ETH zürich

Zurich Information Security and Privacy Center (ZISC)

Annual Review 2021

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

Welcome

During the year 2021 the world began to gradually recover from the Covid-19 pandemic that had disrupted many aspects of our lives. For the ZISC researchers this meant both continued remote working but also partial return to in-person meetings and office life. Under such varied circumstances, the ZISC center continued to deliver excellent results on both of its main mandates: applied research projects with the industry partners and long-term basic research.

Regarding applied research projects that are defined and customized together with our partner companies, we worked on multiple topics during 2021. To mention two examples, the ZISC researchers conducted research on self-securing networks together with armasuisse and studied the effectiveness of phishing emails and training in large organization together with the Swiss Post. More about these applied research projects, and other similar collaborations, will be explained in the following pages of this report.

ZISC researchers worked also on research projects that address fundamental challenges in information security and privacy. To mention one example, ZISC researchers lead by Prof. Shweta Shinde discovered new attack techniques against hardware-assisted trusted execution environments. As another example, Alexander Viand and Anwar Hithnawi developed new techniques for the adoption of fully homomorphic encryption, a fundamental security primitive that has various applications including more secure cloud computing. You can read more about such research projects and highlights later in this report.

The ZISC center also made contributions to projects that have societal importance beyond academia and industry. For example, The ZISC faculty members Prof. David Basin, Prof. Srdjan Capkun, and Prof. Adrian Perrig have collaborated with the International Committee of Red Cross (ICRC) in a project that is, informally said, defining "red cross for cyberspace". The results of this research may be anchored in International Humanitarian Law that could change (for the better) how cyberwars are fought. As another example, ZISC researchers lead by Prof. Kenneth Paterson analyzed the popular Telegram messaging application and identified critical vulnerabilities that have now been fixed. Such research helps millions of users of encrypted messaging services worldwide.

During 2021, the ZISC center also continued its long-standing ZISC lunch seminar tradition with a series with bi-weekly research talks. At the moment, most of the the seminar talks take place online, which makes it easy for our industry partners to join. The online format also allows us to host excellent speakers located around the world. Next year, we hope to complement online seminar talks with in-person presentations that can be accompanied by a social lunch gathering.

The ZISC center wishes all its partners and collaborators a relaxing holiday season and we are looking forward to working with you again in 2022!

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 future Internet architectures.

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

Research Highlights 2021

Zeph: Cryptographic Enforcement of End-to-End Data Privacy

As increasingly more sensitive data is being collected to gain valuable insights, the need to natively integrate privacy controls in data analytics frameworks is growing in importance. Today, the most integral parts of existing data protection systems are security controls such as data access control and data encryption which protect data by guarding it and limiting access to authorized parties. In this model, authorized parties still have unfettered access to data that accounts for much of the data misuse today.

Recently, end-to-end encrypted systems have emerged in which data is encrypted at the source such that even cloud storage or service providers never see data in the clear. However, authorized utility providers from «one end» still get access to the underlying data. Therefore, end-to-end encryption alone is not sufficient to meet the needs of adequate privacy control.

We ultimately need to ensure that users› privacy preferences are respected even by authorized entities. Existing measures generally do not compose well with end-to-end encrypted systems. Privacy solutions that control the extent of what can be inferred (i.e., data minimization and purpose limitation) from data and protect individuals› privacy (i.e., differential privacy) are crucial if we continue to extract utility from data safely. In this project, we work towards a data privacy design that marries end-to-end encrypted systems and privacy solutions. **We want to simultaneously ensure the confidentiality of data from unauthorized parties and provide strong privacy guarantees for data accessed by authorized parties.**

Despite recent advancements in privacy technologies, privacy frameworks remain shaped by regulatory requirements that predominantly focus on the notion of notice and consent. We identified three shortcomings in the current status quo. (i) In the current model, privacy controls are implemented and enforced by data curators who have full access to data in the clear. There are no assurances that data processing complies with the stated privacy policies. (ii) Though privacy regulations mandate services to grant users more control over their data, the materialization of this has been disappointing in practice. Users cannot exert their data privacy preferences except to give blanket consent if they choose to use the service. (iii) Today›s privacy solutions are mostly ad hoc efforts rather than an integral part of the data processing ecosystem. We need a cohesive end-to-end approach to data privacy that follows data from source to downstream.

