, one important aspect has been neglected so far: real devices do not have ideal time references, but rely on inaccurate clocks. To obtain accurate distance measurements, the chips must compensate for clock errors whenever they receive a signal. In a recent publication at USENIX Security 2023 [1], we show that compensation mechanisms supposed to prevent measurement errors can also be abused to cause them; we demonstrate two novel distance reduction attacks, namely “Mix-Down” and “Stretch-and-Advance”, which affect the current UWB standard (IEEE 802.15.4z) and the draft of the future one (IEEE 802.15.4ab). Additionally, we propose countermeasures to enable secure ranging secure in the presence of non-ideal clocks.

The first attack, Mix-Down, exploits the explicit compensation of clock errors based on the carrier frequency offset of the incoming message. We show that a simple analog circuit can manipulate this estimate and induce distance reductions of several meters. The attack works even if the packets contain unpredictable pulse sequences, a common security feature in secure ranging protocols. We conclude that, at present, Mix-Down can only be prevented by performing additional, bi-directional measurements and changing the way the time of flight is computed.

In contrast to Mix-Down, Stretch-and-Advance is only possible if the messages exchanged by ranging devices are very long (i.e., several milliseconds), as in the proposal for IEEE 802.15.4ab. Even small clock inaccuracies produce large variations in the duration of long messages, and a receiver must take this into account. This tolerance allows a sophisticated attack, in which the attacker transmits a specially crafted signal containing parts of the legitimate ranging message. On a high level, this message convinces the receiver that the transmitter’s clock is too slow, which opens a window for considerable distance reduction attacks. As IEEE 802.15.4ab is still under development, and compliant hardware did not exist at the time of writing, we developed a theoretical analysis. We show that distance reductions of dozens of meters are possible, unless transceivers implement specific countermeasures that we propose.

Future work on the design, analysis, and application of ranging systems will have to consider the consequences of non-ideal clock references that we have investigated in our research. While we have focused our analysis on UWB, our results are generally applicable to time-of-flight-based ranging.

Further information

Publication:
[1] Claudio Anliker, Giovanni Camurati, Srdjan Čapkun "Time for Change: How Clocks Break UWB Secure Ranging". 32nd USENIX Security Symposium, USENIX Security 2023, August 9–11, 2023, Anaheim, CA, USA, 2023.

Researchers

Claudio Anliker (ETH)
Dr. Giovanni Camurati (ETH)
Prof. Srdjan Čapkun (ETH)

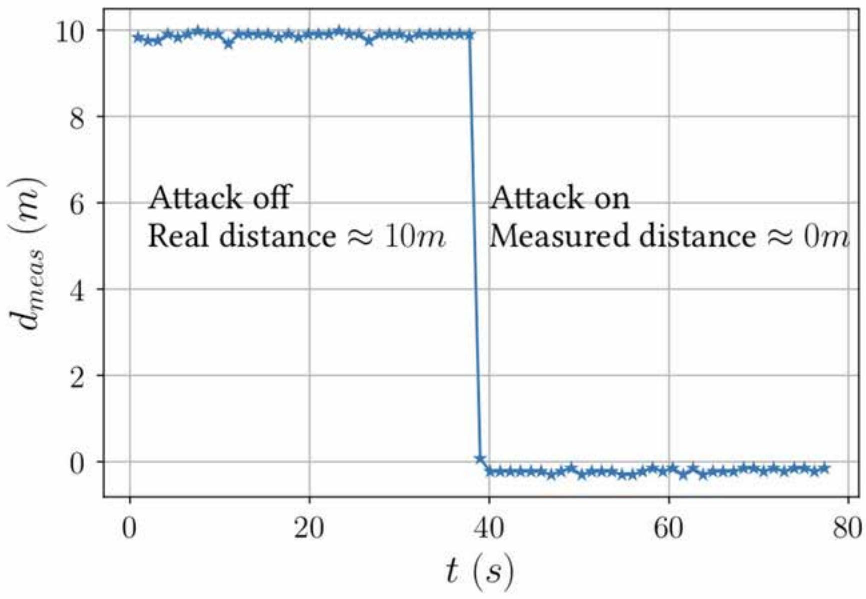

Research Highlights 2023

Studying the security of HMAC

Matteo Scarlata

Cryptographic Protocols

Cryptography plays a fundamental role in our daily digital interactions, even though it usually operates discreetly behind the scenes. Whether we're browsing the internet, sending an SMS, or unlocking our cars, these actions rely on "cryptographic protocols": communications protocols which make use of cryptography to achieve some desired security goal(s).

TLS, the protocol protecting our web browsing (recognizable by the familiar "padlock" icon in our browsers), employs encryption algorithms, authentication codes, digital signatures and key exchange mechanisms to deliver the guarantee that the information we see (or input) in a webpage travels securely all the way to the intended servers, and that nobody on the path between us and the servers can read or modify that information.

These protocols are, in the best cases, carefully designed and studied for years before being adopted en-masse. Analyses of the protocols help relate the security of this "ensemble" of cryptographic primitives to the security of the primitives themselves, and establish confidence that a protocol will stay secure in the years to come. However, even the most detailed analysis can miss some detail – and gaps between the stated security goals and the proved properties emerge.

A Gap in HMAC

Our research focused on one such gap. A very common primitive used to build cryptographic protocols – namely HMAC, "Hash-Based Message Authentication Code" (Figure 1) – was assumed in many security analyses in the literature to have some properties that had never actually been proved.

HMAC is omnipresent in Internet protocols: it is used in TLS, Signal, SSH, and Wireguard, and as a building block of HKDF, it is often employed in the critical step of deriving cryptographic keys from some input key material. Should HMAC's security have proven inadequate, the security guarantees of numerous applications could have been undermined.

When Messages Are Keys...

The precise notion of security that HMAC was assumed to achieve is named "Dual-PRF Security". In short, PRF Security of a function f(key, message) guarantees that the output of the function will "look random" (in a cryptographic sense), independently of the message, provided that a secret key is given in input.

Figure 3 contains the formal definition of a PRF security "game": an adversary can call Fn, and receives either a real PRF output (d=1) or a randomly sampled output (d=0).

The function F is PRF secure if the adversary above can't tell the two apart – that is, if it has negligible probability of guessing whether d is 0 or 1.
The "Dual-PRF" notion is stronger: we want the output to "look random" both when the function is called with a key in the first argument (as before, f(key, message)), and when the function is used in a "swapped" fashion, with the key in the second argument (i.e. f(message, key)).

HMAC had been proven a secure PRF, but no proof of its Dual-PRF security had ever been given. Worse than that – some trivial attacks show that HMAC cannot hope to achieve Dual-PRF security in a general sense: the way keys are padded (see Figure 2) means that it is easy to construct different messages for which the output stays the same in the swapped fashion, negating the needed "random-looking" property of HMAC.
Yet, HMAC is commonly used as a Dual-PRF – so should we be worried?

...And When Keys Are Keys

Further investigation showed that not only had HMAC not been proven secure in the Dual-PRF setting, but also that the proofs in the literature for its PRF security only hold for "full-length keys".

If we look at the HMAC definition in Figure 1, we can see that it only defines the behaviour of HMAC on keys that are block-sized, where the block size depends on the underlying hash function H. But the HMAC standard actually allows arbitrary length keys:

$$\mathsf{HMAC}_b(K_b, M) = \mathsf{H}((K_b \oplus \texttt{opad})\|\mathsf{H}((K_b \oplus \texttt{ipad})\|M))$$

Figure 1 "Textbook" HMAC definition, for a hash function H, and assuming block-sized keys.

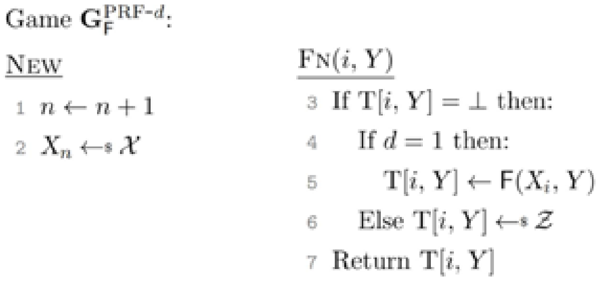

```
Game $\mathbf{G}_{\mathsf{F}}^{\mathrm{PRF}\text{-}d}$:
NEW
  $n \leftarrow n+1$
  $X_n \leftarrow\!\!\$\; \mathcal{X}$
FN$(i, Y)$
  If $\mathrm{T}[i, Y] = \perp$ then:
    If $d = 1$ then:
      $\mathrm{T}[i, Y] \leftarrow \mathsf{F}(X_i, Y)$
    Else $\mathrm{T}[i, Y] \leftarrow\!\!\$\; \mathcal{Z}$
  Return $\mathrm{T}[i, Y]$
```

Figure 3: PRF security game. A function F(X,Y) is "Dual-PRF" secure if both F and F'(X,Y) = F(Y, X) are PRF-secure

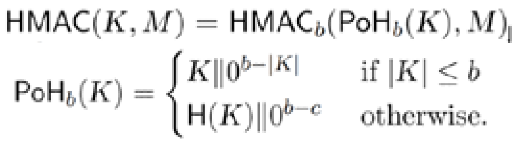

$$\mathsf{HMAC}(K, M) = \mathsf{HMAC}_b(\mathsf{PoH}_b(K), M)$$

$$\mathsf{PoH}_b(K) = \begin{cases} K \| 0^{b-|K|} & \text{if } |K| \le b \\ \mathsf{H}(K) \| 0^{b-c} & \text{otherwise.} \end{cases}$$

Figure 4: "Full" HMAC definition, with a PoH function used to preprocess input keys

as Figure 4 shows, the key is first passed through a "Pad-or-Hash" (PoH) function that zero-pads the key if it's shorter than the block size, or hashes it if it is longer.

Proofs in the literature do not take the PoH function into account. To our surprise, this does not cover the use of HMAC in practical protocols: the typical required "full" length of an input key is 512 bits, with many applications only providing 256 or 128 bits keys, requiring PoH to be applied. The gap widens!

Provable Security to The Rescue

What guarantees do we actually have from HMAC? Can our findings lead to new attacks? To answer these questions, we decided to start afresh and try to prove the security of HMAC in the Dual-PRF setting, and with arbitrary length key inputs.

The detailed proof steps involved decomposing HMAC into its prime components – sequences of compression functions (small "h" in the picture), chained in "cascades", where the output of one is the input of the next – and rethinking the assumptions on these components.

