ETH zürich

Zurich Information Security and Privacy Center (ZISC)

Annual Review 2022

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

Welcome

During the year 2022, the ZISC center continued to deliver excellent results on both of its main mandates: applied research projects that are jointly defined and customized to the needs of our industry partners, and long-term basic research.

Regarding applied research projects that are defined and together with our partners, we worked on multiple topics during 2022. To mention a few examples, the ZISC researchers conducted research on how to improve network security through programmability together with armasuisse and developed novel blockchain-based payment methods in collaboration with NEC. The research work on highly-available communication for financial networks was continued together with SIX and ZKB, as part of the larger deployment of the SCION Internet deployment. With Swiss Post, we continued to explore the problem of phishing and challenges related user training, and together with Zurich Insurance, we continued our research on privacy-preserving machine learning and its applications in the insurance industry. More about such applied and collaborative research projects will be explained in the following pages of this report.

The ZISC researchers worked also on research projects that address fundamental challenges in information security and privacy. To mention a couple of example, the research group lead by Prof. Ueli Maurer explored dynamic tradeoffs in cryptographic protocols and the work by Prof. Kenny Paterson and his group created new research results in the area of searchable encryption. You can read more about such research projects and research highlights later in this report.

The ZISC center also made substantial contributions to projects that have societal importance beyond academia. For example, The ZISC faculty members Prof. David Basin, Prof. Srdjan Capkun, and Prof. Adrian Perrig continued the collaboration with the International Committee of Red Cross (ICRC) in a project that is, informally said, defining "red cross for cyberspace". Researchers lead by Prof. Kenneth Paterson analyzed popular cloud encrypted storage services and identified critical vulnerabilities, while researchers lead by Prof. Srdjan Capkun studied commercial ranging technologies and demonstrated practical distance reduction attacks. Such projects are examples of research, where the ZISC researchers have shown how secure, or insecure, popular systems and services used by millions of people every day actually are. Such research also paves the way for practical development of more secure systems.

During 2022, the ZISC center also continued its long-standing ZISC lunch seminar tradition with a series with bi-weekly research talks that are accompanied by a social lunch gathering. We were privileged to host world-renowned researchers such Virgil Gligor from Carnegie Mellon University and Jens Groth from DFINITY as our seminar speakers. After a few years of the pandemic, it was great to see the excitement when people where able to meet for the lunch seminar in person again.

The ZISC center also continued to grow in 2022. Prof. Florian Tramèr joined the ZISC Faculty as an Assistant Professor in August. His research focuses on the safety and privacy of machine learning which is a highly topical and exciting research area. We warmly welcome Florian to the center and look forward to working with him in the years to come.

During 2022, the ZISC researchers won numerous awards for their work. To mention a few examples, the research paper titled "Victory by KO: Attacking OpenPGP Using Key Overwriting" won the Distinguished Paper Award at CCS'22, the research paper titled "Four Attacks and a Proof for Telegram" won the Distinguished Paper Award at IEEE S&P'22, the SCION research paper won the prestigious Test-of-Time Award at IEEE S&P'22, and the paper "Automating Cook Consent and GDPR Violation Detection" won the Distinguished Artifact Award at USENIX Security'22. The Platypus research project was nominated for the ETH Spark Award as one of the top inventions with commercial potential at ETH Zurich in 2022.

The ZISC center wishes all its partners and collaborators a relaxing holiday season and we are looking forward to working with you again in 2023!

About ZISC

Information Society of Tomorrow

The world is undergoing a dramatic transformation from the industrial society of the 20th century to the information society of the 21st. New information technologies and services emerge at a rapid pace and these innovations have a significant impact on our social, political, and economic lives. The change does not come without risks. Interruption of services can threaten lives and properties, corruption of information can disrupt the work of governments and corporations, and disclosure of secrets can damage individuals as well as institutions. These threats are no longer limited to hobbyists hackers; instead we witness attacks from organized crime, terrorists and governments. To counter such risks in the constantly evolving information technology landscape, we need a thorough understanding on the theoretical foundations of information security, as well as practical attacks and countermeasures.

Research Center

The Zurich Information Security and Privacy Center (ZISC) is an industry-supported research center of ETH Zurich, founded in 2003. The goal of ZISC is to bring academia and industry together to solve the information security challenges of tomorrow. In ZISC, PhD students and senior researchers perform academic research under the supervision of ETH Zurich faculty members. Many ZISC research projects are done in co-operation with an industry partner.

Education

Besides research, ZISC provides world-class academic education in information security. This includes training through projects, classes at ETH Zurich, and workshops for ZISC researchers and industry partners.

Why a Security Center in Zurich?

Zurich is a center of global banking and insurance, two industries that have particularly strong security needs and whose success inherently depends on their reputation as being secure. Zurich also hosts many leading technology companies that develop novel security and privacy solutions. Finally, Zurich is centrally situated in the heart of Europe. The goal of ZISC is to establish a critical mass of information security talent and research in Zurich that benefits academia, economy and society.

Partners

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

Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra

armasuisse

NEC

Zürcher Kantonalbank

ZURICH

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 Privacy and Security Lab whose research currently focuses on understanding and improving the worst-case behavior of machine learning systems.

Research Highlights 2022

Dynamic trade-offs in Cryptographic Protocols

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 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)

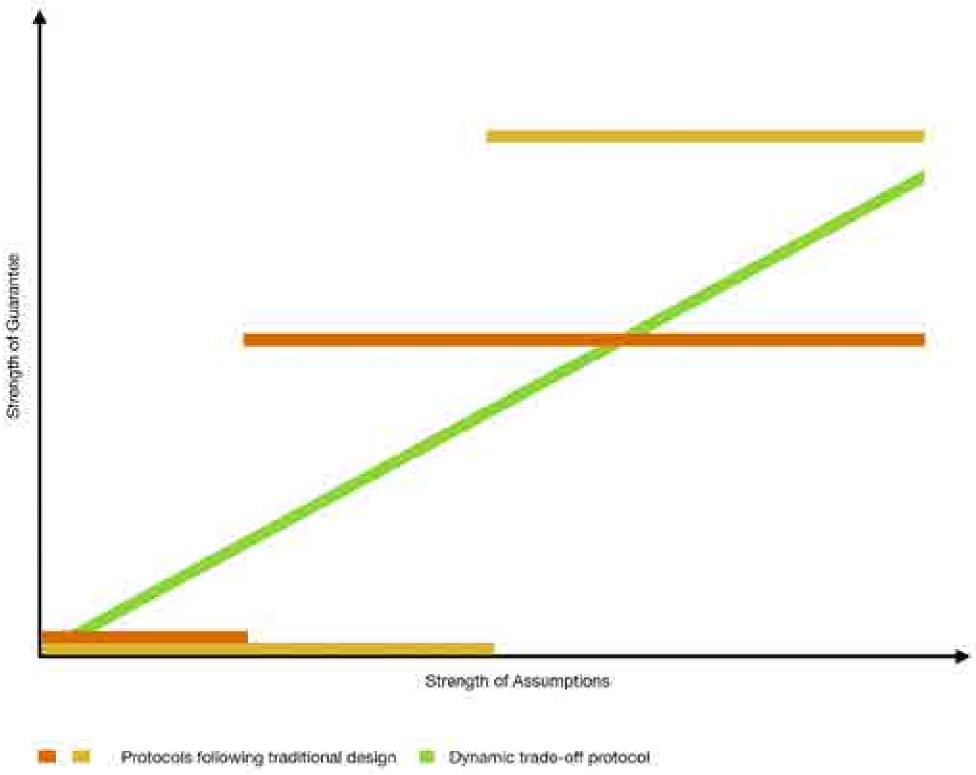

Research Highlights 2022

Security Analysis of MEGA Cloud Storage

Introduction: The clouds are here to stay. Offloading data storage to the cloud has become increasingly prevalent in recent years, and much of our personal data is now uploaded by default from our devices to the servers of companies such as Google, Microsoft and Apple. Users also choose to place their files in the cloud to benefit from features such as multi-device access, data back-up, sharing and collaboration. The price consumers pay for the (often free) storage and services is the loss of privacy; the providers gain direct access to the data of their customers and can use it to perform analytics, train algorithms and serve targeted advertising.

Together with the rising awareness of privacy issues, the public demand for privacy-preserving services has risen. In response to this, a number of providers have appeared which offer "secure cloud storage", in the sense that user files are encrypted by the users before they are offloaded to the cloud. The largest such provider to date is MEGA. While still comparatively small in contrast to e.g. Google Drive, they have a sizeable consumer base of more than 250 million registered users, with over 10 million active users daily and a total storage exceeding 1000 petabytes (or 1 million TB) [1].

MEGA advertises themselves as "the privacy company" and state on their website that "by properly applying end-to-end encryption, MEGA achieves actual privacy by design". Further, they say that their system has been designed with "user-controlled end-to-end encryption", a term which they explain with "All your data on MEGA is encrypted with a key derived from your password; in other words, your password is your main encryption key. MEGA does not have access to your password or your data." In summary, MEGA aims to provide an alternative to unencrypted cloud storage services, in which the confidentiality and integrity of user data is guaranteed even in a threat model where the provider itself is not trusted.