To address these shortcomings for streaming data, we introduced a new design called Zeph that provides the means to extract value from encrypted streaming data safely while ensuring data

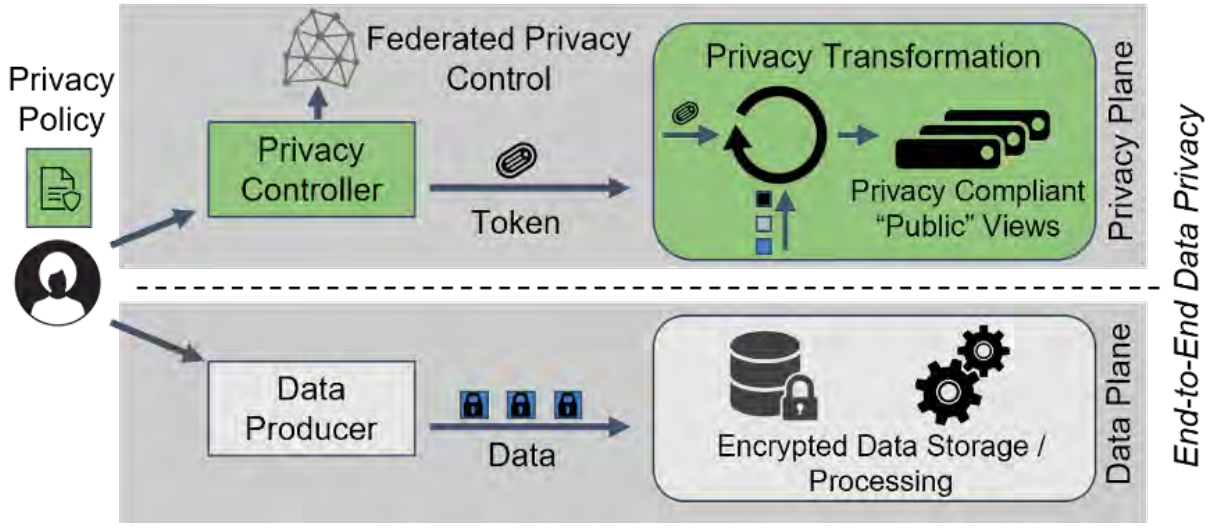

confidentiality and privacy by serving only privacy-compliant views of the data. Zeph augments an existing encrypted data processing system with a privacy plane that allows users to authorize services to access privacy-compliant data securely (See Figure). Zeph builds on two key ideas: (i) a user-centric privacy model that enables users to express their privacy preferences. In Zeph, a user can authorize services to access raw data or privacy-compliant data securely. This aligns with data sharing practices claimed in privacy policies today: e.g., «we share or disclose your personal data with your consent» or «we only provide aggregated statistics and insights». In addition to this commonly referenced aggregation policy, Zeph supports more advanced privacy-compliant data transformations. For example, transformations that restrict what can be inferred from the data (e.g., generalization techniques) or ensure differential privacy. (ii) Zeph cryptographically enforces privacy compliance and executes privacy transformations on-the-fly over encrypted data, ensuring that the generated transformed views conform to users› privacy policies.

To enable privacy-compliant data transformations on encrypted data, we presented: (i) a new approach for encryption that decouples data encryption from privacy transformations. Data producers remain oblivious to the transformations and do not need to encrypt data towards a fixed privacy policy. (ii) The separation between data and privacy plane requires fundamentally different designs, as traditional encrypted solutions are too heavily interwoven cryptographically to allow this split. Therefore, we introduced the concept of cryptographic transformation tokens to realize flexible data transformations. These tokens are, in essence, the necessary cryptographic keying material that enables the respective transformation on encrypted data. Our system creates these tokens via a hybrid construction of secure multi-party computation (MPC) and partially homomorphic encryption schemes. We designed our system so that privacy tokens (which can be combined with encrypted data to release transformed data as per privacy policies) can be generated independently of the data producers. We leverage Homomorphic Secret Sharing (HSS) to achieve this logical separation while still allowing us to homomorphically compute on keying material to construct privacy tokens. Outputs of privacy transformations over encrypted data at the server-side are then released by combining the output results with the corresponding transformation tokens.