Our security proof involves "game hopping": a method for proving the security of a scheme by considering the probability of success of an adversary in subverting the security goal in a sequence of games. We move from a "real world", in which the scheme is actually executed, to a "random world", in which the outputs of the scheme are replaced with random strings, and it is obvious that no adversary can succeed. For each step, we gradually replace the computation with random sampling of the results, and we bound the probability that the attacker can distinguish the changes in each step. If the bounds we prove are small, then combining all the steps we can infer that no adversary has a large advantage in breaking security in the real world.

Our proof allowed us to confirm that HMAC is indeed secure as a variable-key length PRF, and we managed to do so with stronger bounds than in the previous literature and without using strong additional assumptions.

When it comes to Dual-PRF security, the attacks show that we cannot achieve it. However, we proved that there exists a class of "safe" messages for which HMAC is Dual-PRF secure, and all the practical applications fall into this class.

Indeed, HMAC was used by applications (such as TLS, KEMTLS and Post-Quantum Wireguard) as a Dual-PRF to "combine" two keys – that is, to obtain f(key1, key2), where key1 could actually be known to an adversary and all the "randomness" must come from key2. In this setting the message (or "key1") input is of fixed length, which prevents our trivial attacks. We show that in general, for fixed length "key1" inputs HMAC achieves Dual-PRF security.

Key Combiners on The Horizon

With Post-Quantum security being today a hard requirement for new protocols, and "hybrid" schemes that combine Post-Quantum and classical algorithms becoming very commonplace, the need for primitives to "combine" keys is at an all time high. HMAC was just "secure enough"

Further Information

Backendal, M., Bellare, M., Günther, F., Scarlata, M. (2023). "When Messages Are Keys: Is HMAC a Dual-PRF?". In: Handschuh, H., Lysyanskaya, A. (eds) Advances in Cryptology – CRYPTO 2023. CRYPTO 2023. Lecture Notes in Computer Science, vol 14083. Springer, Cham. https://doi.org/10.1007/978-3-031-38548-3_22

Researchers

Matilda Backendal, Dr. Felix Günther, **Matteo Scarlata** - Applied Cryptography Group, ETHZ
Prof. Dr. Mihir Bellare - UCSD

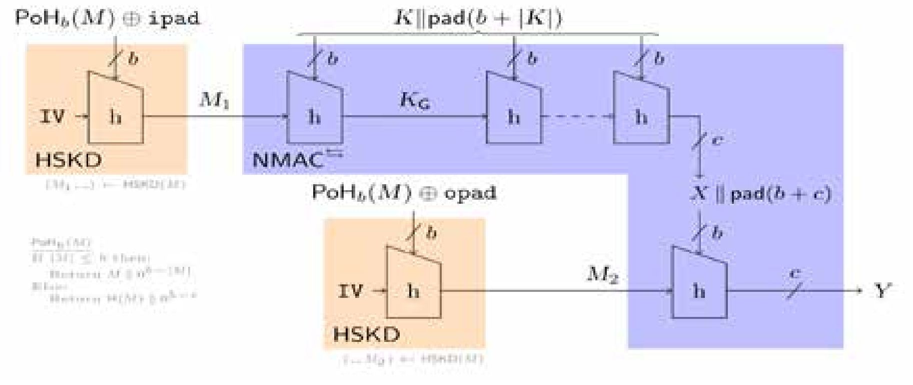
$\mathsf{PoH}_b(M) \oplus \mathtt{ipad}$
$K \| \mathsf{pad}(b + |K|)$
b
b
b
b
IV
h
M_1
h
K_G
h
h
c
HSKD
$\mathsf{NMAC}^{\leftrightarrows}$
$\mathsf{PoH}_b(M) \oplus \mathtt{opad}$
$X \| \mathsf{pad}(b + c)$
b
b
IV
h
M_2
h
c
Y
HSKD

Figure 2: Illustration of HMAC, when used in a "swapped" fashion

Research Highlights 2023

Collaborations with the International Committee of Red Cross (ICRC)

Prof. David Basin

Prof. Basin's group. In times of armed conflict, the emblems of the red cross, red crescent, and red crystal are used to mark physical infrastructure. This enables military units to identify assets as protected under international humanitarian law to avoid attacking them. In 2020 and in context of our work in the Centre for Cyber Trust, funded by the Werner Siemens-Stiftung, we challenged ourselves with the novel security problem of how to extend such protection to digital, network-connected infrastructure through a digital emblem. In 2021, we finalized a first version proposal for a digital emblem, that we call ADEM: An Authentic Digital Emblem. Since 2021, our design has been evaluated by domain experts, in a series of meetings at the invitation of the International Committee of the Red Cross, and our design was finalized and accepted for publication at the top-tier security conference ACM CCS 2023.

ADEM provides a unique combination of security requirements, namely, authentication, accountability, and a property that we call covert inspection. Covert inspection states that those wishing to authenticate assets as protected must be able to do so without revealing that they may attack unprotected entities. Moreover, ADEM was designed to fit the context of international diplomacy as it allows nation states to make sovereign decisions through a decentralized design.

In 2023, we also initiated the knowledge transfer phase of this research project, by developing ADEM prototypes in collaboration with the ICRC. Currently, we work on deploying ADEM in the ICRC's network to showcase that ADEM is both practical and scales to large organizations. We are optimistic that the ICRC's network can present the digital emblem to anyone interacting with it by the end of 2023.

Prof. Srdjan Capkun

Prof. Capkun's group. Relying on cloud infrastructures requires trust in the cloud service provider (CSP). Currently, this trust is necessary because, the CSP has physical access to the machines in which the data resides or is being processed, and control over supervisor software. While the CSP intentions might not be actively malicious, it might be forced to employ these attacks to comply with a lawful order to do everything necessary to access or tamper with customers' data. Given the attacker capabilities of the CSP and the possibility of these lawful requests, international organizations are usually faced with the choice of either having to fulfill their mission or employing CSP services. For instance, the International Committee of the Red Cross (ICRC) regularly visits war prisons to verify whether human rights are being violated.
The information collected as part of these visits could give an edge to the parties involved in the conflict. Therefore the ICRC is allowed to visit on the condition that information is kept secure and inaccessible to the other party. This guarantee cannot be reasonably given if the CSP is under the jurisdiction or sphere of influence of a country involved in the conflict. Thus, current CSPs cannot provide services for such organizations. In this project, we are exploring technical solutions that aim at bridging this gap. In particular, we are exploring solutions that would give a data owner, i.e., the ICRC the guarantee that the CSP can never access or tamper with their data while still benefitting from a cloud deployment.

Prof. Adrian Perrig

Prof. Perrig's group. The ICRC relies on digital infrastructure in order to fulfill its mission. As an International humanitarian organization, it operates in contexts of armed conflicts and violence. Thanks to its neutral role and diplomatic immunities, it has access to highly confidential data. Such information represents a high value target for state actors involved in conflicts, and therefore requires strong data protection measures. In addition, the migration of workloads to public clouds makes it more challenging to keep data under the same jurisdiction and protected by the organisation immunities. With this shift, Internet connectivity between the organisation branches, users and cloud datacenters becomes even more critical, especially when it comes to guaranteeing confidentiality, sovereignty, availability and protection from state surveillance.

The ICRC collaborates with ZISC and the Network Security Group in order to tackle such challenges while leveraging the SCION next generation Internet Architecture. Joint research efforts focus on several aspects of securing Internet communication. We showcased how SCION provides strong routing security, protecting traffic from route hijacks, that are common on today's BGP-based internet and are often exploited by threat actors to eavesdrop communications. Additional sovereignty guarantees are provided thanks to SCION's path awareness, so that Internet traffic can be "geofenced" and exclusively routed on trusted infrastructure.

ICRC

Education 2023

Security in School education

The Center of Computer Science Education (ABZ) of ETH Zurich was established with the goal to introduce computer science as a subject into school education. The main activities of ABZ include developing text-books and online platforms for teaching computer science on all levels of schools and testing them in school, training teachers, popularization of computer sci-ence in the whole society, and supporting pupils for different CS competitions like Olympiad in Informatics, Informatics Beaver, ACM Programming Contests.

The main achievements are establishing "informatics" as a mandatory subject in Lehrplan 21 for obligatory schools as a result of long-term projects in more than 500 schools involving more than 5000 teachers in training, 19 textbooks for teaching computer science in all age groups from kindergarten to high school, and more than 400 appearances in the media.

The main contributions of the last year are: 1) the textbook "Algorithmen und künstliche Intelligenz" with the focus on data management, design of algorithms and machine learning of winning strategies (360 pages with detailed explanations, motivating challenges and projects).
2) Organization of the competition Informatics Beaver with swiss finals at ETH in all 6 age categories.
3) Teacher training in cryptography for about 60 high school teachers.
4) about 150 schoolprojects (8-20 lessons per class) including security issues.
5) about 40 projects for gifted children indifferent Swiss cantons.

The ZISC center is proud to support this project!

Main Research Areas

Sovereign Smartphone

Prof. S. Shinde

The majority of smartphones either run iOS or Android operating systems. This has created two distinct ecosystems largely controlled by Apple and Google—they dictate which applications can run, how they run, and what kind of phone resources they can access. Barring some exceptions in Android where different phone manufacturers may have influence, users, developers, and governments are left with little control. Specifically, users need to entrust their security and privacy to OS vendors and accept the functionality constraints they impose. Given the wide use of Android and iOS, immediately leaving these ecosystems is not practical, except in niche application areas.

We are building a new smartphone architecture that securely transfers the control over the smartphone back to the users while maintaining compatibility with the existing smartphone ecosystems. Our architecture, named TEEtime, implements novel TEE-based resource and interrupt isolation mechanisms which allow the users to flexibly choose which resources (including peripherals) to dedicate to different isolated domains, namely, to legacy OSs and to user's proprietary software. We have shown the feasibility of TEEtime design via a prototype on ARM platform and are working towards building a fully functional phone.

Foundations of Cryptography

Prof. D. Hofheinz

Cryptographic building blocks (such as encryption schemes or zero-knowledge protocols) ensure the secrecy and integrity of information, and help to protect the privacy of users. Still, most actually deployed cryptographic schemes are not known to have any rigorously proven security guarantees.

Our goal is to provide practical cryptographic building blocks that come with rigorously proven security guarantees. These building blocks should be efficient enough for the use in large-scale modern information systems, and their security should be defined and formally analyzed in a mathematically rigorous manner. Specifically, we are interested in the foundations of theoretical cryptography, and in general ways to derive constructions and security guarantees in a modular fashion.

Future Internet Architecture SCION

Prof. A. Perrig

SCION is the first clean-slate Internet architecture designed to provide route control, failure isolation, and explicit trust information for end-to-end communications. SCION organizes existing ASes into groups of independent routing sub-planes, called isolation domains, which then interconnect to form complete routes. Isolation domains provide natural isolation of routing failures and human misconfiguration, give endpoints strong control for both inbound and outbound traffic, provide meaningful and enforceable trust, and enable scalable routing updates with high path freshness.