End-to-end encrypted cloud storage. Providing secure cloud storage in the untrusted provider-setting comes with a slew of challenges. First of all, users must be responsible for the management of the keys used to encrypt their files. This, already, is an issue whose difficulty should not be underestimated. Given the lack of understanding of cryptography among the general population, as well as the continuous evidence that users find it difficult to handle even human-memorable secrets such as passwords, entrusting users with the care of cryptographic keys is asking a lot. Second, to provide features such as multi-device access, the encryption keys must be synced between user devices. To avoid tasking the customers with out-of-band key transport, this syncing should ideally be performed by the untrusted cloud service, without violating the end-to-end encryption guarantees.

At the time of writing, there is still no standard for encrypted cloud storage to provide guidance on how to overcome these – and other – challenges. Hence every provider who wishes to offer this type of protection for their customers must design their own protocol, including MEGA. In recent work, a team of researchers from the Applied Cryptography Group analyzed the cryptographic design of MEGA and found several severe issues. Together, the discovered vulnerabilities lead to a complete break of the confidentiality and integrity of user files in the MEGA cloud storage in the untrusted provider setting.

The Attacks: MEGA's cryptographic design.
MEGA's approach to end-to-end encryption starts with the user password, from which a cryptographic key K is derived using the password-based key derivation function PBKDF2. Independently, a number of cryptographic keys are locally generated for the user, including a 128-bit AES key called the master key, an RSA key pair called the sharing keys and a set of 128-bit AES node keys; one for each file uploaded by the user. The sharing keys are used to share files between users, and MEGA stores the public keys of all customers.

To allow multi-device access, the master key is encrypted with the key K derived from the password and uploaded to MEGA's servers. This allows the user to sign into their account from any device using their password, rederive k_e and download and decrypt the master key. The master key in turn is used by the client to encrypt all other keys, including the private RSA sharing key and the node keys. The encryption mode used is AES-ECB, without any integrity protection of the resulting ciphertexts. The thus encrypted keys are also offloaded to MEGA for safekeeping.

When a user logs in to their account, they receive the encrypted keys from the server, together with an RSA-encrypted session ID. The client decrypts the private RSA key and uses it to decrypt the

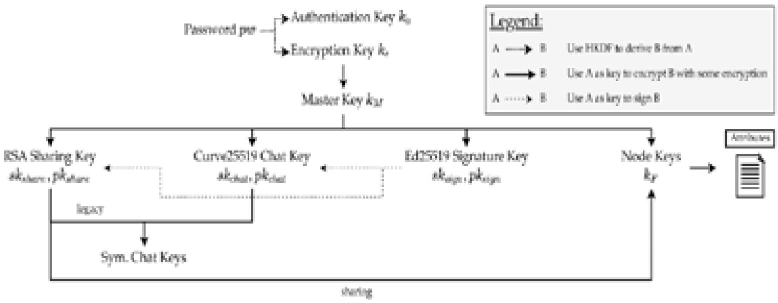

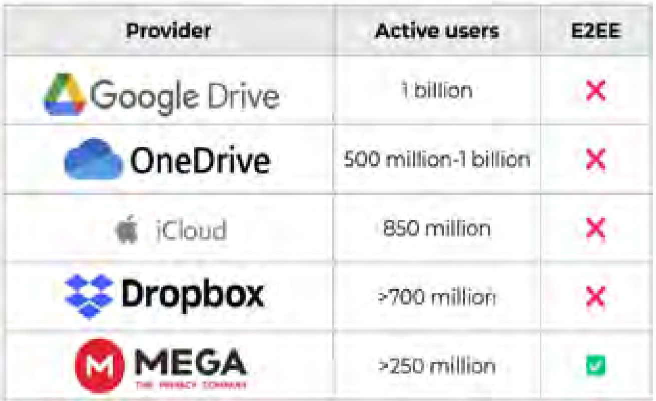

Provider	Active users	E2EE
Google Drive	1 billion	✗
OneDrive	500 million-1 billion	✗
iCloud	850 million	✗
Dropbox	>700 million	✗
MEGA THE PRIVACY COMPANY	>250 million	✓

Figure 1: MEGA's key hierarchy.

Figure 2: Consumer cloud storage providers. Approximate number of users and end-to-end encryption.

session ID, which it then includes in subsequent requests to the server as a form of authentication.

RSA private key recovery.

The team's first attack on MEGA's cryptographic design uses the session ID exchange to recover the private RSA key of a user. The attack exploits the missing integrity protection of the outsourced key material to overwrite the encrypted RSA private key. The malicious service provider then chooses an adversarial value of the session ID, such that the client's response enables a binary search for one of the prime factors of the RSA modulus. With the factorization of the modulus, the adversary can recover the secret exponent and, thus, the complete RSA private key. With the compromised RSA key, the adversary can break the confidentiality of all files shared with the victim. Additionally, knowledge of the RSA private key enables further attacks which eventually allow the adversary to decrypt arbitrary AES-ECB ciphertexts encrypted with the user's master key.

AES-ECB plaintext recovery.

With knowledge of the RSA private key, an attacker can once again use the session ID exchange to instantiate a decryption oracle for AES blocks encrypted with the master key. The attack recovers two blocks of plaintext per login attempt by the user. To do this, the adversary overwrites part of the RSA private key ciphertext with the two target ciphertext blocks. Using properties of RSA-CRT – the method by which the RSA private key is decrypted by the client – the plaintext corresponding to the two target blocks can be recovered from the session ID sent back to the server by the victim. Since all user keys are encrypted with AES-ECB under the master key, an attacker exploiting this vulnerability can decrypt all user keys. For instance, one client login suffices to recover a node key after the RSA private key was recovered, after which the corresponding file can be trivially decrypted. Hence, this attack completely breaks the confidentiality of user data.

Additional vulnerabilities.

In addition to the two attacks described above, the research team's analysis also showed that the format used by MEGA to store encrypted node keys allows an attacker to overwrite a node key with an all-zero key, without the client noticing. This allows an adversary to break the integrity of user data, since the attacker can encrypt a file of their choosing with the all-zero key and then upload it to the user's storage together with the modified node key. The unsuspecting user will decrypt the file to find potentially harmful or compromising material.

This integrity attack can be made even less detectable by more careful use of the AES-ECB decryption oracle from the second attack. By decrypting an arbitrary AES block encrypted under the master key, the adversary can "create" a node key, which it can then use to encrypt arbitrary files of its choosing and place in the user's storage. This attack is completely undetectable and could be used to, for example, frame a user with possession of illegal material.

Lastly, a variant of Bleichenbacher's attack [2] adapted to the padding format used by MEGA for RSA encryption allows an attacker to recover plaintexts without knowing the RSA private key. Hence, this attack complements the first attack and allows the adversary to recover (among other things) node keys shared with the user, without performing the RSA key recovery attack.

For more details on the attacks, see https://mega-awry.io/ and the resulting research paper [3].

Disclosure and migitations.

The research team disclosed the vulnerabilities to MEGA on March 24, 2022. They proposed extensive mitigations to protect against the discovered attacks. MEGA patched their system in June, choosing a different fix to the ones proposed, which nonetheless prevent the direct chain of attacks described here. Time will tell how robust the selected fix is.

The coordinated publication of the attacks and fix saw significant interest from the community, including press coverage in Ars Technica, The Register and heise.de, and discussion on social media.

Further information

Website: https://mega-awry.io/
Publication:
"MEGA: Malleable Encryption Goes Awry", IEEE Symposium on Security & Privacy, to appear, 2023.

Researchers

Prof. Kenny Paterson, Matilda Backendal, Miro Haller. Applied Cryptography Group, ETH Zurich.

[1] https://mega.io/

[2] D. Bleichenbacher. Chosen Ciphertext Attacks Against Protocols Based on the RSA Encryption Standard PKCS #1. CRYPTO 1998: 1-12.

[3] M. Backendal, M. Haller and K.G. Paterson. MEGA: Malleable Encryption Goes Awry. IEEE Symposium on Security & Privacy 2023, to appear. Paper available from https://mega-awry.io/.

Research Highlights 2022

Rethinking Searchable Symmetric Encryption

Searchable Symmetric Encryption

Database encryption is a key enabler for secure storage-as-a-service, wherein clients can securely outsource the storage and processing of large databases to (potentially untrusted) third party servers. Searchable symmetric encryption (SSE) is a special subclass of database encryption that aims to efficiently support search queries over symmetrically encrypted databases. The core functionality enabled by SSE is the following: given an encrypted document collection in which each document is tagged with keywords, find the set of all documents tagged with a given keyword w.

Structure-only vs End-to-end

There are two main design paradigms in constructing an SSE scheme: structure-only or end-to-end. In the structure-only approach, one focuses on building a secure and efficient search index, which the client can query with any keyword w and retrieve all document identifiers for documents that contain the keyword w. These document identifiers can then be used to fetch the actual documents in a subsequent operation. The retrieval of the actual documents is not part of a structure-only scheme. Instead, one assumes that there is a secure and efficient way to retrieve the actual documents given the document identifiers. Most schemes in the literature take this approach. On contrary, in the end-to-end approach, one builds a scheme that supports the retrieval of the actual documents natively.

Leakage

Most of the SSE constructions achieve the efficiency goal by allowing the server to learn some amount of information about either the database itself or the queries made by the client. This information is typically called "leakage" in the literature. We say that a scheme is secure if its leakage is benign (e.g. if the scheme leaks the size of the database and nothing else).