We have built a prototype of Zeph that is interfaced with Apache Kafka. Our evaluation results show that Zeph can serve real-time privately transformed streams in different applications with a 2x to 5x latency overhead compared to plaintext.

Further information

Zeph: Cryptographic Enforcement of End-to-End Data Privacy
Lukas Burkhalter*, Nicolas Küchler*, Alexander Viand, Hossein Shafagh, Anwar Hithnawi
USENIX OSDI 2021
https://www.usenix.org/conference/osdi21/presentation/burkhalter.
A prototype of Zeph is available at https://github.com/pps-lab/zeph-artifact.

Researchers

Lukas Burkhalter, Nicolas Küchler, Alexander Viand, Anwar Hithnawi

Research Highlights 2021

LTrack: Stealthy Tracking of Mobile Phones in LTE

People consider their current location as one of the most private information. Defenses against attacks that allow continuous tracking of people are therefore considered one of the most critical. We saw that tracking is possible in a wide range of technologies, e.g., WiFi, but only a few of them are as widespread as cellular networks. LTE is the most widely deployed and used cellular technology. It was designed to not only enable communication but also to protect the security and privacy of users by encrypting communication between a mobile phone and a base station. Unlike the user›s data, LTE low-level layer control messages are transmitted unencrypted. Instead, LTE relied on unique user identifiers (IMSI) being replaced with temporary identifiers (TMSI) to protect users› privacy.

In order for an attacker to be able to exploit LTE technology for large-scale, stealthy tracking of mobile phones, the attacker needs to:

1. determine the location of the mobile phone,
2. obtain a mobile phone›s identifier that links observed locations into a trace, and
3. avoid detection.

Until now, no attack fulfills all of the above at the same time.

Researchers from the System Security Group introduce LTrack, a new tracking attack on LTE that allows an attacker to stealthily extract user devices› locations and permanent identifiers (IMSI).

To remain stealthy, the localization of devices in LTrack is fully passive, relying on a new uplink/downlink sniffer. The sniffer records both the times of arrival of LTE messages and the contents of the low-level control messages, based on which LTrack calculates locations. LTrack is the first to show the feasibility of a passive localization in LTE through implementation on software-defined radio.

Attacks that obtain permanent identifiers IMSIs are called IMSI Catchers. These devices exclusively work by setting up a fake base station that acts like a real base station. To get the mobile phones to connect to the fake base station (which is a requirement of the attack), the attacker needs to transmit continuously at high power and can therefore be detected by law enforcement and operators. LTrack overcomes this challenge by introducing