As a result, the SCION architecture provides strong resilience and security properties as an intrinsic consequence of good design principles, avoiding piecemeal add-on protocols as security patches.

Secure Positioning and Localization

Prof. S. Capkun

In our daily lives, we increasingly rely on location and proximity measurements for important applications. Already today, people issue contactless payments simply by bringing their credit card close to a card reader. Modern cars automatically detect the key in proximity to unlock the doors. We see first instances of autonomous transportation drones navigate using GPS or Galileo.

The security of many existing and future applications surrounding navigation, access control, and secure routing critically depends on the correctness of the underlying location and proximity estimates.

Trusted Execution Beyond CPUs

Prof. S. Shinde

Modern data centers have grown beyond CPU nodes to provide domain-specific accelerators such as GPUs and FPGAs to their customers. From a security standpoint, cloud customers want to protect their data. They are willing to pay additional costs for trusted execution environments such as enclaves provided by Intel SGX and AMD SEV. Unfortunately, the customers have to make a critical choice—either use domain-specific accelerators for speed or use CPU-based confidential computing solutions.

To bridge this gap, we are building datacenter scale confidential computing that expands across CPUs and accelerators. Having wide-scale TEE-support for accelerators presents a technically easier solution, but is far away from being a reality. Instead, we aim to provide enclaved execution guarantees for computation distributed over multiple CPU nodes and devices with/without TEE support, which presents security, scalability, and performance challenges.

Machine learning Security

Prof. F. Tramèr

Machine learning systems are becoming critical components in various industries, yet they face clear security and privacy challenges. Attacks on a machine learning models data can destroy the integrity of the entire system; deployed models can memorize and leak sensitive training data; and models themselves can be copied and stolen.

In our research, we study the behavior of machine learning systems in adversarial settings, to better understand the current limitations and risks of this nascent and booming technology. We then draw on this knowledge to propose new defense mechanisms to safeguard machine learning applications and their users.

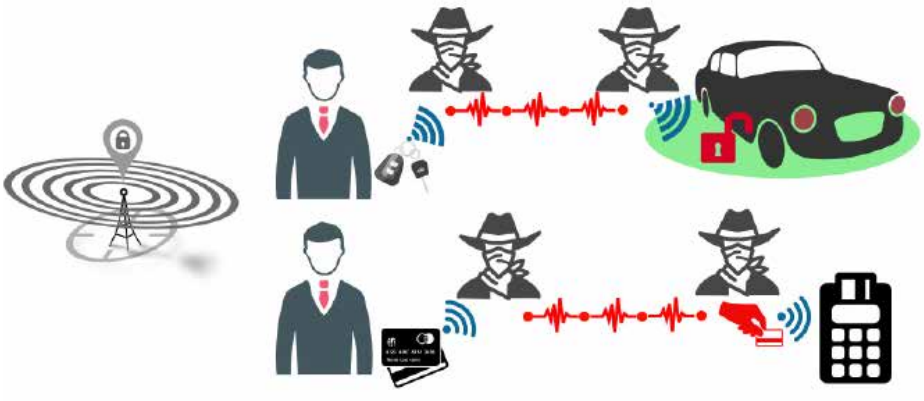

Main Research Areas

Access control

Prof. D. Basin

Access control has wide range of applications, from physical access control, e.g., to buildings, over to access control in single applications, to enterprise-wide access control solutions for an entire company's data management.

Our work covers multiple facets of the challenges in effective access control. We work on languages for efficiently analyzable yet expressive access control policies, techniques for ensuring that no security-critical access control queries are omitted, mining access control policies, and modern systems for access control in physical systems.

Constructive Cryptography

Prof. U. Maurer

Constructive cryptography is a new paradigm for defining the security of cryptographic schemes. It allows to take a new look at cryptography and the design of cryptographic protocols.

One can give explicit meaning to various types of game-based security notions of confidentiality, integrity, and malleability, one can design key agreement, secure communication, certification, and other protocols in a modular and composable manner, and one can separate the understanding of what cryptography achieves from the technical security definitions and proofs, which is useful for didactic purposes and protocol design.

Applied Cryptography

Prof. K. Paterson

Cryptography provides a fundamental set of techniques that underpin secure systems. It includes basic techniques to enable services such as confidentiality and integrity of data in secure communication systems, as well as much more advanced methods such as cryptographic schemes that enable searches over encrypted data.

It draws broadly from theoretical computer science (algorithms, complexity theory), mathematics (number theory, probability) and engineering (both electronic- and software-engineering). Our research in Applied Cryptography brings all of these strands together to produce impactful research that improves the security of today's and tomorrow's cryptographic systems.

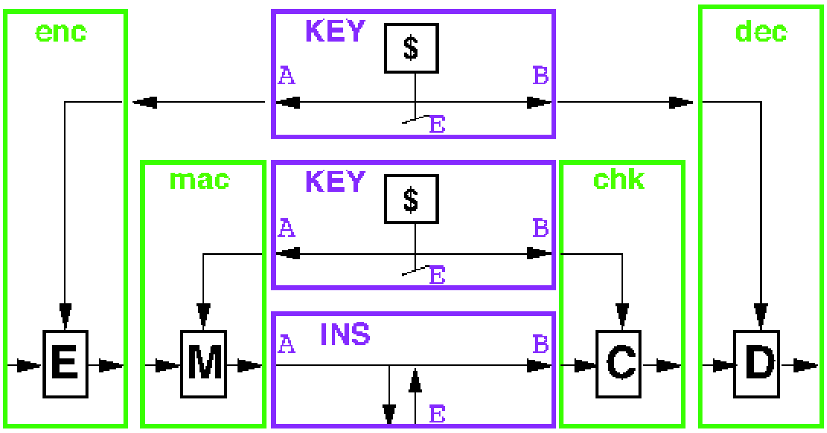

Security protocol verification

Prof. D. Basin

We develop tools for security protocol analysis in the symbolic model, particularly the Tamarin-prover. Security protocols are well-known to be error-prone, thus having a formal analysis enhances trust in the protocol's correctness. Our tools have theoretic foundations guaranteeing their conceptual correctness.

The symbolic model has a high level of abstraction compared to bitstrings over the wire, but provides automation and the ability to consider all modes of operation of a protocol at once. Practical results and impact has been seen in the standardization of TLS v1.3 and the analysis and discovery of serious vulnerabilities in the EMV protocol used for worldwide electronic payments.

Blockchain Technology

Prof. S. Capkun

Blockchain technologies promise many attractive advantages for digital currencies, financial applications and digital society in general. These advantages include reduced trust assumptions, increased transparency, reduced costs and improved user privacy. However, the current state-of-the-art solutions suffer from significant limitations.

In our research we investigate the limitations of current blockchain solutions and develop novel systems with improved security, privacy and performance guarantees. Examples of our research results include new types of smart contract execution environments, improved solutions for client privacy and novel digital currencies with regulatory support.

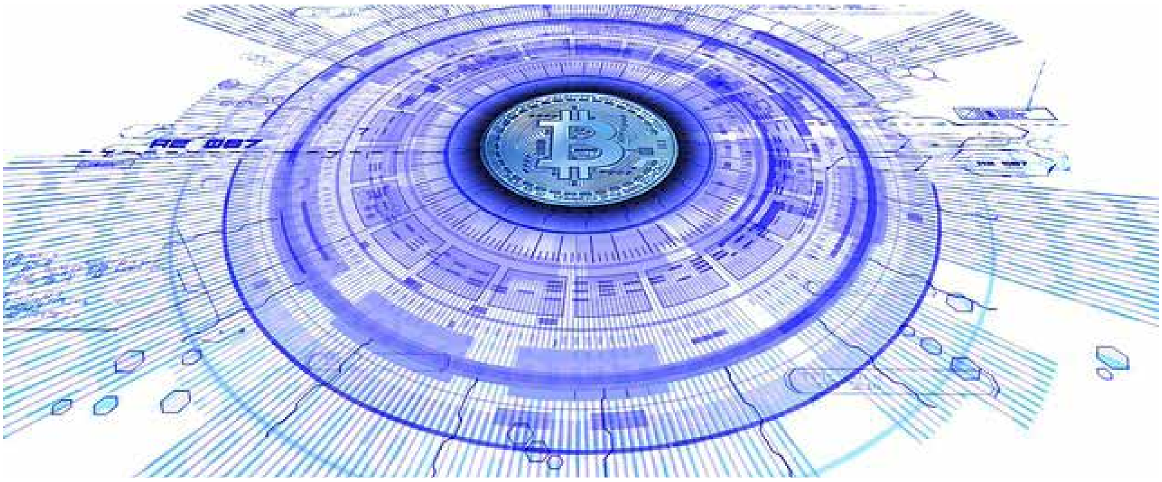

Research Projects

Highly Available Communication for Financial Networks

Communication, in particular for critical infrastructures, requires a high level of availability that remains available despite earthquakes, power outages, misconfigurations, or network attackers. One example is the financial industry, which has high requirements on availability to ensure that up-to-date trading information is accessible, that financial transactions are executed within short time windows, and that end customers can execute banking applications online.

The financial industry is generally a prime target for network attackers, mainly because of the importance of availability for banking applications. The importance of availability has lead to several extortion attacks in the past, where banks would sometimes pay for attack termination in the short term, rather than protecting their networks for the long term.

To provide high availability and thus to make the financial industry robust against the aforementioned attacks and calamities, we investigate the deployment of multi-path connectivity between branches of one of our ZISC banking partners. In the context of the future Internet architecture SCION, designed and developed at ZISC, we focus on connecting branches over multiple existing links at the same time.

We are deploying a multi-path communication system that automatically selects multiple independent, high-quality paths to avoid outages even if some of the independent paths fail.

To further increase the resilience against attacks, in particular in the context of DDoS defense and IoT security, our architecture offers the option to hide paths from the public and thus to prevent attackers from flooding such invisible paths. Moreover, we have developed a scalable bandwidth-reservation scheme that protects inter-domain communication by establishing fine-grained resource allocations to ensure no links between networks are saturated.

Further information

A. Perrig, P. Szalachowski, R. M. Reischuk, L. Chuat.
SCION: A Secure Internet Architecture
Springer International Publishing AG, 2017.

Benjamin Rothenberger, Dominik Roos, Markus Legner, and Adrian Perrig.
PISKES: Pragmatic Internet-Scale Key-Establishment System.
In Proceedings of the ACM Asia Conference on Computer and Communications Security (ASIACCS) 2020.