Leakage-abuse Attacks

Crucially, the structure-only schemes were only proven secure with respect to the search index but not the whole system (i.e. a system that includes the document retrieval step too). This naturally raises the following question:

Do structure-only SSE schemes result in secure end-to-end SSE systems when system-wide leakage is taken into account?

In our work, we answer this question in the negative. We show that all the structure-only SSE schemes, including the state-of- the-art ones, incur damaging system-wide leakage when used to construct end-to-end SSE systems in the natural way. Concretely, we show an efficient and highly scalable query recovery attack targeting the vast majority of end-to-end SSE schemes built from the existing structure-only schemes and a naive document retrieval protocol.

System-Wide Leakage Migitation

A natural approach to mitigating system-wide leakage in end-to-end SSE systems built from structure-only schemes is to deploy a more secure document retrieval protocol.

Intuitively, this enhances resistance to system-wide attackers, but potentially degrades efficiency. In this context, we now ask the following question:

Can existing leakage suppression techniques efficiently mitigate system-wide leakage?

We also answer this question in the negative. We demonstrate experimentally that it is practically infeasible to use the structure-only schemes themselves to build an efficient document retrieval protocol. We also rule out the feasibility of building an efficient document retrieval protocol from other data retrieval primitives such as oblivious RAM and private information retrieval.

Rethinking SSE

On one hand, our query reconstruction attacks on end-to-end SSE schemes built naively from structure-only schemes demonstrate that structure-only schemes are insufficient in building an end-to-end SSE scheme. On the other, our experiments on the alternative document retrieval protocols suggest that we do not have a secure and efficient construction for SSE yet.

This leads us to believe that we need to:

- Rethink security definitions for SSE while taking into account the impact of system-wide leakage.
- Rethink how we construct an SSE scheme. An SSE scheme should be built as a system-wide secure scheme from the beginning.

Further information

Publication: "Rethinking Searchable Symmetric Encryption", IEEE Symposium on Security & Privacy, to appear, 2023. Preprint available from: https://eprint.iacr.org/2021/879

Researchers

Dr. Zichen Gui (Applied Cryptography Group, ETH Zurich)
Prof. Kenneth Paterson (Applied Cryptography Group, ETH Zurich)
Dr. Sikhar Patranabis (formerly with the Applied Cryptography Group, ETH Zurich, now with IBM Research India).

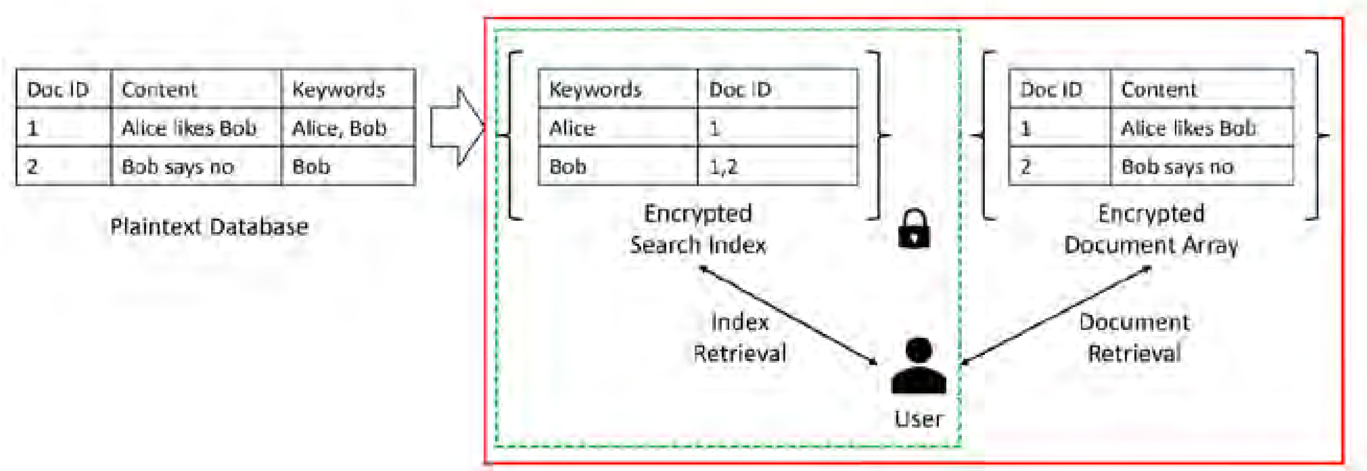

Doc ID	Content	Keywords
1	Alice likes Bob	Alice, Bob
2	Bob says no	Bob

Keywords	Doc ID
Alice	1
Bob	1,2

Doc ID	Content
1	Alice likes Bob
2	Bob says no

Green dotted box: construction and security considered by structure-only SSE schemes
Red solid box: end-to-end SSE schemes built from a naive extension of structure-only schemes are broken in this system-wide view.

Research Highlights 2022

Elasticlave: An Efficient Memory Model for Enclaves

Trusted Execution Environments
Isolation, commonly through the use of the process abstraction provided by an OS, is a cornerstone for security. It allows us to isolate and limit software compromises to one fault domain within an application and is the basis for applying the design principle of privilege separation.

Trusted execution environments (TEEs) isolate user-space applications into secure enclaves without trusting the OS. In the last few years, user-level enclaves have become available in commodity CPUs that support TEEs. Conceptually, enclaves are in sharp contrast to processes in that they do not assume a trusted OS, promising a drastic reduction in the trusted computing base (TCB) of a fault domain. The enclaved TEE design is of fundamental importance to security because they offer a new isolation primitive for software.

Need for a better memory model
Existing TEE memory models are rigid—they do not allow an enclave to share memory with other enclaves. This lack of essential functionality breaks compatibility with several constructs such as shared memory, pipes, and fast mutexes that are frequently required in data intensive use-cases. We revisit one of the key abstractions provided by enclaved TEEs—their memory model. Several existing TEEs follow what we call the spatial isolation model.

When applied to enclaves, the spatial isolation model is a simple but rigid model that is insufficient for memory sharing. Its underlying principle breaks compatibility with the most basic of data sharing patterns where the enclave needs to compute privately on some data before making it public or sharing it externally.

If we want to support memory sharing between enclaves on spatially isolated memory, we require additional trusted coordinator enclaves together with cryptographic secure message passing channels. Without these additional mechanisms, achieving secure shared memory is fraught with challenges in managing ownership and access rights of the shared region, as attack vectors like permission re-delegation, confused deputy, malicious races, and TOCTOU attacks have shown.

Secure and flexible memory model
Elasticlave is a new TEE memory model which allows sharing by striking a balance between security and flexibility in managing access permissions. It allows enclaves to share memory across enclaves and with the OS, with more flexible permissions than in spatial isolation. While allowing flexibility, it does not make any simplistic security assumptions or degrade its security guarantees over the spatial isolation model. We view enclaves as a fundamental abstraction

for partitioning applications in this work, and therefore assume that enclaves do not trust each other and can become compromised during their lifetime. Our implementation of Elasticlave on RISC-V achieves performance overheads of about 10% compared to native (non-TEE) execution for data sharing workloads. In contrast, a similarly secure implementation on a rigid TEE design incurs 1-2 orders of magnitude overheads for these workloads. Thus, Elasticlave enables cross-enclave data sharing with much better performance.

Further information

Website
https://github.com/jasonyu1996/elasticlave

Publication:
Elasticlave: An Efficient Memory Model for Enclaves
Jason Zhijingcheng Yu, Shweta Shinde, Trevor E. Carlson, Prateek Saxena
In Proceedings of the USENIX Security Symposium (USENIX Security 2022).

Researchers

Prof. Shweta Shinde

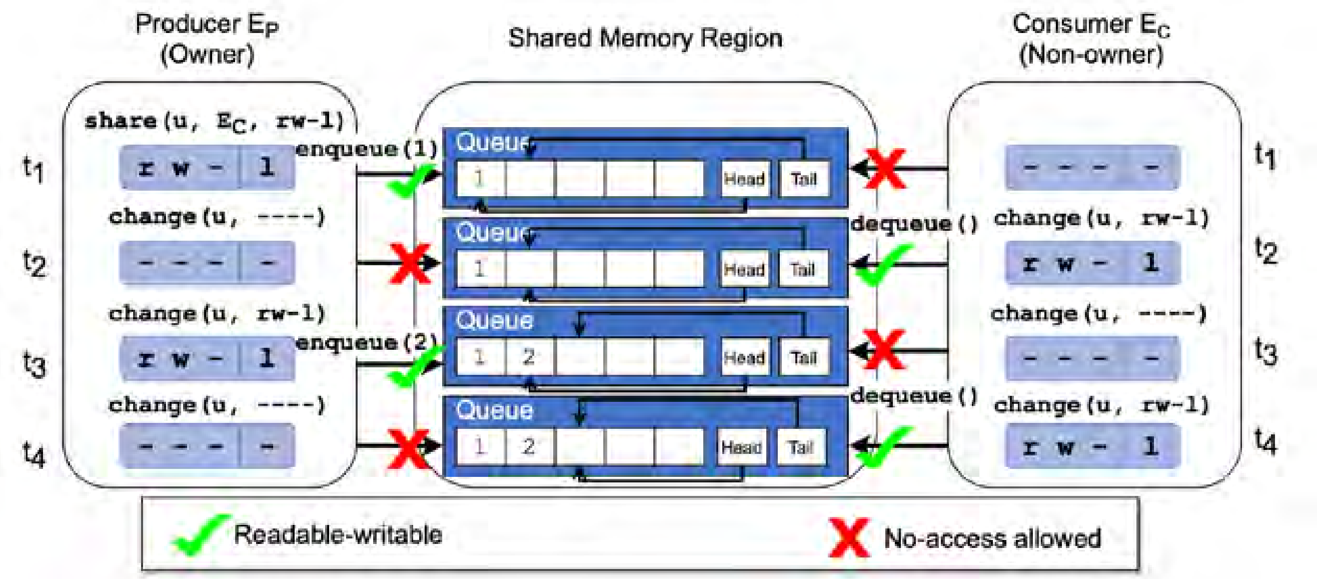

Research Highlights 2022

Evaluating DoS Susceptibility of VPN Implementations

In today's Internet ecosystem, enterprise networks often stretch across several locations, such as corporate branches, data centers, and infrastructure hosted by cloud providers. VPN systems are an integral part of these setups, serving as the glue that securely connects the different locations by encrypting and authenticating traffic between pairs of endpoints over an untrusted network such as the Internet.