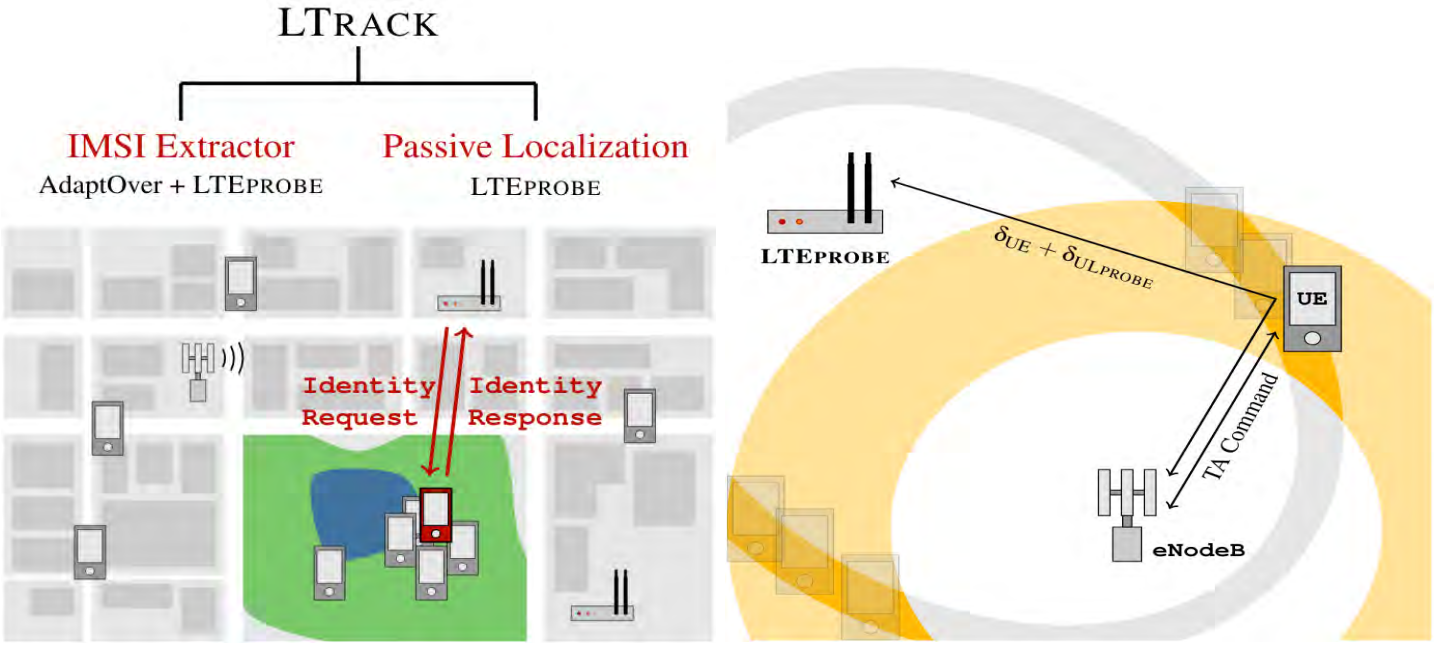

and implementing a new type of IMSI Catcher named IMSI Extractor. It extracts a device›s IMSI and binds it to its current TMSI. Instead of relying on fake base stations like existing IMSI Catchers, IMSI Extractor relies on our uplink/downlink sniffer enhanced with surgical message overshadowing. The attacker precisely synchronizes with a base station and transmits the adversarial message with a slightly higher power. Both the honest and the adversarial messages are received by phone at the same time. The adversarial message is sent with a higher power, resulting in mobile phone decoding it instead of the honest message. Overshadowing with passive sniffing makes our IMSI Extractor the stealthiest IMSI Catcher to date.

Researchers evaluate LTrack through a series of experiments and show that in line-of-sight conditions, the attacker can estimate the location of a phone with less than 6m error in 90\% of the cases.

Furthermore, they successfully tested our IMSI Extractor against a set of 17 modern smartphones connected to our industry-grade LTE testbed. They further demonstrated the uplink/downlink sniffer and IMSI Extractor in a test facility of an operator on a real LTE network.

Further information

LTrack: Stealthy Tracking of Mobile Phones in LTE
Martin Kotuliak, Simon Erni, Patrick Leu, Marc Röschlin, and Srdjan Capkun
in Usenix Security 2022