Cyrill Krähenbühl, Seyedali Tabaeiaghdaei, Christelle Gloor, Jonghoon Kwon, David Hausheer, Aadrian Perrig, and Dominic Roos.
Deployment and Scalability of an Inter-Domain Multi-Path Routing Infrastructure.
ACM Conference on emerging Networking EXperiments and Technologies (CoNEXT) 2021.

Researchers

Various members of the Network Security Group.

Industry partner

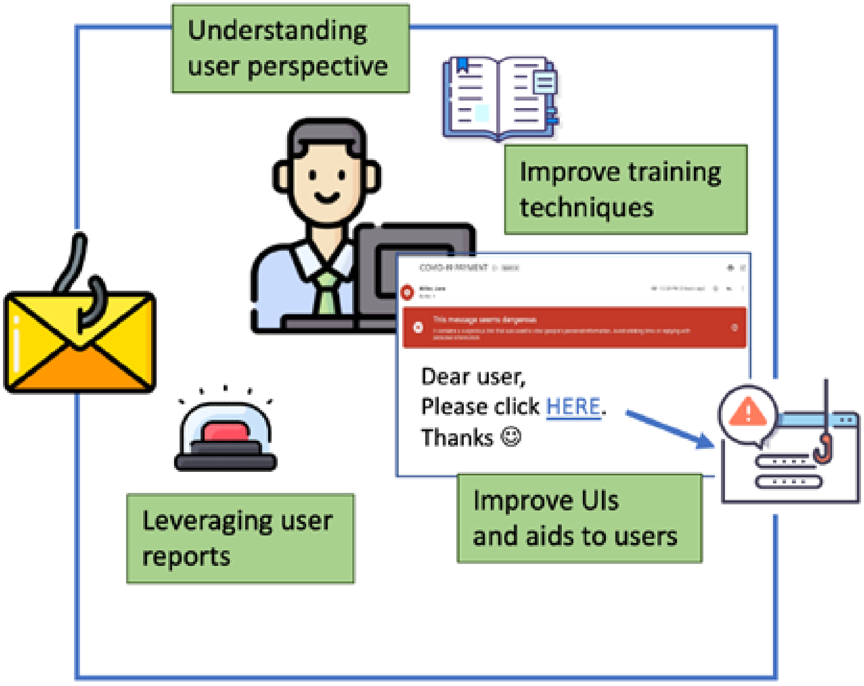
Understanding user perspective
Improve training techniques
This message seems dangerous
Leveraging user reports
Improve UIs and aids to users

Dear user,
Please click HERE.
Thanks ☺

Phishing in Large Organizations

Phishing emails are deceptive messages made for data stealing and malware propagation. In this type of attack, miscreants pose as legitimate organizations (e.g., banks and financial institutions, delivery companies, shopping websites), and send emails crafted to look like the impersonated organization's. These emails try to solicit a sense of urgency by prompting users to act swiftly, such as changing compromised passwords. Links in these emails lead to deceptive websites that often include login pages to try and trick the user into submitting their credentials. Other means for the attack can be opening malicious attachments or drive-by downloads of malware.

Phishing is a real threat to corporations: employees falling for phishing and revealing corporate credentials can be the first step for further attacks and data breaches. Phishing leads to significant economic losses: estimates put the yearly cost of phishing attacks in the order of millions of dollars for companies that fall victim. For this reason, it is of paramount importance to understand the most effective ways to protect users from phishing attacks.

In this project, in partnership with the Swiss Post, we aim to understand phishing in large organizations from the point of view of employees and IT departments. On employees, our measurements are improving how phishing training is delivered and understood, and we are developing novel user interfaces to help people spot potential attacks. On IT departments and defenders, we are analyzing novel countermeasures to deploy in organizations for early detection of phishing attacks.

Further information
D. Lain, K. Kostiainen, S. Capkun.
Phishing in Organizations: Findings from a Large-Scale and Long-Term Study.
In IEEE S&P 2022. San Francisco, CA, USA.

Researchers
Daniele Lain (ETH)
Kari Kostiainen (ETH)
Prof. Dr. Srdjan Capkun (ETH)

Industry partner

SWISS POST

Research Projects

Blockchain and Cloud Security

In this project, NEC and ETH aim to address various issues in cloud and blockchain security to improve their security and scalability. In blockchain technology, our project focuses on the security and privacy of different blockchain technologies and on developing new protocols and systems to enhance functionality.

As the first research contribution, we have proposed a new approach to protect the privacy of lightweight clients in blockchain systems like Bitcoin. Our main idea is to leverage commonly available trusted execution capabilities, such as SGX enclaves. We have designed and implemented a system called BITE where enclaves on full nodes serve privacy-preserving requests from lightweight clients. Because a naive method of serving client requests from within SGX enclaves still leaks user information, BITE integrates several privacy measures that address external leakage and SGX side channels. The resulting solution provides strong privacy protection and improves the performance of current lightweight clients.

As the second research contribution, we have designed and developed a new method to allow for the execution of expressive smart contracts on legacy cryptocurrencies, such as Bitcoin, that do not natively support a Turing complete scripting language. Our system, called Bitcontracts, allows the smart contract creator to designate a set of so-called service providers that are responsible for executing the contract off-chain. The contract state is stored in on-chain transactions, and the service providers can collectively authorize state changes by using multi-signature transactions signed by a quorum of them.

As the third research contribution, we have investigated the problems with mining centralization and analyzed approaches that try to solve these issues with decentralization of mining pools. We have found that mining centralization provides several advantages for individual miners compared to decentralized solutions and thus miners are incentivized to prefer centralized mining pools. To mitigate some of the issues that arise from current centralized mining pools, we have proposed a novel mining solution using trusted execution environments.

As the fourth research contribution, we have investigated the censorship-resilience of fast blockchain payments. Permissionless blockchains are known to be too slow for applications like point-of-sale payments. While several techniques have been proposed to speed up blockchain payments, none of them are satisfactory. In particular, existing solutions like payment channels require users to lock up significant funds, and schemes based on pre-defined validators enable easy transaction censoring. We have developed a system called Quicksilver that works with practical collaterals and is fast, censorship-resilient, and confidential at the same time.

Researchers

Kari Kostiainen (ETH)
Ghassan Karame (NEC)

Further information

Censorship-Resilient and Confidential Collateralized Second-Layer Payments
Kari Kostiainen, Sven Gnap, Ghassan Karame, eprint, November 2022
https://eprint.iacr.org/2022/1520.pdf

Industry partner

NEC

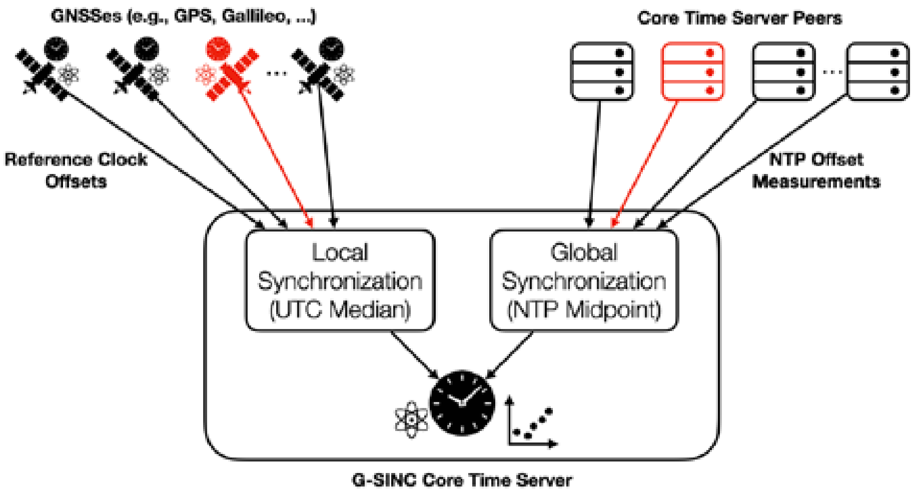
GNSSes (e.g., GPS, Gallileo, ...)
Core Time Server Peers
Reference Clock Offsets
NTP Offset Measurements
Local Synchronization (UTC Median)
Global Synchronization (NTP Midpoint)
G-SINC Core Time Server

G-SINC: Global Synchronization Infrastructure for Network Clocks

Secure and dependable time synchronization is an essential prerequisite for many industries with applications in finance, telecommunication, electric power production and distribution, or environmental monitoring.

Current best practice to achieve large-scale time synchronization relies on global navigation satellite systems (GNSSes) at the considerable risk of being exposed to outages, malfunction, or attacks against availability and accuracy. Natural disasters like solar superstorms also have the potential to hit and severely impact GNSSes.

It is therefore all too apparent that time synchronization solely based on GNSSes as global reference clocks does not fulfill fundamental dependability requirements for systems that serve indispensable functionalities in our society. Facing these concerns, governments have issued mandates to protect critical infrastructure services from disruption to GNSS services, including a 2020 US Executive Order. Operators and equipment manufacturers are encouraged to intensify research and development of alternative technologies in this space.

Aiming to join these efforts, we are developing G-SINC: a novel global, Byzantine fault-tolerant clock synchronization approach that does not place trust in any single entity and is able to tolerate a fraction of faulty entities while still maintaining accurate synchronization on a global scale among otherwise sovereign network topologies. G-SINC can be implemented as a fully backward compatible active standby solution for existing time synchronization deployments.

This is achieved by building on the solid body of fault-tolerant clock synchronization research dating all the way back to the 1980s and the SCION Internet architecture providing required resilience and security properties at the network level as an intrinsic consequence of its underlying design principles.

Besides the possibility to use multiple distinct network paths in parallel for significantly improved fault-tolerance, we highlight the fact that SCION paths are reversible and therefore symmetric. Hence, they help to increase time synchronization precision compared to clock offset measurements over the often asymmetric paths in today's Internet.