Given this role, VPN protocols are as ubiquitous in emerging SD-WAN deployments as they are in more traditional site-to-site connections. Moreover, the recent shift to remote work has led to a surge in VPN use, as they are also often deployed as an access control mechanism to protect network segments and services. Since VPN endpoints are usually exposed to the public Internet, they represent an attack surface for adversaries. Combined with the critical reliance on the endpoints, they are an attractive target for DoS attacks.

All major VPN protocols integrate DoS defenses, such as source-binding cookie mechanisms. However, developing a VPN implementation with strong DoS resilience is challenging. While prior studies analyzed the performance of different VPN implementations, their measurements do not consider adversarial settings.

We addressed this lack of adversarial testing by evaluating the DoS resilience of four widely used state-of-the-art VPNs: WireGuard, strongSwan (IPsec), OpenVPN, and Cisco Any Connect (IPsec and SSL). To obtain quantitative measurements of each VPN's resilience, we have implemented various attacks and measured their impact on a high-performance server. The results were surprising: a few hundred Mb/s of attack traffic are sufficient to deny essential functionalities of any evaluated VPN implementation.

Through an in-depth performance analysis of the implementations, we identified significant inefficiencies primarily related to multi-core synchronization. Furthermore, we discovered vulnerabilities in strong Swan's implementation, which an attacker can easily exploit to block the establishment of new connections. Additionally, we found a vulnerability in the OpenVPN protocol that allows the disruption of an already-established connection with a single attack packet.

These findings have important implications for real-world deployments on various infrastructures, including critical infrastructures, rely on functioning VPNs. Additionally, the insights highlight that rigorous adversarial testing is crucial for creating more DoS-resilient network setups.

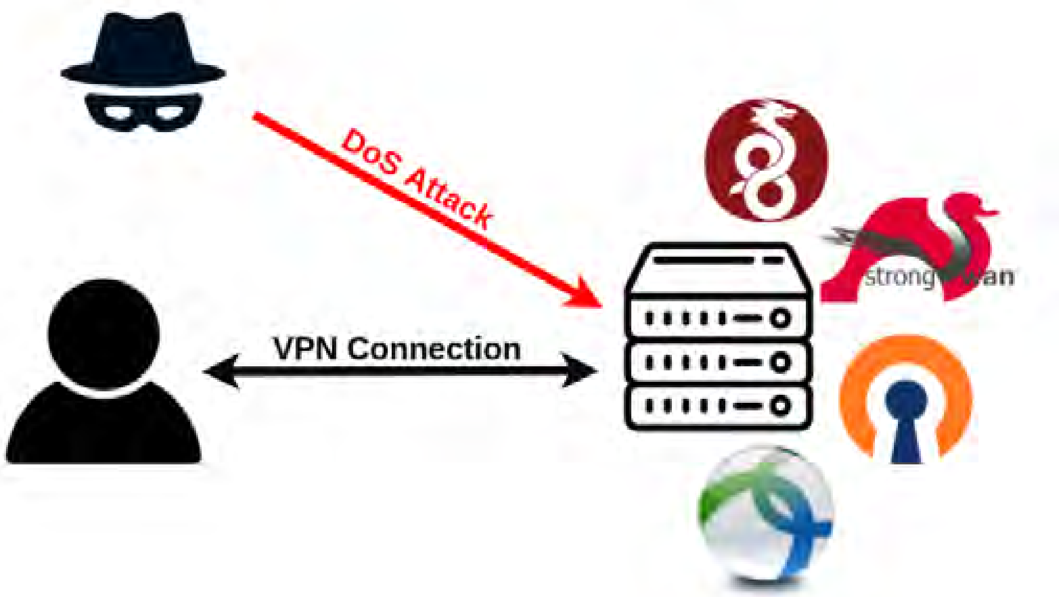

Further information

Recent Publication:
Evaluating Susceptibility of VPN Implementations to DoS Attacks Using Adversarial Testing.
Fabio Streun, Joel Wanner, and Adrian Perrig
in the proceedings of Network and Distributed System Security (NDSS) 2022

Researchers

Fabio Streun, Joel Wanner, Prof. Adrian Perrig

Research Highlights 2022

From data poisoning to privacy leakage in machine learning

A central tenet of computer security is that data privacy is often impossible without some form of data integrity. In cryptography, for example, there are many attacks that succeed in decrypting ciphertexts by leveraging an adversarys ability to modify the ciphertext in transit. In a recent collaboration with Google, the National University of Singapore, and Oregon State University, we showed that similar vulnerabilities apply to the training of machine learning models.

Our project connects two long and independent lines of work that study attacks on the integrity and privacy of training data in machine learning. Modern machine learning models are often trained on extremely large datasets, which are aggregated from a multitude of potentially untrusted and/or sensitive sources. In a **poisoning attack**, an adversary targets the integrity of a models data collection process to degrade model performance. Such an adversary has the ability to inject malicious data into a training set, and might abuse this ability to fool the model into misclassifying specific inputs. In a **privacy attack**, an adversary instead aims to extract private information about the model's training set by interacting with a trained model. Such attacks exploit the tendency of machine learning models to memorize specificities of the data that they were trained on.

Our work asks whether an adversary can exploit the ability to poison individual training samples in order to maximize the privacy leakage of other unknown training samples. In other words, **can an adversary's ability to «write» into the training dataset be exploited to arbitrarily «read» from other (private) entries in this dataset?**

We show that this is indeed the case! An attacker that can control a very small fraction of training data points can cause the leakage of other users› data to grow by multiple orders-of-magnitude. As a concrete example, consider the case of a company training a large language model on text written by its end-users (e.g., to build a customized client chatbot). We show that if an adversary injects 64 special sentences into the training data of this model, they can boost their ability to extract secret number sequences (e.g., credit card numbers, or phone numbers) contained in the training set, by a factor of 40x.

The main take-away from our work is that worst-case privacy guarantees should matter to everyone. Indeed, prior work had found that it is mainly data outliers that are at risk of privacy attacks. In, other words, users whose data is not «abnormal» seemed to be safe, by virtue of a «hiding-in-the-crowd» effect. Yet, being an outlier is a function of not just the user›s own data, but also its relation to other users› data in the training set. What our work shows is that only a small number of poisoned points suffice to transform any users data into an outlier, whose data privacy is then put at risk.

On the positive side, we find that privacy defenses that aim at minimizing the information leakage of outlier data (such as differential privacy) can effectively mitigate our attacks. Such defenses are thus warranted for machine learning models deployed in settings where training data can be both untrustworthy and privacy sensitive.

Further information

Recent Publication:
Truth Serum: Poisoning Machine Learning Models to Reveal Their Secrets.
Florian Tramèr, Reza Shokri, Ayrton San Joaquin, Hoang Le, Matthew Jagielski, Sanghyun Hong, Nicholas Carlini.
Proceedings of the ACM Conference on Computer and Communications Security (CCS). November 2022.

Researchers

Prof. Florian Tramèr

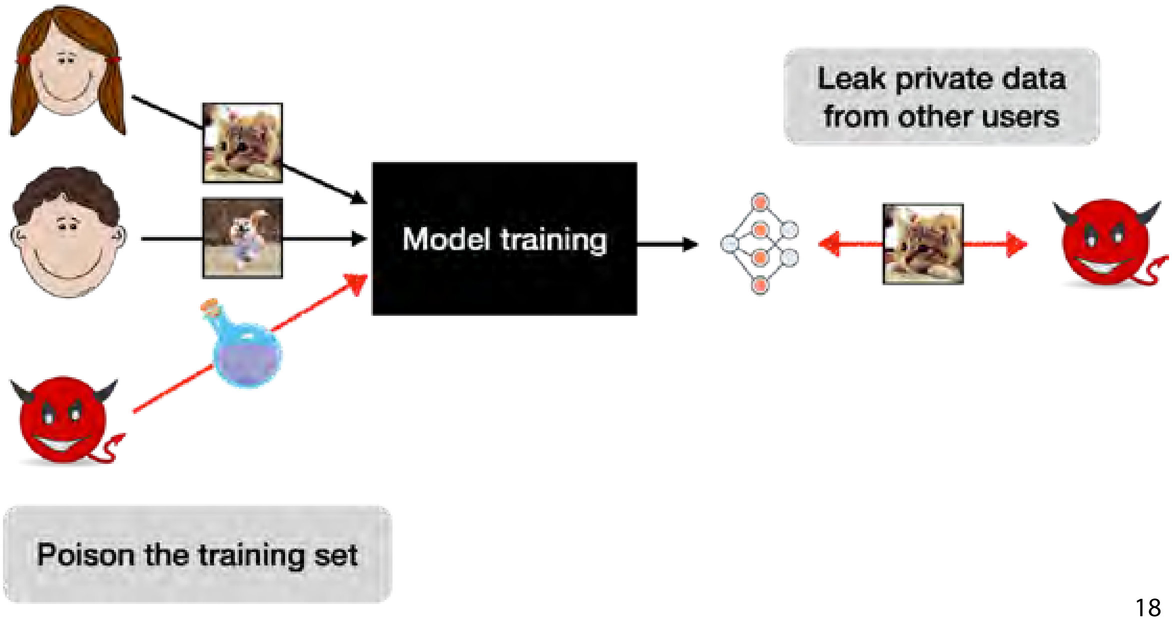

Research Highlights 2022

Ghost Peak: Practical Distance Reduction Attacks Against HRP UWB Ranging

Secure ranging consists in measuring the distance between two devices correctly even in presence of an external attacker. One of its main applications is distance-based access control. For example, the Passive Keyless Entry and Start (PKES) system of a car only grants access when the owner is at short distance. Secure ranging is also useful for mobile payments, asset tracking, and other applications requiring secure and reliable distance measurements.