Researchers

Martin Kotuliak, Simon Erni, Patrick Leu, Dr. Marc Röschlin, Prof. Srdjan Capkun

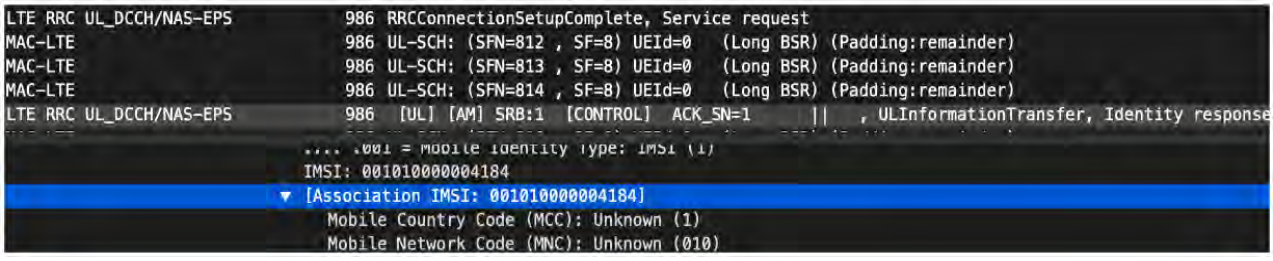

Research Highlights 2021

Making Fully Homomorphic Encryption Accessible

Privacy and security are gaining tremendous importance across all organizations, as public perception of these issues has shifted and expectations, including regulatory demands, have increased. This in turn has translated into organizations adopting stronger security and privacy protections. While industry best practices like in-transit and at-rest encryption provide important protection for user data, the need to decrypt data to compute on it still exposes the data to a wide variety of threats. Secure computation techniques like Fully Homomorphic Encryption (FHE), which allows a third party to perform arbitrary computations on encrypted data, address this gap. Effectively, FHE allows delegating computation without needing to grant access to the underlying data. In recent years, we have seen a leap in performance in FHE driven by a series of breakthroughs and advancements, which has propelled FHE into the realm of practical applications.

Today, performance is no longer the major barrier to adoption. Instead, it is primarily the complexity of developing an efficient FHE application that currently limits deployment. The intricacy of the underlying schemes means that achieving state-of-the-art results requires significant experience, limiting the development of FHE-based applications predominantly to experts. Additionally, FHE imposes a fundamentally different programming paradigm. This arises not only since the security guarantees imply programs must be data-independent, but also because FHE ciphertexts 'deteriorate' during homomorphic operations, which must be carefully managed. Finally, many schemes feature powerful inherent parallelism, where one can pack many messages into a single ciphertext. However, fully exploiting this feature requires significant rethinking and redesign of applications and algorithms as the characteristics of FHE do not match existing parallel computation paradigms. As a result of these challenges, there is a vast gap between state-of-the-art performance results and what non-experts can achieve themselves.

The next leap towards broader adoption of FHE necessitates designing and building a development ecosystem for FHE that facilitates FHE application development. This requires designing and developing tools and compilers that provide the right abstractions and automatic optimizations to tame the current complexity of FHE development. Therefore, the primary goal of our work is to build this ecosystem by developing techniques, abstractions, and tools that address the subtle complexities of working with FHE. Beyond this, we focus on FHE cost models and code generation techniques for heterogeneous platforms that employ FHE hardware accelerators.

Towards this, we developed HECO, an end-to-end compiler and toolchain for FHE that aims to make developing secure and efficient FHE applications accessible to non-experts. At its core is a program transformation logic that

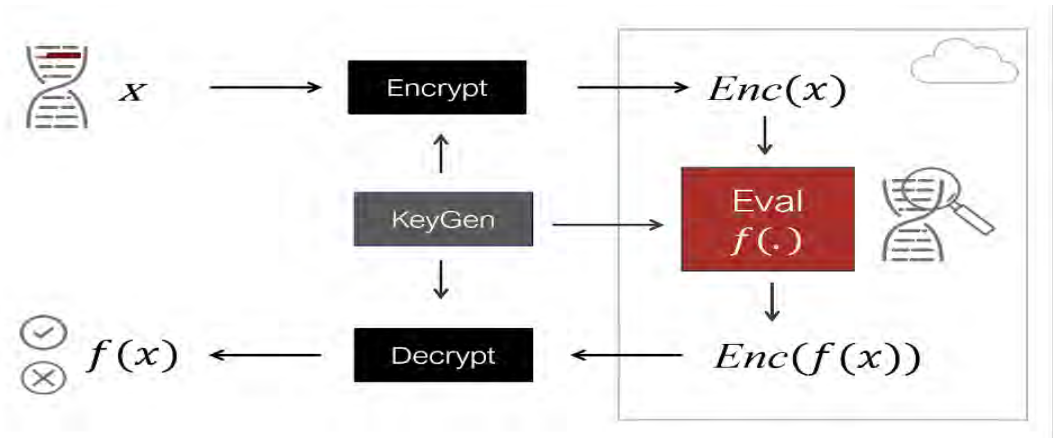

"Fully Homomorphic Encryption (FHE) allows evaluating a function over encrypted data. For example, FHE can be used in Genome Wide Association Studies (GWAS) to compute whether a gene mutation is associated with a specific disease."

maps standard high-level imperative code to the unique programming paradigm of FHE. Our compiler can generate, from unoptimized high-level input, code that matches that written by an expert. This can have dramatic performance benefits, as the performance differences between unoptimized code and code that has been adapted to the FHE paradigm can be significant, with differences by more than one Order of Magnitude being common.