Researchers
Marc Frei (ETH)
Dr. Jonghoon Kwon (ETH)
Seyedali Tabaeiaghdaei (ETH)
Marc Wyss (ETH)
Prof. Dr. Adrian Perrig (ETH)
Dr. Christoph Lenzen (CISPA)

Research Projects

Improving Network Security Through Programmability

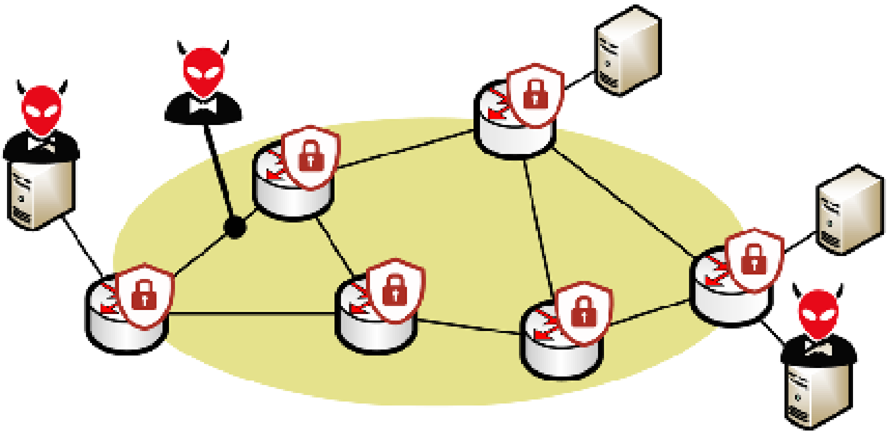

In this project, we argue that the network itself should be able to detect and mitigate attacks instead of relying purely on perimeter-based protection provided by dedicated appliances. To do so, we plan to leverage recent advances in network programmability which enable both the control plane and the data plane to be reprogrammed on-the-fly.

The goal of this project is to leverage recent advances in network programmability to make the network able to defend itself against: (i) anonymity and privacy attacks, performed by attackers which can eavesdrop on and modify traffic; and (ii) more general attacks (e.g., denial-of-service, data exfiltration), performed by attackers sitting at the edge of the network, on compromised hosts.

Protecting networks from in-network attackers: This part of the project aims at designing and developing a network-based anonymity and privacy framework targeted specifically at enterprise networks. Being network-based, the framework will enable to secure any connected devices (even unforeseen ones) and internal communications, without complex setup. To develop this "securing" network, we will actively leverage the new programmability primitives offered by Software-Defined Networks (SDN) in both the control plane (OpenFlow) and the data plane (P4).

Protecting networks from edge attackers: In this part of the project, we focus on attackers that get access to the network via one or more infected hosts. After infecting at least one host, such attackers usually initiate a "reconnaissance" phase in which they scan the network in search of high value targets. Network programmability enables to efficiently distribute the task of scan detection on the network devices and provides the ability to source traffic on the network device in order to implement advanced deception techniques in which the attacker is presented with fake information (e.g., fake IP addresses).

Further information

Ege Cem Kirci, Maria Apostolaki, Roland Meier, Ankit Singla, Laurent Vanbever. "Mass Surveillance of VoIP Calls in the Data Plane". ACM SOSR 2022. (Online October 2022).

Roland Meier, Vincent Lenders, Laurent Vanbever. "ditto: WAN Traffic Obfuscation at Line Rate". NDSS Symposium 2022. San Diego, CA, USA (April 2022).

For more details, see: https://nsg.ee.ethz.ch

Researchers

Roland Meier (ETH)
Laurent Vanbever (ETH)

Industry partner

Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra

armasuisse

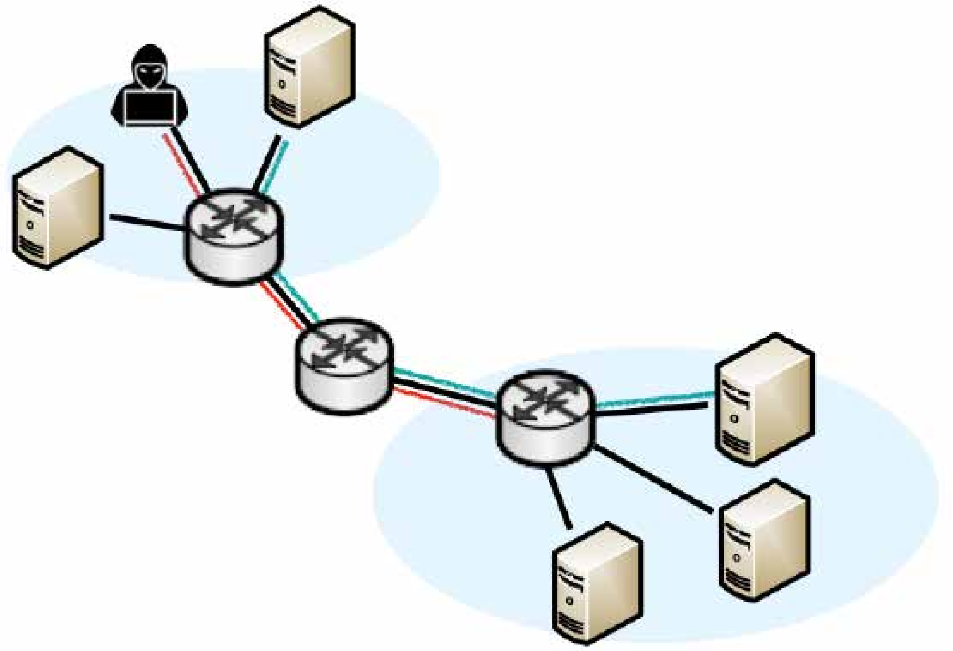

Self-securing Networks

The goal of this project is to build data-driven network infrastructures that can autonomously protect, detect and defend themselves against attacks. We intend to develop network-specific learning and inference algorithms that can run directly in the data plane, in real-time, to perform tasks that are difficult to solve today such as (encrypted) traffic classification and fine-grained anomaly detection. To implement these learning and inference algorithms, we intend to leverage the newly available capabilities of programmable data planes to run complex forwarding logics. Specifically, we will use these capabilities to: (i) extract representative network data; (ii) train learning models; and (iii) drive forwarding decisions accordingly— at line rate.

Traffic classification: In a first package, we intend to build in-network online classification mechanisms. Traffic classification is a key building block when securing today's networks. Classifying traffic directly in the network enables network devices to adapt their forwarding decisions according to the application types. For instance, it enables network switches to direct specific flows to dedicated boxes for further processing. It also enables switches to drop traffic (or possibly de-prioritize it) as soon as it enters the network.

Anomaly detection: In a second package, we intend to investigate methods and tools on top of programmable data planes to perform anomaly detection network-wide, ideally on all the traffic. While performing large-scale anomaly detection is highly challenging and requires fundamental research contributions, one can use simpler, detection mechanisms in the data plane, and compensate for their lack of precision (i.e.. false positives) with lightweight confirmation stages.

Data-driven defenses: In a third package, we intend to consider the problem of active, data-driven network defenses.

Intuitively, while the two first packages consider the problem of sensing the network, this work package will consider the problem of actuating the network accordingly, i.e. closing the control loop. Here we plan on developing several techniques to confirm and mitigate alleged attacks.

Further information

Albert Gran Alcoz, Martin Strohmeier, Vincent Lenders, Laurent Vanbever. "Aggregate-Based Congestion Control for Pulse-Wave DDoS Defense". ACM SIGCOMM 2022. Amsterdam, Netherlands (August 2022). For more details, see: https://nsg.ee.ethz.ch

Researchers

Albert Gran Alcoz (ETH)
Laurent Vanbever (ETH)

Industry partner

Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra

armasuisse

Research Projects

VerifiedSCION

Full-Stack Verification of Secure Inter-Domain Routing Protocols

Inter-domain routing is a part of the Internet's core infrastructure. The currently used Border Gateway Protocol suffers from attacks leading to severe disruptions of the Internet. This prompted the development of the secure Internet architecture SCION. In this research project, we examine the SCION protocols in detail and formally verify their desired functional and security properties. We do this both at the modeling and the implementation level. Our goal is to gain a better understanding of the underlying properties of the SCION protocols and routing protocols in general, and to improve on the state of the art for the verification of concurrent, object-oriented programs. Moreover, this work will contribute to the first Internet protocol suite that has been verified from the ground up.

We have developed a combined model-and-code verification technique, which works as follows. We start by formalizing the protocols and their security properties. We then use several refinement steps to derive more concrete protocol models from which we automatically extract program specifications expressing the implementation's desired I/O behavior (called I/O specifications). All these steps are formalized in the interactive theorem prover Isabelle/HOL.

In particular, we have proved the soundness of the translation from protocol models to I/O specifications, which links two quite different formalisms. We then use a code verifier to prove the functional correctness of the implementation (i.e., memory and crash safety, race freedom, and adherence to the I/O specification). We additionally prove security-related properties of the implementation like secure information flow.

To facilitate the verification of dataplane protocols, we have created a parametric verification framework in Isabelle/HOL. Based on this framework, we have verified the latest version of the SCION data plane protocol and automatically generated an I/O specification for the router implementation. This work is reported in a 2023 journal paper, extending our 2021 conference paper.

The SCION router is written in the Go programming language. We have developed "Gobra", a Go verifier capable of handling Go's advanced language features, namely channel-based concurrency and interfaces. Using Gobra, we have proved memory and crash safety for SCION's border router implementation. We are now in the process of proving that the router code also satisfies its I/O specification.

Recently, we have also published a paper about a technique and a tool to automatically generate I/O specifications from Tamarin protocol models. This enables the future use of the Tamarin prover to automatically verify SCION protocols and link them to an implementation.

Further information

Linard Arquint, Felix A. Wolf, Joseph Lallemand, Ralf Sasse, Christoph Sprenger, Sven N. Wiesner, David Basin, and Peter Müller. Sound Verification of Security Protocols: From Design to Interoperable Implementations. SP 2023.

Tobias Klenze, Christoph Sprenger, and David Basin. IsaNet: A framework for verifying secure data plane protocols. Journal of Computer Security, 2023.

Felix A. Wolf, Linard Arquint, Martin Clochard, Wytse Oortwijn, João C. Pereira, and Peter Müller. Gobra: Modular Specification and Verification of Go Programs. CAV 2021.

Researchers

Prof. David Basin (ETH)
Prof. Peter Müller (ETH)
Prof. Adrian Perrig (ETH)
Dr. Christoph Sprenger (ETH)
Dr. Ralf Sasse (ETH)
Sofia Giampietro (ETH)
Linard Arquint (ETH)
Felix Wolf (ETH)
João Carlos Mendes Pereira (ETH)
Dionysios Spiliopoulos (ETH)

SBAS: Bridging the Gap to SCION

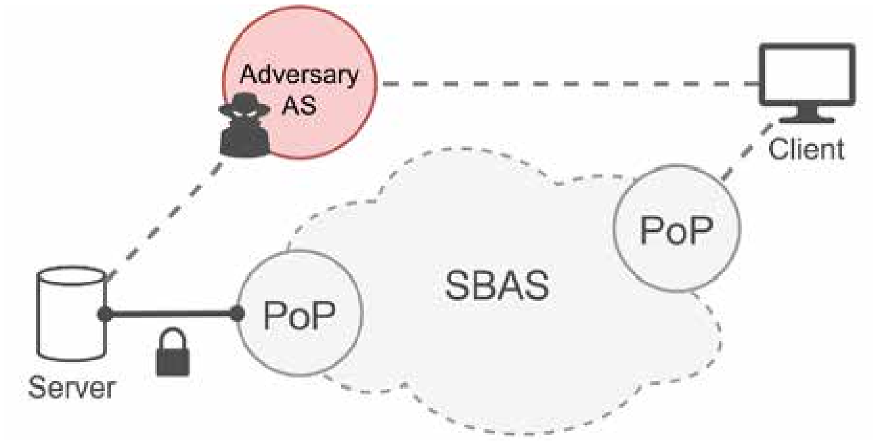

The recent Facebook outage went on record as one of the largest outages for a major application provider. With the root cause for Facebook, Instagram, and WhatsApp going offline being the BGP routing protocol, there is more awareness than ever that more reliable approaches are required to route Internet traffic.