Ultra Wide Band (UWB) chips measure the distance between two devices based on the time-of-flight of radio packets. In a typical ranging sequence, the initiator sends a ranging packet to the responder, which replies after a short time. The distance is then computed as the time-of-flight multiplied by the speed of light.

Recently, devices implementing the High-Repetition Pulse (HRP) mode of the IEEE 802.15.4z standard have been widely deployed in consumer devices. For example, Apple integrates its U1 chip in most of its products (e.g., iPhone, AirTag, HomePod, Apple Watch). UWB chips are also now available in Android phones (e.g., Samsung Galaxy S21+). Some manufacturers have already made available or announced the use of UWB-enabled phones to conveniently open cars.

At the logical layer, HRP UWB offers promising security features. Using an out-of-band channel (e.g., Bluetooth), the two devices involved in distance measurement establish a shared cryptographic key. Every ranging packet contains an unpredictable field, called Scrambled Time Sequence (STS), which is generated with AES in counter mode using the shared secret key. An attacker, unable to predict the STS, cannot transmit ranging packets earlier than the victim. Clearly, this should prevent distance-reduction attacks that would trick an access control system into falsely believing that the user is close.

Unfortunately, HRP UWB is vulnerable at the physical layer. This problem can be understood by first analyzing the physical layer of a receiver. In realistic conditions, the radio signal is reflected and scattered by objects in the environment, causing multiple copies of the ranging packet to arrive at the receiver at different times and with different power levels. This phenomenon is known as multipath. In addition, the receiver might not be in line of sight. The receiver uses a known template of the expected signal to estimate the Channel Impulse Response (CIR) containing one peak for each copy of the signal. Because of constructive and destructive interference among copies and of obstacles shadowing the line of sight, the earliest peak corresponding to the shortest path between receiver and transmitter is not necessarily the strongest one. Therefore, the receiver first detects the highest peak, which corresponds to the strongest path, and then goes back in time, looking for the earliest peak above the noise. Since the strength of a peak depends both on the number of correct STS bits and on these complex multipath effects, it is fundamentally hard to decide if a small peak comes from a legitimate packet.

An attacker can transmit a packet with high power and random STS bits. The resulting noise in the CIR has a certain non-negligible probability of being

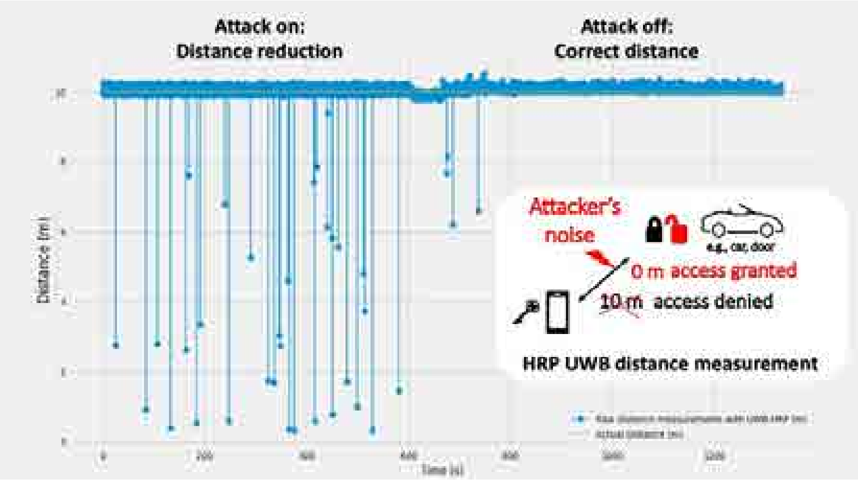

misclassified as a valid peak having arrived earlier ("Ghost Peak"), causing a distance reduction and ultimately breaking access control.

The success probability is considerably higher than one would expect at the logical layer (the probability to guess STS bits correctly). In a recent publication at USENIX Security 2022, we have demonstrated a practical attack on Apple U1 chips, showing distance reductions of up to 12m with up to 4% success rate. Further experiments have proven Qorvo chips vulnerable, too.

As of now, there is no known way to formally prove the security level of HRP UWB receivers. The fact that the standard does not specify any implementation details and the closed-source nature of proprietary implementations further complicate the analysis. For defense, mitigations could be envisioned (e.g., checking the consistency between preamble and STS fields). More fundamentally, efforts should be made such that the new standard (currently under discussion at IEEE) will include a well-defined security mechanism with proven security guarantees.

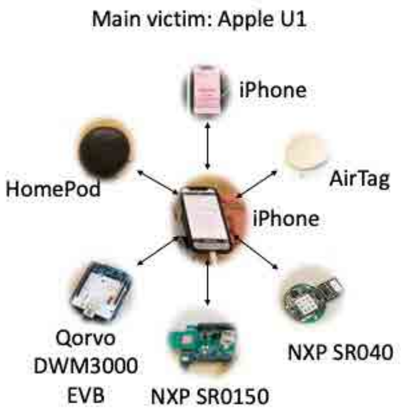

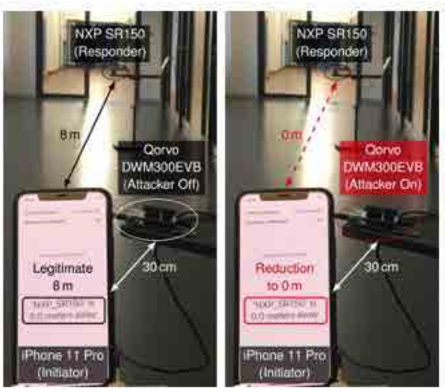
NXP SR150 (Responder)
8 m
Qorvo DWM300EVB (Attacker Off)
Legitimate 8 m
30 cm
iPhone 11 Pro (Initiator)
NXP SR150 (Responder)
0 m
Qorvo DWM300EVB (Attacker On)
Reduction to 0 m
30 cm
iPhone 11 Pro (Initiator)

Website

[https://securepositioning.com/ghost-peak/https://securepositioning.com/ghost-peak/https://securepositioning.com/ghost-peak/(https://securepositioning.com/ghost-peak/)](https://securepositioning.com/ghost-peak/)

Further information

Publication:
Patrick Leu, Giovanni Camurati, Alexander Heinrich, Marc Roeschlin, Claudio Anliker, Matthias Hollick, Srdjan Capkun, Jiska Classen "Ghost Peak: Practical Distance Reduction Attacks Against HRP UWB Ranging". In Proceedings of the 31st USENIX Security Symposium, USENIX Security 2022, Boston, MA, USA, August 10-12, 2022.

Researchers

Patrick Leu*, Giovanni Camurati*, Alexander Heinrich**, March Roechlin*, Claudio Anliker*, Matthias Hollick**, Srdjan Capkun*, Jiska Classen**
* ETH Zurich
** TU Darmstadt

Research Highlights 2022

Collaborations with the International Committee of Red Cross (ICRC)

Prof. Basin's group. Protected Parties (PPs) offer humanitarian services in regions of armed conflict and are granted special protection under international humanitarian law (IHL). They may advertise their protected status by the well-known emblems of the red cross, red crescent, and the red crystal. As part of this project, we proposed a scheme, An Authenticated Digital EMblem (ADEM), to distribute digital emblems, which mark entities as protected under IHL in an analogy to the physical emblems. ADEM avoids the need for a central authority and follows a distributed approach by leveraging certificate chains. We designed ADEM with versatility, usability, and security in mind. It applies to any digital entity, scales to small and large organizations, and both deployment and verification can be automated. Emblems and the respective public keys can be cryptographically verified as authentic, and we hardened our scheme against a wide range of attacks, even against secret key compromise. In 2020, the international committee of the red cross (ICRC) reached out to ETH Zurich to explore the technical feasibility of a digital emblem. As a response, ADEMs development started in late 2020. Early 2021, we finalized a first proposal that received positive feedback from the ICRC. Ultimately, our goal is to standardize ADEM globally to meet the challenge of signaling protection under IHL in the digital realm.

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. 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.