Looking to the future, we have a variety of exciting work lined up to deliver on our vision of an accessible FHE ecosystem. We are proposing rich Intermediate Representations that can capture both high-level program information and low-level details of the different targets, with the aim to achieve an industry standard to unify the ecosystem. In order to broaden the scope of FHE compilers, we have also started working on targeting heterogeneous hardware and recently started collaborating with Intel and Microsoft in their joint effort to design dedicated FHE hardware accelerators.

Further information

Viand et al. Marble: Making Fully Homomorphic Encryption Accessible to All. Proceedings of the 6th Workshop on Encrypted Computing & Applied Homomorphic Cryptography (2018), 49–60.

Viand et al. SoK: Fully Homomorphic Encryption Compilers.
IEEE Symposium on Security and Privacy (SP) (2021), 1166–1182.

A prototype of the FHE compiler is available at https://github.com/MarbleHE/ABC

Researchers

Alexander Viand, Anwar Hithnawi

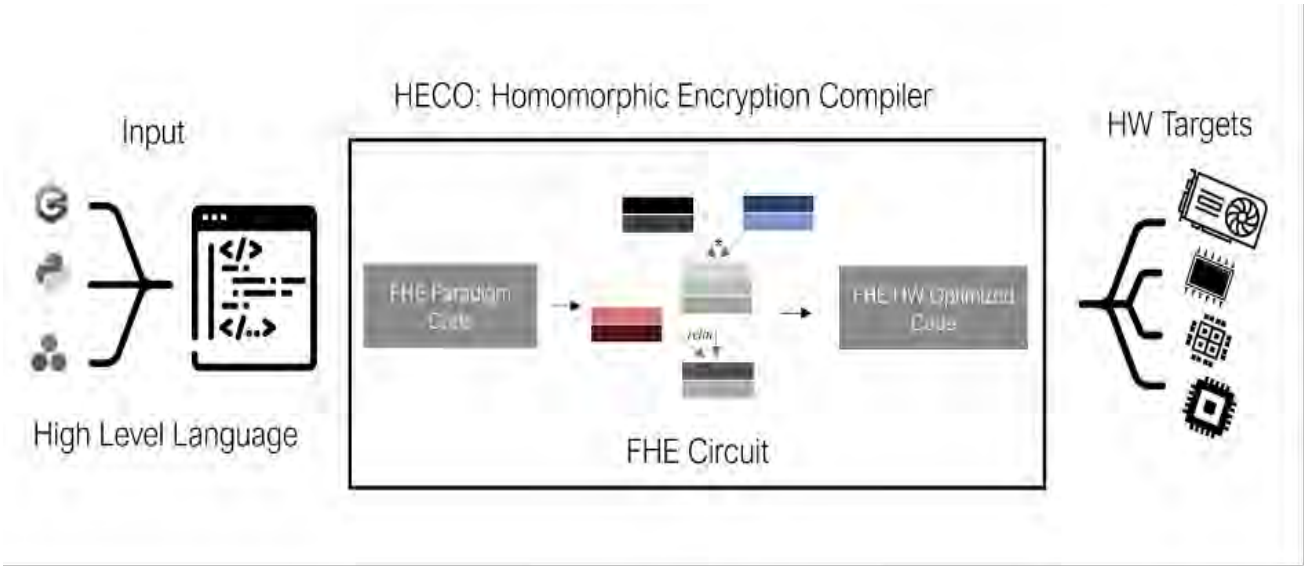

Research Highlights 2021

Colibri: scalable and secure Internet bandwidth reservations

SLOs for the Internet

With the consolidation of computation in data centers, the tension created by co-location of multiple tenants sharing resources has resulted in numerous separation mechanisms to provide resource guarantees. Consequently, service-level objectives (SLOs) have been established at various levels: the CPU provides isolation to prevent information leakage; hypervisors can be configured to guarantee minimal amounts of resources to individual virtual machines, and the local network optimizes communications within the data center. The final frontier is the Internet, raising a fundamental research question: can we provide SLOs in an open Internet—even under the threat of denial-of-service (DoS) attacks—in a scalable manner?