Today, many products are offered that enable connectivity over a globally deployed private backbone such as Cloudflare. However, with such networks, customers seeking higher reliability and security for their internet connectivity are placing their trust in a single entity.

The inter-domain routing security provided by SCION enables a different approach: to construct a federated backbone consisting of a group of entities. In our project, we are developing the Secure Backbone AS (SBAS), a system that both leverages and drives partial deployment of SCION. It can be used to provide immediate benefits for legacy Internet hosts today. Crucially, SBAS requires minimal additions for Internet Service Providers (ISPs) that already deploy SCION and is compatible with standard BGP practices.

The SCION architecture is already serving a variety of use cases today. However, without SBAS, it is not possible to carry the benefits of SCION out into the wider Internet: a service hosted on a SCION endpoint will not offer improved security to customers of ISPs that do not deploy SCION. Using SBAS, the space for use cases is much larger: even endpoints that are not aware of the system can benefit from it, thanks to the seamless bridge between SCION and BGP provided by SBAS. At a small additional cost, ISPs can therefore deploy SBAS to tap into novel offerings for their customers, such as hijack-resilient server addresses or carbon-optimized Internet connections.

The goal of the SBAS project is to design and implement the system in a way that incurs minimal costs to the participating ISPs, in order to provide the financial incentives required for real-world deployment.

Moreover, after initial prototype implementations and experiments in academic network testbeds, the SBAS team is currently driving several efforts to set up a deployment with ISPs and customers.

Researchers

Joel Wanner (ETH)
Dr. Jonghoon Kwon (ETH)
Prof. Dr. Adrian Perrig (ETH)

Henry Birge-Lee (Princeton)
Grace Cimaszewski (Princeton)
Dr. Liang Wang (Princeton)
Prof. Dr. Prateek Mittal (Princeton)
Prof. Dr. Yixin Sun (Virginia)

Research Projects

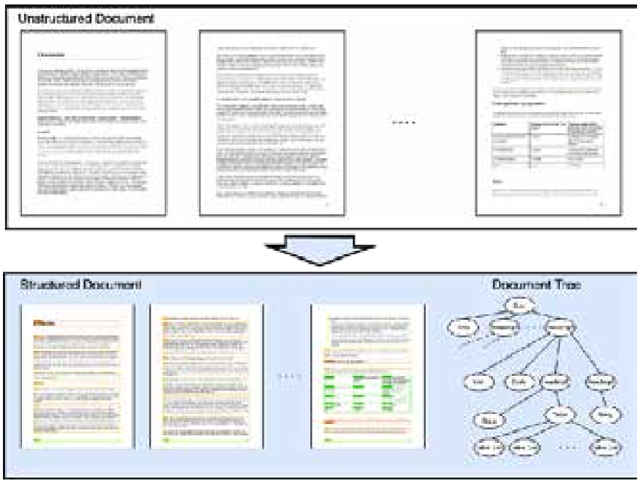

Enhancing Document Processing with Hierarchical Structure

Automated information retrieval techniques are powerful tools to build knowledge bases from data available in PDF documents, both in private organizations, the public sector, and the sciences. However, while these pools of data contain valuable information, they are typically unstructured, which poses a major obstacle to extracting useful insights using state-of-the-art information retrieval methods. This leads to the need for humans to manually go through hundreds of documents, a process that is not scalable and thus results in large amounts of data being left unexploited.

The goal of this project is to build an AI system that brings structure into these documents and thus enables further downstream processing by information retrieval engines. The system takes as input PDF documents and produces structured, intermediate representations of the documents. To achieve this goal, multiple challenges have to be overcome.

First, due to a lack of publicly available large-scale datasets, we build a system that can annotate the hierarchical structure of MS Word, LaTex, RTF, and other document formats at scale. To that end, we make use of structural information extracted from the source code of the documents. By crawling the web for these file types, we use our annotation system to create the first large-scale open dataset with a diverse range of annotated documents, reflecting the distribution of real-world documents composed by humans.

The second challenge is to design and train a large document analysis model which has a general "understanding" of document layouts, their content, and relations between different elements of the documents. As documents are inherently multi-modal, the model design needs to account for this and make use of recent progress in natural language processing, computer vision, and document analysis research.

The third challenge is to design a pipeline that allows researchers and practitioners to fine-tune our models on specific types of documents. As this process often requires organizations to provide infrastructure providers with their data and to respect privacy concerns, we need to develop a technique that enables anonymization, while still maintaining the layout and semantic meaning of the elements present in the documents.

Researchers

Prof. Ce Zhang (ETH)
Gero Gunkel (Zurich Insurance)
Maurice Weber (ETH)

Industry partner

ZURICH

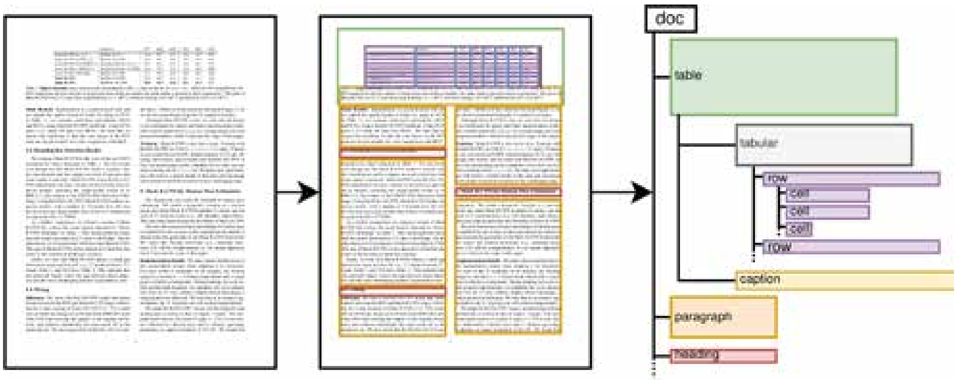

Automatic Visual Document Parsing

Automatic information retrieval methods are powerful tools to build structured knowledge bases from large datasets of real-world documents in science, industry and the public sector. The system we are building automatically produces an intermediate representation for a diverse range of documents that can be used by such information retrieval methods. It takes as input PDF documents or document images and translates them into JSON files containing the natural semantic hierarchy representing a document. These JSON files can be queried using a document database, and be used as a uniform document representation by downstream information extraction engines.

A major obstacle in using information retrieval methods on documents in PDF format is the lack of machine-readable structure information, e.g. document sections, tabular contents, lists, etc. Due to this challenge, ad-hoc code typically has to be written to correctly extract document contents for differently formatted documents. This approach often fails to generalize over varying document formats and code has to be re-written to cope with even minor format changes.

Instead of manually extracting contents from PDF raw data, we leverage the visual document representation for more robust content retrieval, similar to how a human reader would process the information. A convolutional neural network that operates on the rendered PDF documents is applied in our system. The network is trained for the task of page entity detection, e.g. the prediction of the locations of figures, tables and contained table cells and captions.

We pretrain the neural network in a weakly-supervised fashion on a large dataset of annotated documents that was automatically created from publicly available scientific articles. This weak supervision strategy greatly reduces need for manual annotation and allows for efficient adaptation of our system to new document types. In a subsequent step, structural relationships between detected page entities are automatically identified in order to produce the full hierarchical structure for document pages.

Researchers

Ce Zhang (ETH),
Johannes Rausch (ETH)
Gero Gunkel (Zurich)

Industry partner

ZURICH

Research Projects

Prioritizing Cybersecurity Controls based on Coverage of Attack Techniques

The objective of this project is to study existing mappings of controls in NIST Special Publication 800-53 to attacks in MITRE ATT&CK and to suggest an automatic or semi-automatic methodology to integrate attack frequencies and impacts to produce a priority ordering on controls to implement. This methodology will be implemented as a software prototype that allows domain experts to query it using various criteria.

As a first step, we design a Cyber Threat Intelligence feed model that, based on a a recent vulnerability statistics report or a vulnerability database, automatically computes attack frequencies. Our model proposes how to extend existing data structures used to exchange threat intelligence reports. We also review and evaluate existing prototypes that attempt to extract used techniques from textual descriptions of campaigns based on natural language processing.

We then devise an automated methodology to integrate the attack frequencies from the threat intelligence feed into the mapping in order to prioritize most impactful controls. Several algorithms are explored to allow an automatic computation of sets of controls to implement to optimize a risk minimization criteria (for instance, which controls mitigate the most frequently occurring attack techniques).

Last, we preliminary evaluate the methodology by instantiating it using information of malware campaigns from the past years and various example queries. We further discuss how our methodology can be evaluated in future work by means of an analysis of data breaches occurrences vs. implemented controls.

Researchers

Silvia La (ETH)
Dr. Martin Ochoa (ETH)
Vivien Bilquez (Zurich)

Industry partner

ZURICH

Interactively exploring 3D scanned dynamic environments

Swiss Post is active in areas that touch many parts of our daily lives, be it communication through mail, transportation, banking, and not least as a large employer in Switzerland. The goal of this project is to showcase the diversity of Swiss Post as a workplace through immersive, realistic and representative 3D experiences that people may discover and explore using emerging technologies, including Virtual Reality headsets and interactive 3D experiences on tablets and mobile devices. These experiences will give people unfamiliar with many of the activities of Swiss Post novel opportunities for insight into the daily lives of Swiss Post employees and customers across a variety of divisions. The immersive 3D experiences we are creating in this project are based on actual 3D scans of Swiss Post environments, fully interactive and ready to be explored to understand the World of Swiss Post. A second goal of this project is to use the rich captures of daily procedures performed by Swiss Post employees for training purposes of new personnel, thereby moving away from text-based instructions to immersive 3D scenarios that will aid learning on the job.