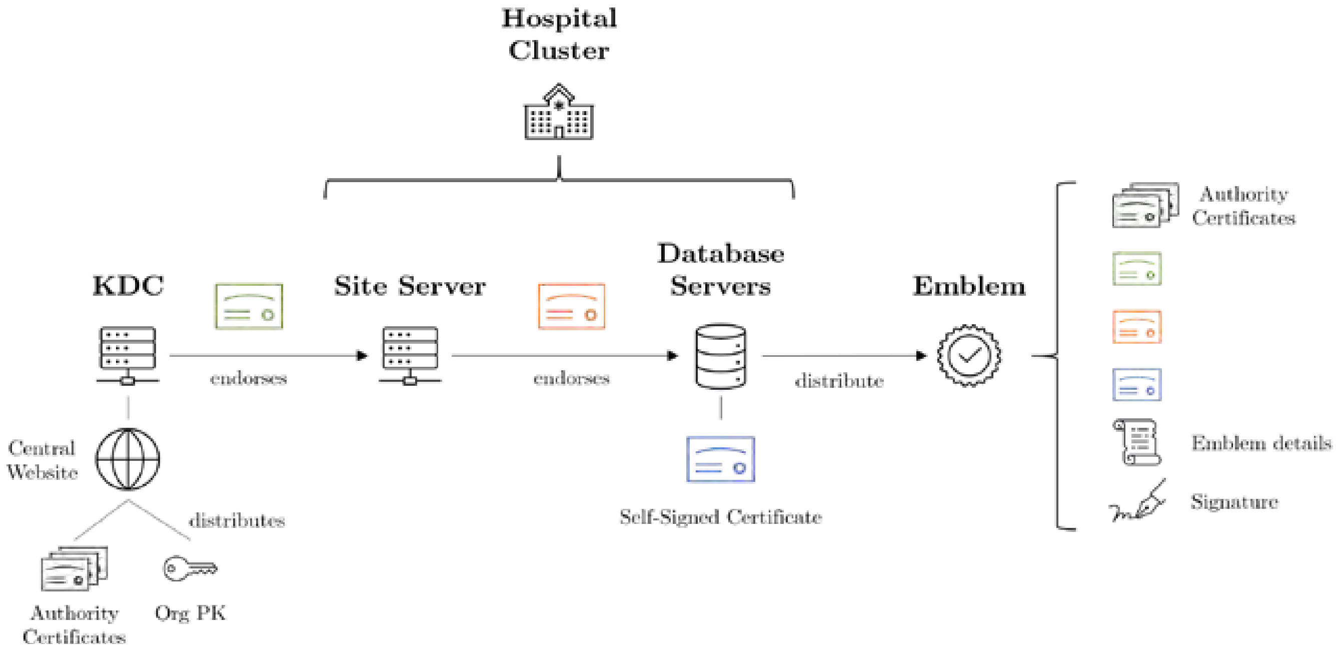

Education 2022

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 Be-aver, 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 «Data Science und Sicherheit» with the focus on Security and Data Science for high school (460 pages with detailed explanations, motivating challenges and projects). 2) Organization of the competition Informatics Beaver (37 000 pupils) with semifinals (850 pupils) and finals (119 pupils) at ETH with several tasks related to security. 3) Teacher training in cryptography for40 high school teachers. 4) 198 schoolprojects (8-20 lessons per class)including security issues. 5) 47 projects(697 lessons) 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 model›s 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.

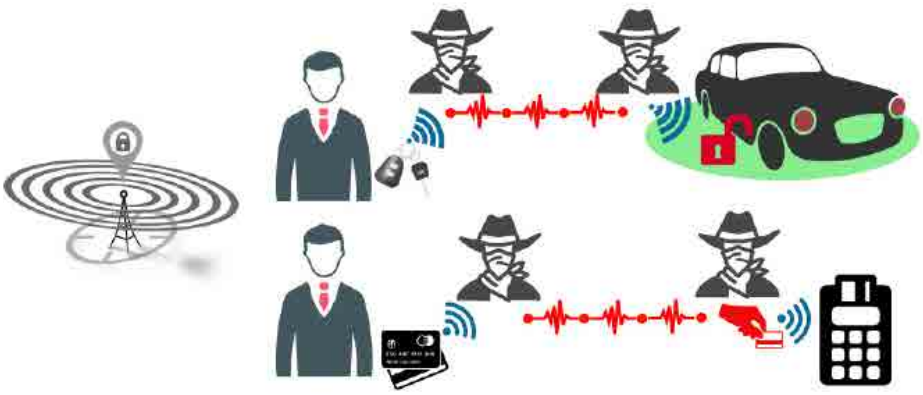

Main Research Areas

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.

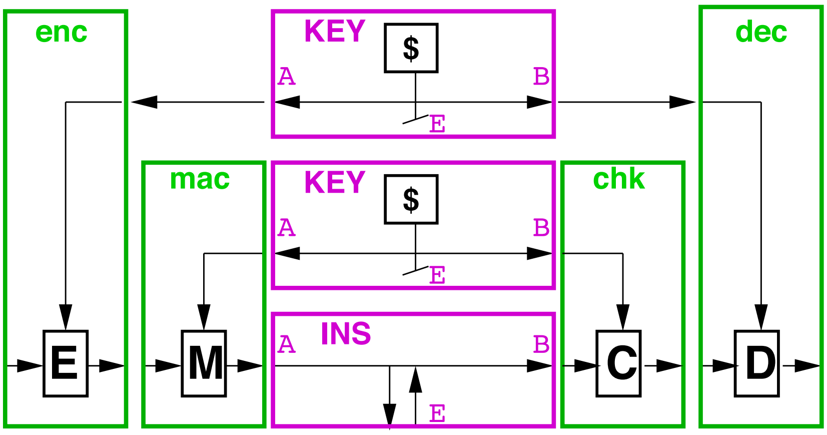

Security protocol verification

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 of the next-generation 5G mobile communication key exchange protocol 5G-AKA.

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.