While centralized control facilitates SLOs (e.g., inside a data center), there are numerous interacting entities on the heterogeneous Internet, complicating the task of providing decentralized, scalable SLOs.

Trade-offs of the past

Of the many systems proposed to achieve guarantees for global communication, the two archetypal and most widely deployed architectures are Integrated Services (IntServ) and Differentiated Services (DiffServ). They constitute the two extreme points in the trade-off spectrum between scalability and strength of offered guarantees: strong guarantees overload the network precluding Internet-wide scalability, while in lightweight systems such guarantees are sacrificed. Further, neither DiffServ nor IntServ protects against adversarial action.

Colibri: bandwidth reservations as Internet SLOs

Various ETH researchers from the Network Security group (Prof. Adrian Perrig) have introduced Colibri, a concrete design and implementation of a collaborative lightweight inter-domain bandwidth-reservation infrastructure for the global Internet that overcomes the scalability–quality trade-off and can provide SLOs in the form of worst-case minimum bandwidth guarantees.

Colibri is made possible through the confluence of several recently developed technologies: path-aware networking with SCION, providing path choice and stability; NTube, a fair and scalable resource-allocation system; the DRKey global symmetric-key distribution system enabling efficient per-packet authentication; an efficient replay-suppression system; and LOFT, an overuse-flow-detection system.

The core property of Colibri is to provide minimum bandwidth guarantees between any pair of ASes on a given path, irrespective of distributed-denial-of-service (DDoS) attacks or other allocations. In the common (non-attack) case, allocations will be much higher, but worst-case guarantees enable Internet-scale SLOs.

In summary, the researchers developed a concrete, scalable system design to solve the open problem of Internet-scale SLOs, bridging the gap between the data center and the end-user, and completing the final step towards end-to-end-protected services. They also implemented and integrated Colibri to SCIONLab. This is an important step to providing strong Internet-wide SLOs, thus increasing the efficiency, resilience, and profitability of inter-domain communications.

Further information

Colibri: A Cooperative Lightweight Inter-domain Bandwidth-Reservation Infrastructure
In Proceedings of the ACM Conference on emerging Networking EXperiments and Technologies (CoNEXT) 2021.

Researchers

Giacomo Giuliari, Dominik Roos, Marc Wyss, Juan A. Garcia-Pardo, Dr. Markus Legner, Prof. Adrian Perrig