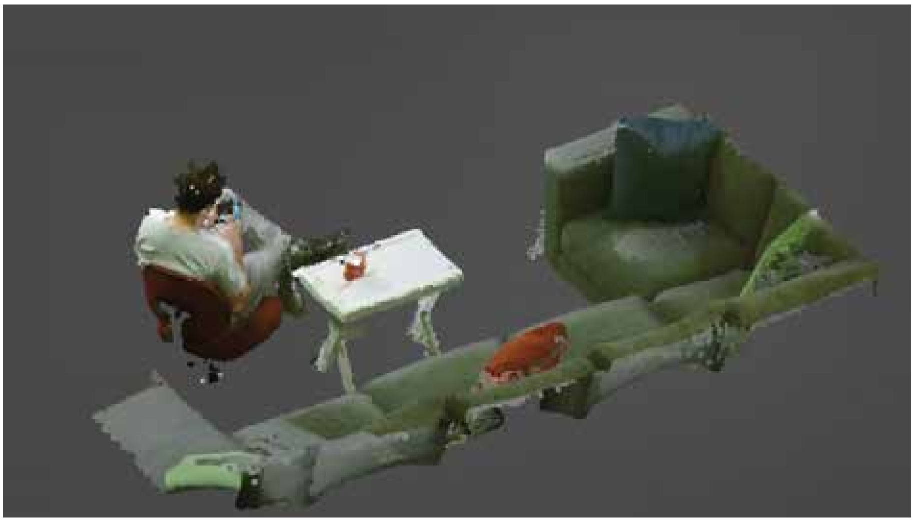

Approach

Solving the problems mentioned above and creating immersive 3D experiences based on scanned dynamic environments at Swiss Post requires processing technologies that fuse depth maps and textures from multiple high-resolution RGB and depth cameras into a coherent model. Post-processing needs to fuse the resulting point clouds into high-quality 3D meshes, removing artifacts and temporal inconsistencies, so as to render meshes in 3D for interactive consumption. To this end, we will build on our frameworks for fusing multi-camera input in conjunction with emerging point-cloud processing techniques and deep learning-based methods for scene understanding. Building on this will be a layer of interactivity, where elements of the 3D scene come to life and respond to user input. Using our experience in creating immersive 3D experiences, we will build and evaluate suitable interaction techniques for end users to interact with these 3D experiences, either in Virtual Reality or through touch controls on mobile devices.

Researchers

Prof. Christian Holz (ETH CS)
Dr. Andreas Fender (ETH CS)
Sensing, Interaction & Perception Lab, ETH Zürich

Industry partner

SWISS POST

Research Projects

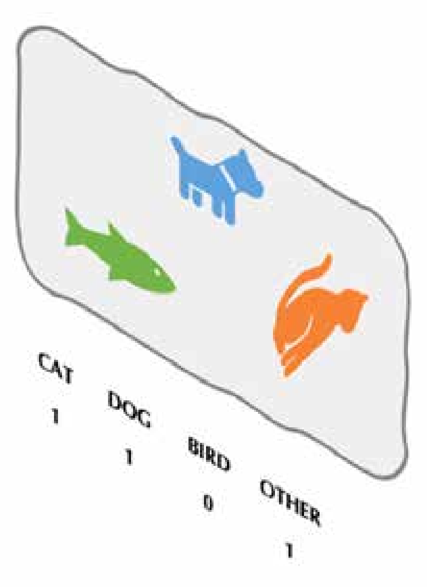

Multi-label classification

This research project aims at shedding a new light on multi-label classification and consists of two main goals: (i) developing a comprehensive, up-to-date benchmark on multi-label classification for two data modalities, and (ii) improving a multi-label classification system for an email forwarding task used by our industry partner.
Multi-label classification is a generalization of the common machine learning problem of multi-class classification. The distinction is that each data sample in a multi-label dataset can belong to multiple classes simultaneously, as opposed to only one in a multi-class regime. The number of real-world applications of multi-label classification has increased drastically in recent years --- from automatic categorization of lengthy emails and news articles, to tagging objects in images. In recent years we witnessed an opening of several new horizons of tools used in machine learning, from pre-processing, training, optimizing to post-processing. Knowing which method to deploy on a given task becomes harder with more tools being introduced.

At the same time, we observe very little research on multi-label classification tasks, even though they introduce fundamentally different questions than those in the multi-class regime, e.g., interdependent labels within complex label structures and often highly unbalanced datasets.

In our project, we start by performing a comprehensive study of existing methods. On a carefully curated list of datasets from two data modalities that are ubiquitous in modern machine learning (computer vision and natural language processing), we explore a cross product of numerous baselines, state-of-the-art machine learning methods and feature extraction strategies, and commonly-used classifiers, all evaluated through various metrics. This study is challenging due to the plenitude of available paths that one can take in designing a multi-label classifier. It yields new insights and aims at providing guidance for future applications.

In the second part of the project, we perform a similar experimental study, this time on a real-world dataset provided by our partner, building on top of their system that is currently in use. This email-forwarding task is quite interesting due to the sheer particularities of the task and the dataset --- a small number of samples, highly imbalanced label distribution, all combined with several constraints on the labels. We explore which properties of the above study can be transferred to this application and develop components that can be integrated in the system of our partner.

Researchers
Ce Zhang (ETH)
Luka Rimanic (ETH)
Gero Gunkel (Zurich)

Industry partner

ZURICH

Enhanced 5G Security

Security and privacy in 5G are highly challenging. As 5G connects everyone to everything everywhere, the 5G network is a rich source of critical information, from personal data and business assets, to mission-critical sensor data. To protect highly valuable information, 3GPP specifies the security aspects of the 5G system. The most significant 5G security enhancements compared to the previous generations are access-agnostic primary authentication, secure key establishment and management, and service-based architecture security.

Network slicing is the foundation of 5G security enhancements. 5G network slicing splits shared network resources into logical or virtual networks to satisfy specific service requirements that adhere to a Service Level Agreement (SLA). Each slice has isolation from the other network slices, achieving higher security with precise access control. To this end, different mechanisms may be envisioned for the logical network isolation, e.g., VLAN, Openflow, or other NFV mechanisms. Yet, no network slicing mechanism has been proposed, which suits for 5G environment.

The goal of this project is to leverage network programmability and cryptographic features that the next-generation Internet architecture delivers to enable:

i) dynamic network isolation at UE (User Equipment)-granularity, ii) network isolation continuity across remote edge networks even through the public Internet, iii) highly secure access control in network slice transit with cryptographic protection, and iv) scalable key establishment and management mechanisms.

Further information

Jonghoon Kwon, Taeho Lee, Claude Hähni, and Adrian Perrig.
SVLAN: Secure & Scalable Network Virtualization.
In Proceedings of the Symposium on Network and Distributed System Security (NDSS) 2020.
Jonghoon Kwon, Claude Hähni, Patrick Bamert, and Adrian Perrig.
MONDRIAN: Comprehensive Inter-domain Network Zoning Architecture.
In Proceedings of the Symposium on Network and Distributed System Security (NDSS) 2021.

Researchers

Jonghoon Kwon (ETH)
Prof. Dr. Adrian Perrig (ETH)

Industry partner

NEC

Research Projects

Design of Bug Bounty Schemes

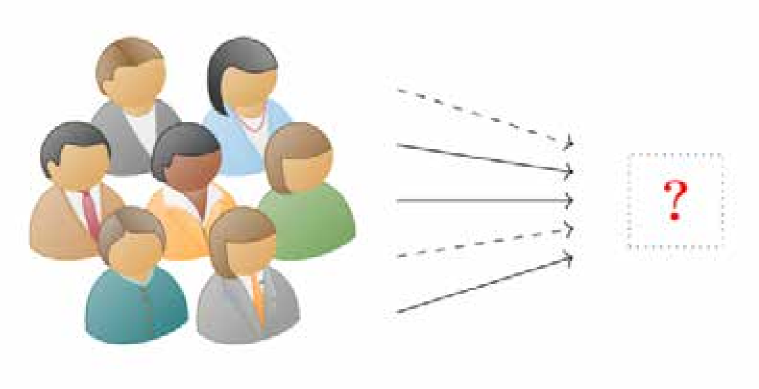

Systems and blockchains often have security vulnerabilities and can be attacked by adversaries, with potentially significant negative consequences. Therefore, infrastructure providers increasingly rely on bug bounty programs, where external individuals probe the system and report any vulnerabilities (bugs) in exchange for rewards (bounty).

We develop a simple contest model of bug bounty in which a group of individuals of arbitrary size is invited to undertake a costly search for bugs. The individuals differ with regard to their abilities, which we capture by different costs to achieve a certain probability to find bugs if any exist. Costs are private information. We study equilibria of the contest and characterize the optimal design of bug bounty schemes. In particular, the designer can vary the size of the group of individuals invited to search, add a paid expert, insert an artificial bug with some probability, and pay multiple prizes. We obtain the following results. First, we characterize the equilibria, establishing that any equilibrium strategy must be a threshold strategy, i.e. only agents with a cost of search below some (potentially individual) threshold participate in the bug bounty scheme. Second, we provide sufficient conditions for the equilibrium to be unique and symmetric. Third, we show that even inviting an unlimited crowd does not guarantee that the bug, if it exists, is found, unless there are agents which have zero costs, or equivalently have intrinsic gains from participating in the scheme. It may even happen that having more agents in the pool of potential participants lowers the probability of finding the bug. Fourth, adding a paid expert can increase or decrease the efficiency of the bug bounty scheme. Fifth, we demonstrate that in a model with multiple prizes, having one prize (winner-takes-all) achieves the highest probability of finding the bug. Sixth, we identify circumstances when asymmetric equilibria occur. Lastly, we illustrate how our baseline model can be extended to allow for multiple bugs, multiple experts, and heterogeneity of agents with respect to cost distributions, search times, and skills.

In the next step of the research, we examine how adding (known) bugs is another way to increase the likelihood that unknown bugs are found. When the additional costs of paying rewards are taken into account, it might be optimal to insert several known bugs, but maybe some only with a certain probability.

Further information

Hans Gersbach, Akaki Mamageishvili and Fikri Pitsuwan
Decentralized Attack Search and the Design of Bug Bounty Schemes
In Proceedings of the 16th International Symposium on Algorithmic Game Theory (SAGT) 2023.