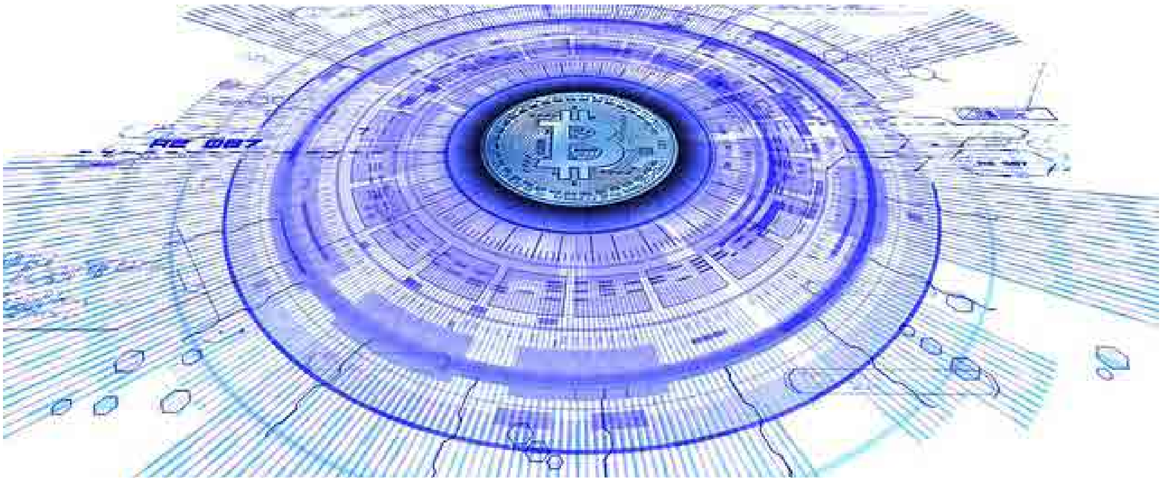

Research Projects

SCION

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

SIX

Zürcher Kantonalbank

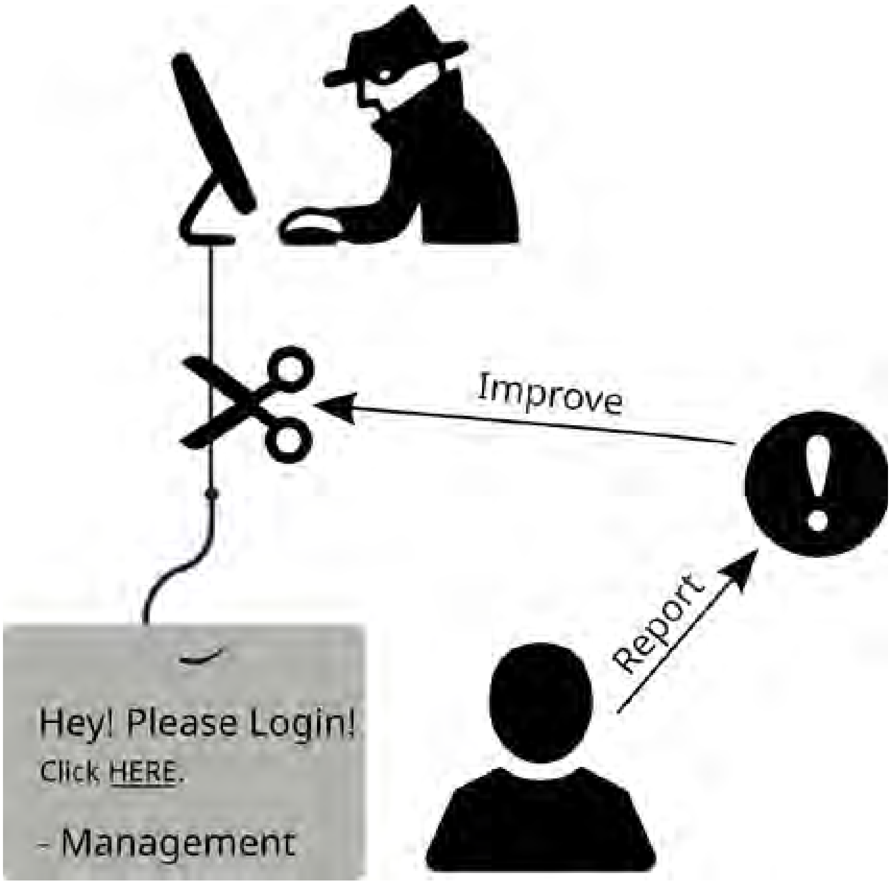

User-Complemented Phishing Protection

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 the user to act swiftly, usually by clicking on a link to change a reportedly compromised password, log in to confirm or update personal data. Such links lead to deceptive websites that are copies of the legitimate ones and 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 for companies that fall victim in the order of million dollars. 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 conduct a large-scale study on phishing prevention, detection, and education. Users will be involved in phishing detection, by having the ability to report suspicious emails and to get feedback by automated analyses and human analysts after their reports. The project aims to find the best ways of involving users in a way that at the same time trains them to recognize phishing emails better. Moreover, we will analyze if user reports can be a useful first line of defense against 0-day phishing, by using reports to train machine learning classifiers that generate rules, instead of relying on burdensome manual creation by human experts.

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)

Industry partner

NEC

G-SINC: Global Synchronization Infrastructure for Network Clocks

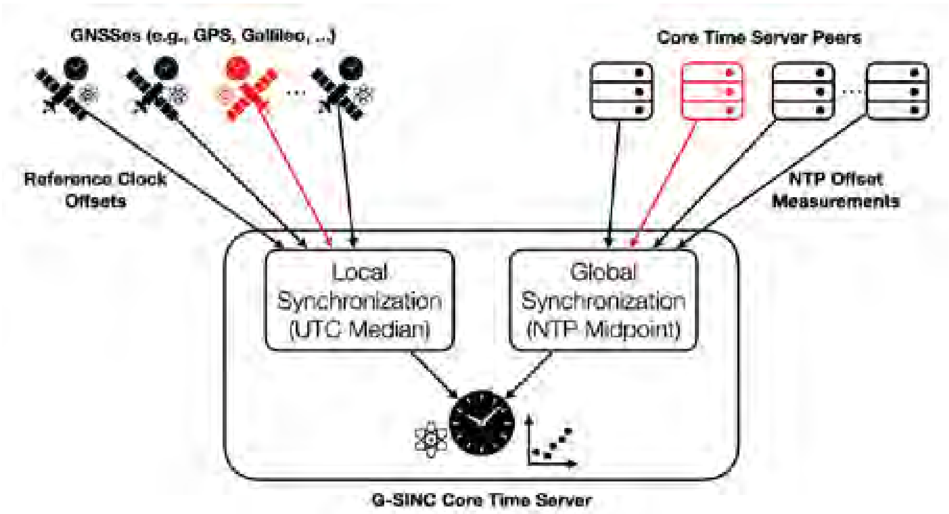
GNSSes (e.g., GPS, Galileo, ...)
Core Time Server Peers
Reference Clock Offsets
NTP Offset Measurements
Local Synchronization (UTC Median)
Global Synchronization (NTP Midpoint)
G-SINC Core Time Server

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

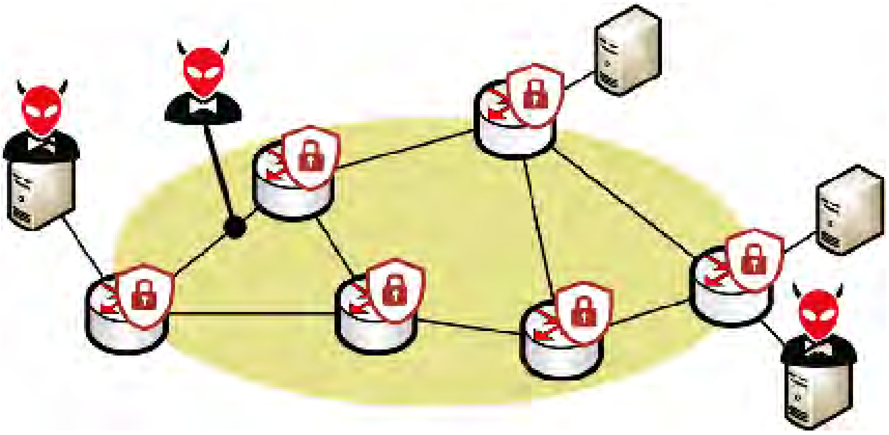

In this project, we argue that the network itself should be able to detect and mitigate attacks instead of relying purely on perimeter-based protection provided by dedicated appliances. To do so, we plan to leverage recent advances in network programmability which enable both the control plane and the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

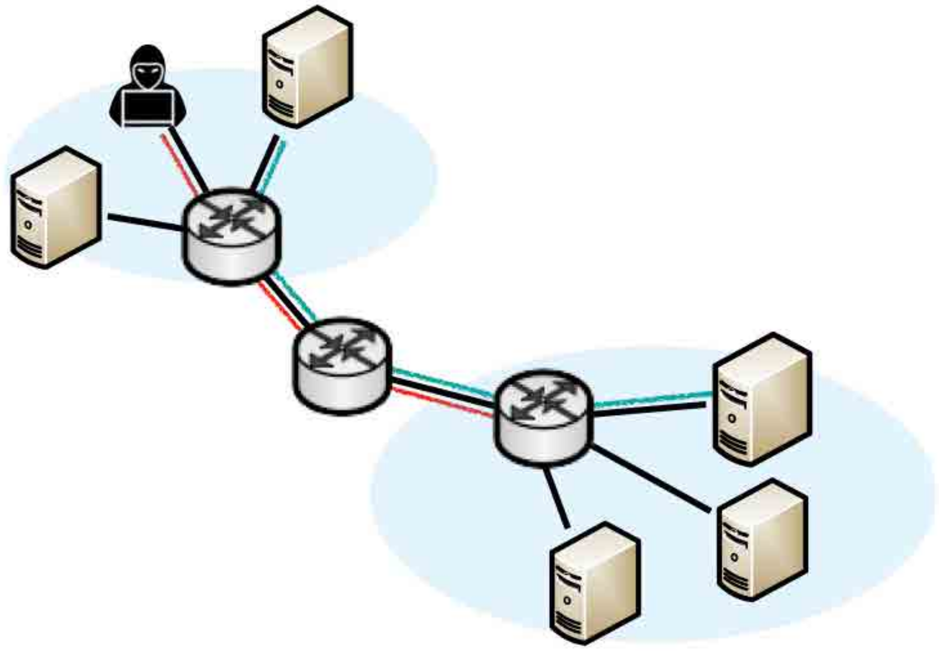

Self-securing Networks

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

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 pro`ject, we examine the SCION protocols in detail and formally verify that they have the desired security properties. We do this both at the modeling and the implementation level. 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 can eventually extract program specifications expressing the implementation's desired I/O behavior. All these steps are formalized in the interactive theorem prover Isabelle/HOL.

We then use this specification to formally verify the Go implementation of the SCION router. In particular, we prove the absence of runtime errors and the implementation's compliance with the specification, i.e., its functional correctness. Additionally, we prove security-related properties of the implementation like secure information flow.

Since the verification effort on the protocol level uses a different formalism than the verification of the code level, a sound link has to be created between them. We realize this link by a refinement step that translates the abstract model into a specification of its IO-behavior. The soundness of this translation is proved in an interactive theorem prover.

Our goal is to gain a better understanding of the underlying properties of the SCION protocol 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.

In 2021, we completed the verification of the latest version of the SCION data plane (cf. Research Highlights) and published a paper on our protocol verification framework. We also made progress towards enriching our models with more details of the SCION protocol, to prepare for a linking to the code using the methodology that we have developed.

On the code verification side, we published a paper about "Gobra", our Go verifier capable of handling Go›s advanced language features, namely channel-based concurrency and interfaces. Using Gobra, we have proved that substantial parts of the data plane implementation of SCION's border router satisfy memory safety and crash safety.