Researchers

Prof. Dr. Hans Gersbach (ETH)
Dr. Fikri Pitsuwan (ETH)

Industry partner

SWISS POST

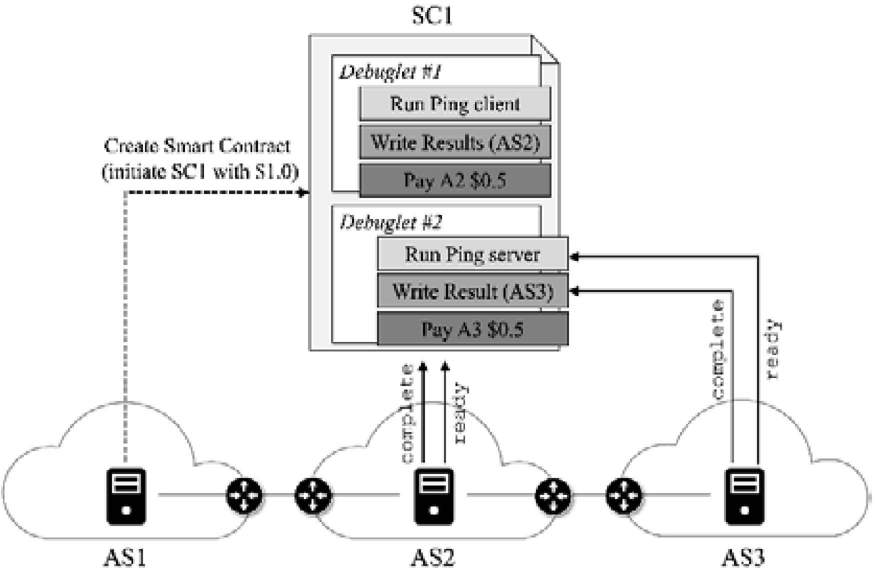

Debuglets: Programmable Network Debugging Infrastructure

Debuglets is an advanced distributed network debugging infrastructure designed to enhance end-user debugging on the Internet. Today's end-user debugging is limited to primitive tools like ping and traceroute, leaving users with insufficient data for isolating network faults and no means of external result validation. Debuglets offers a deployable, incentivized architecture that enables near-path network debugging using real data packets and user-defined code for precise and adaptable network performance measurements.

In this system, Autonomous Systems (ASes) deploy small-scale cloud services for network measurement applications. Users can pay to run their network debugging apps in these environments, and the results can be certified by the deploying AS for third-party verification. Debuglets features three essential planes: Management Plane, Measurement Plane, and Knowledge Plane, ensuring effective contract management, real-time network data collection, and efficient data dissemination.

Debuglets comprises three critical components: the Remote Code Execution (RCE) service, the execution environment for network measurement applications; the Coordinator, facilitating contracts and distributing Debuglets; and the Sanitizer, enforcing information disclosure policies during output dissemination. This infrastructure significantly improves end-user debugging, accelerates network issue identification, and introduces innovative business models for ASes, with even a partial deployment proving beneficial for users.

Researchers

Seyedali Tabaeiaghdaei (ETH)
Dr. Jonghoon Kwon (ETH)
Prof. Dr. Adrian Perrig (ETH)
Patrick Bamert (ZKB)

Industry partner

Zürcher Kantonalbank

Research Projects

Anonymization of Voice Recordings for Privacy-Preserving Emotional Analysis

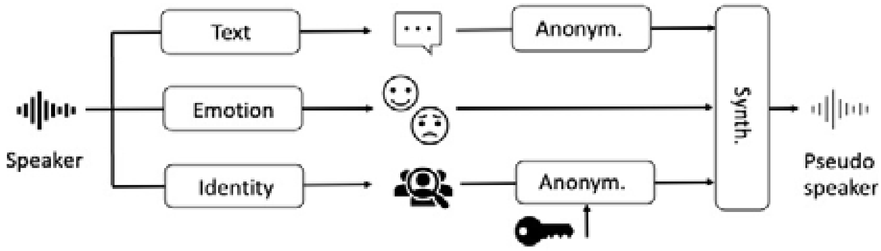

Motivation

Voice samples reveal information about identity, emotion, and other features of the speaker. Voice anonymization is the process to transform the voice sample of a speaker into the voice sample of a different artificial pseudo-speaker. This has many useful applications, each putting different requirements on the anonymization process. For example, when asked to use voice-based authentication, users might want to use a different and revocable pseudo-voice for each service, to remain protected even if their pseudo-voice is leaked. Clearly, this case has strict requirements on the mapping between real and pseudo-voices, but not on the emotional content of the pseudo-voices. In this project, we focus on the different use case of emotional analysis on audio samples from conversations with the customer service. In this case, the anonymization system is used by the company to protect the identity of the customers before storing, analyzing and sharing the audio samples. In this case, the anonymization system must preserve the emotional content, while providing strong guarantees against deanonymization, but many of the features required for authentication are irrelevant.

Related Work

The problem of voice anonymization has been broadly studied in literature. A collection of state-of-the-art techniques has been developed as part of the voice-privacy challenge Competing anonymization approaches are evaluated in terms of utility and privacy with both objective and subjective metrics. All approaches fall into two main categories, corresponding to the two baseline pipelines provided by the organizers. The most promising approach consists in synthesizing a new audio sample after having extracted and replaced the identity features (x-vectors) from the original sample. The alternative approach consists in using traditional signal processing techniques to distort the audio samples. With these approaches, evaluating potential privacy leaks is not trivial. Interestingly, the AltVoice project has taken a different path, with a strong focus on security features. In short, AltVoice converts speech into text, and then synthesizes a new voice based on the text and on an artificial identity. While many security guarantees can be provided for the artificial identity, AltVoice suffers from a total loss of other voice features, for example, emotional content.

Our Approach

In this project, we aim at improving current approaches in two directions. First, from a security and privacy perspective, we want to define a clear threat model and maximize the protection of customer's identity. We also want to understand the fundamental and concrete limits of the privacy guarantees that the anonymization approach can provide. Second, from the perspective of the final application, we want to minimize the loss of sentimental information caused by the anonymization.

Researchers

Simon Spangenberg (ETH), Dr. Kari Kostiainen (ETH), Dr. Giovanni Camurati (ETH), Marta Tolos Rigueiro (Zurich)

Industry partner

ZURICH

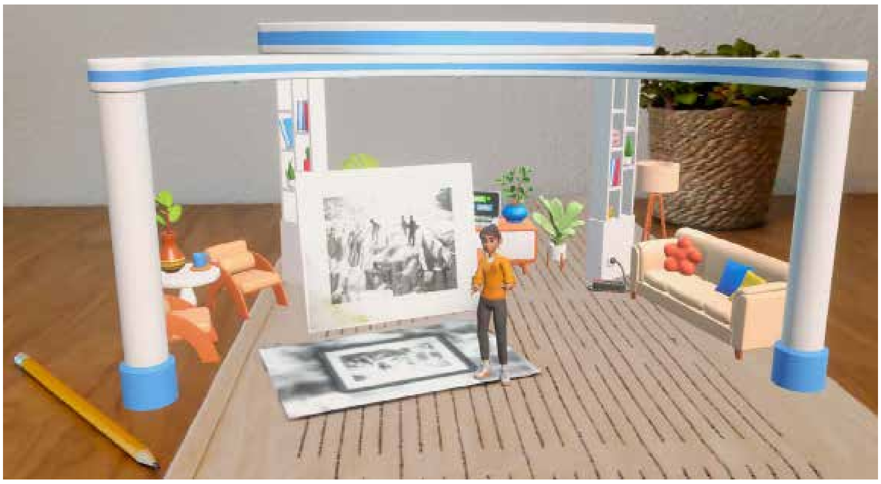

Enhancing Art Engagement with Character-Driven Augmented Reality

The Swiss Post has a long-standing commitment to art and has been collecting works of contemporary art since 1924. Their collection now encompasses over 400 works that are of special relevance to Switzerland and the Swiss population. However, despite the significance of the collection, making it available to the Swiss population remains a challenge. Trends in modern building architecture focus on open plan work spaces and glass materials that limit possibilities for hanging art. And, even when art can be displayed, it is not convenient for most citizens to access it. To address this situation, the Swiss Post and the ETH Game Technology Center engage in a research collaboration to make the Swiss Post's art collection more visible, accessible, and engaging for the Swiss population.

The project aims to develop a prototype application for mobile devices that utilizes novel augmented reality (AR) technology to bring artworks of the Post seemingly into the user's home. On a tablet or smartphone, the user will be able to collect virtual artworks, and view and explore them on the device through AR. A virtual character is present in the scene and explains the art to the user and lets the user interact with it. This animated character will provide deeper context for adults while also making the experience more engaging for children.

Researchers

Börge Scheel (ETH)
Fraser Rothnie (ETH)
Dr. Fabio Zünd (ETH)
Prof. Robert W. Sumner (ETH)
Diana Pavlicek (Die Post)

Industry partner

SWISS POST

Startup Companies

The companies founded by ZISC researcher are listed here*

*Infineon acquired 3db Access AG (3db) in Fall 2023

Affiliated Faculty Members

The ZICS center works in close collaboration with the following ETH faculty members:

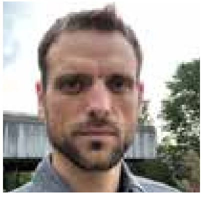

Prof. Stelian Coros leads the Computational Robotics Lab whose research is about robots who understand the physical world and function as skilled co-workers and trusted social companions.

Prof. Hans Gersbach is a professor of Macroeconomics: Innovation and Policy at D-MTEC. His joint research with ZISC focuses on secure governance schemes through assessment voting and vote delegation.

Prof. Christian Holz leads the Sensing, Interaction and Perception Lab whose research covers topics ranging from technical computer-human interaction to wearable sensing and virtual reality.

Prof. Peter Müller leads the Chair of Programming Methodology where the main research objective is to enable programmers to develop correct software.

Prof. Laurent Vanbever leads the Networked Systems Group whose goal is to make the current and future networks, especially the Internet, easier to design, understand and operate.

Prof. Ce Zhang leads the DS3Lab whose research focuses on building data systems to support machine learning and help facilitate other sciences.

Further Information

For more information:
https://zisc.ethz.ch/

How to find us:

Postal address

ETH Zurich
Department of Computer Science
Zurich Information Security and Privacy Center
Universitätstrasse 6
CAB/CNB F
8092 Zurich

Physical address

Entrance to CNB building

ETH Zurich
Department of Computer Science
Zurich Information Security and Privacy Center
Unversitätstrasse 6
Buildings CNB and CAB, floor F (ZISC OpenLab F100.9)
8006 Zurich
Schweiz

phone +41 (0)44 632 86 89

Contact

https://zisc.ethz.ch/

ETH Zurich
Department of Computer Science
Zurich Information Security and Privacy Center
Universitätstrasse 6
CAB/CNB F
8092 Zurich
Schweiz