Further information

T. Klenze, Ch. Sprenger, D. Basin
Formal Verification of Secure Forwarding Protocols, CSF 2021

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)

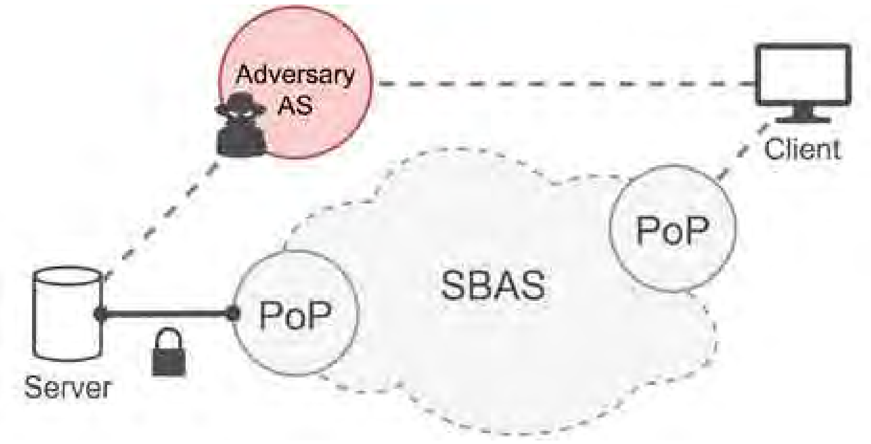

SBAS: Bridging the Gap to SCION

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

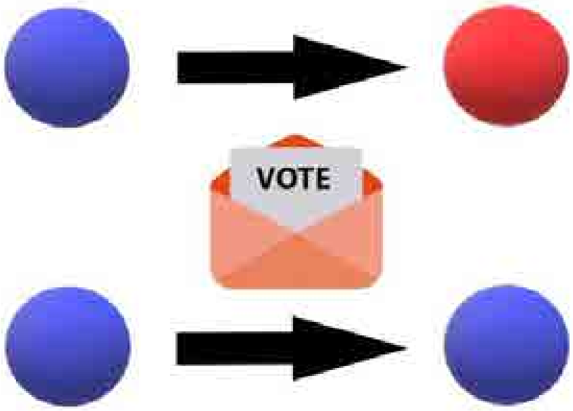

Secure Governance Schemes for Blockchains

Systems based on blockchain technology are promising, as they can be decentralized and rendered robust against attacks. A blockchain is a (distributed) ledger, in which all transactions are recorded sequentially. Because such systems build on distributed consensus –i.e. they require a large number of participants to agree on whether a new transaction should be valid, which they do by holding a copy of the ledger– they function without the need to build trust among its participants or to rely on a trusted third-party.

A blockchain is also governed by a number of parameters such as the block size, the upgrade specifications or the reward systems for validators. Following the decentralization principle underlying distributed consensus, it should be possible for all blockchain stakeholders to have a say on changing these parameters, i.e. to decide about the governance of the blockchain. Yet, most blockchains exclude the majority of stakeholders (participants) from governance.

We develop a new secure voting scheme for the governance of a proof-of-stake blockchain, which we generically call Blockchain Assessment Voting (BAV). Although our focus is on governance, we also expect to reap insights that can be helpful to achieve distributed consensus more efficiently.

BAV schemes consist of two voting rounds. When a proposal is made, some randomly selected stakeholders obtain voting rights in relation to their stakes, but their anonymity is preserved. These stakeholders (simultaneously) vote on the proposal on the table, which is pitted against the status quo. The result of this first voting round is observed by all stakeholders, no matter whether they participated in the first round or not. Upon publication of the first-round results, the proposal may be retracted or amended by its authors, in which case BAV stops.

Alternatively, BAV may also stop if some pre-determined vote threshold has not been reached in this first voting round. If it has not stopped, the scheme continues with the second stage, in which the proposal is put to vote among the remaining stakeholders. The final decision between the proposal and the status quo is taken by adding the votes of the two voting rounds.

In the context of blockchains, one might also want to allow stakeholders to delegate their voting rights to other participants, but this could open possibilities for manipulation. We will use mathematical tools and a blockchain architecture based on proof-of-stake to assess whether and how BAV schemes could be used to improve outcomes in blockchain decisions and how they could prevent manipulation of outcomes by a small coalition.

Researchers
Prof. Dr. Hans Gersbach (ETH)
Dr. Akaki Mamageishvili (ETH)
Manvir Schneider (ETH)

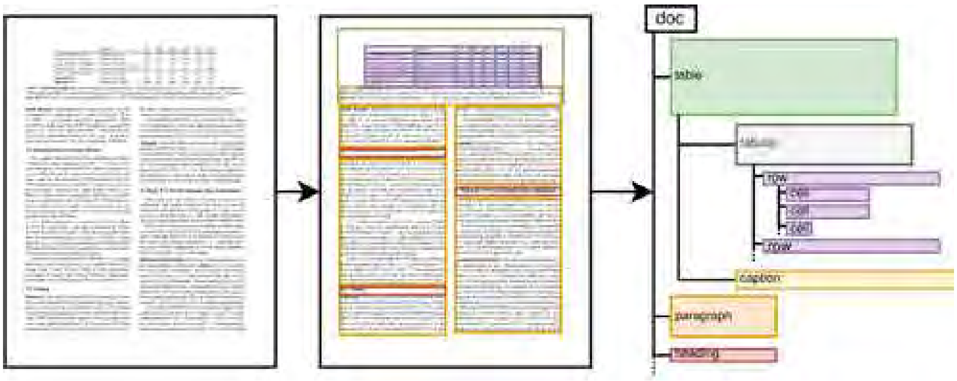

Automatic Visual Document Parsing

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)

Industry partner

ZURICH

Research Projects

Privacy Preserving Machine Learning for Cyber Insurance

A typical cyber insurance product provides coverage against monetary loss caused by cyber attacks or IT failures. Many companies have an increasing need for such protection, and thus this insurance line of business is growing rapidly. Compared to many other traditional areas of insurance, insurers still face challenges with respect to the cyber peril. The level of understanding of cyber risk, i.e. how to thoroughly assess risk, describe the risk, model the risk, is not on the same level as for a number of other risks. One major obstacle insurers are confronted with is the lack of trustworthy and structured data to describe cyber exposures and cyber losses.

Insurers address this problem today by collecting data from the insureds using detailed questionnaires that the customer needs to fill in. Such questionnaires typically include questions regarding security management and security practices of the company, for instance around the software patching process, remote access, backup and recovery practices.

However, many customers are unwilling to reveal full details of their IT systems and security management. Customers are likely to be concerned that honest answers that indicate poor IT security practices could be used to discriminate against them, either at the time of cyber insurance pricing or possible claim handling.

In this project, we explore recent advances in privacy preserving learning methods. In particular, we focus on differentially private gradient boosted decision trees. Differentially private learning methods allow us to learn information about a dataset while withholding information about any specific instance from the dataset. In other words, the influence of every single instance on the learned model is deniable, hence preserving the instance's privacy.

Additionally, we would like to leverage secure enclave environments such as Intel SGX, which would allow participants to verify the correctness of the learning method's source code prior to sharing their own data, and ensure that no single participant has direct access to the whole dataset. Through additional enclave hardening, the learning method would then run completely isolated in this secure enclave, and only release curated statistical information.

Researchers

Dr. Kari Kostiainen (ETH)
Prof. Dr. Esfandiar Mohammadi (University of Lübeck)

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.

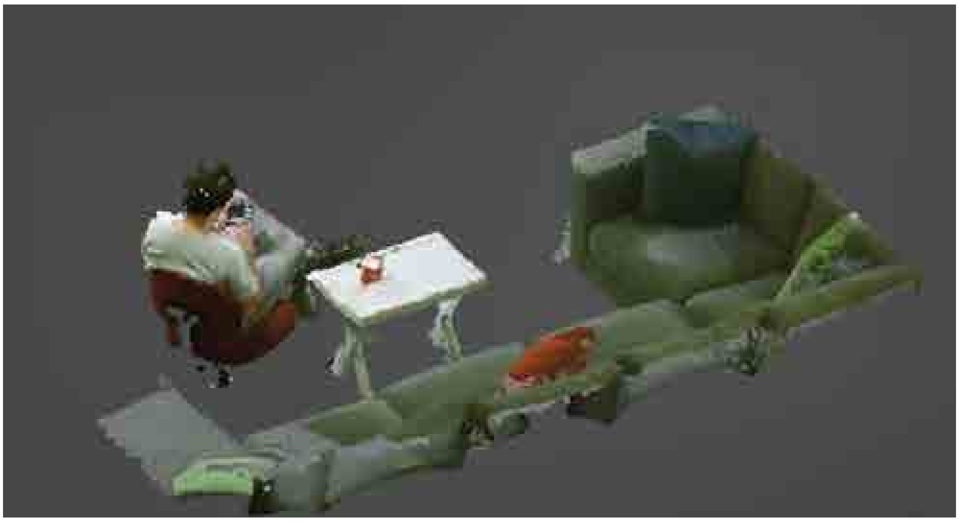

Approach

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

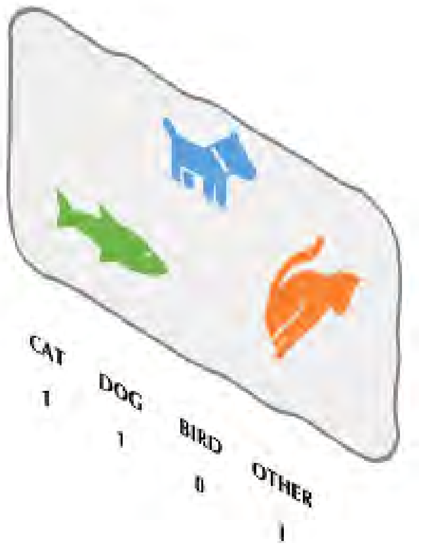

This research project aims at shedding a new light on multi-label classification and consists of two main goals: (i) developing a comprehensive, up-to-date benchmark on multi-label classification for two data modalities, and (ii) improving a multi-label classification system for an email forwarding task 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)

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

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 72 43
fax +41 (0)44 632 11 72

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