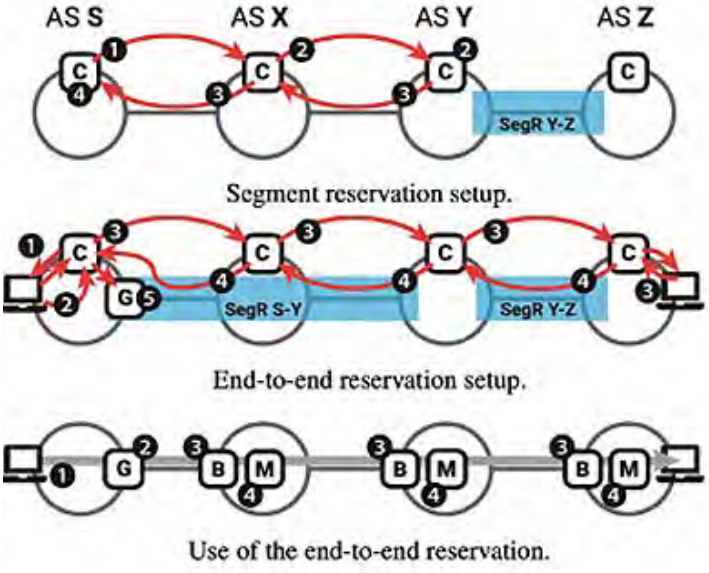

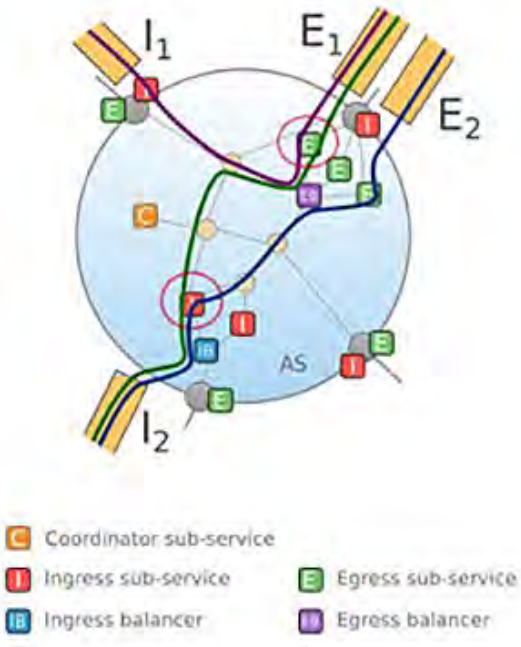

Research Highlights 2021

Formal Verification of the Secure Network Architecture SCION

The current Internet suffers from a deeply-rooted lack of scalability, reliability and security. Despite decades of research, no straightforward solution is in sight, and calls for a clean-slate redesign of the Internet are gaining traction. The SCION architecture was developed by Adrian Perrig, ETH Professor for Network Security, as a response to this challenge. SCION relies on highly efficient cryptographic packet validation for scalable, reliable and secure forwarding.

The SCION packet forwarding protocol has been carefully designed and tested. Testing alone, however, cannot exhaustively cover all possible protocol executions and thus cannot rule out all attacks.

By contrast, mathematical proofs of security hold for all possible protocol executions, assuming that the verified model accurately abstracts the real system and environment. Such proofs can be constructed using formal verification techniques.

This is the domain of ETH researchers headed by David Basin, Professor for Information Security.

The team has used state-of-the-art formal verification techniques and tools to model and verify the SCION's packet forwarding, as part of a larger effort to verify the SCION architecture. This principled approach models system components, the network environment, and the adversary, and formalizes the desired properties. In the case of SCION, the attacker is a colluding Dolev-Yao attacker that controls part of the network.

The properties are path authorization and detectability: The first asserts that each packet only traverses the network along paths authorized by all on-path entities. This ensures that the routing policies of autonomous systems are respected. The second states that a malicious on-path entity cannot hide its presence onthe path.

The properties are formalized in higher-order logic and the model is developed in an interactive theorem prover using refinement. Refinement breaks the proof down into small steps that can be verified by both humans and machines.

The two properties presented above were proven. This process revealed several security-relevant weaknesses that were resolved quickly, and in close collaboration with SCION protocol designers and the implementation team.

The security proofs also aided in the precise understanding of the security mechanisms of SCION and thus facilitated the development of improved variants of forwarding protocols that further strengthen security and achieve even better performance.

In 2021, we were able to complete the security proof for SCION's latest data plane variant. This protocol version includes a new mechanism to authenticate authorized paths. The mechanism makes use of exclusive-or, which is a seemingly simple primitive, but is in fact challenging to formally reason about. We upgraded our verification framework to support this and other features, which allowed us to formally prove SCION's data plane security.

We are currently working on enriching our model with more detail in preparation for linking our protocol verification efforts with an ongoing project verifying the SCION router software. Using the methodology that we developed (cf. Research Project "Full-Stack Verification"), we will soundly connect our models to the code implementing the SCION router.

Further information

Recent Publication:
Tobias Klenze, Christoph Sprenger, David Basin
Formal Verification of Secure Forwarding Protocols
CSF 2021

Researchers

Tobias Klenze, Christoph Sprenger, David Basin

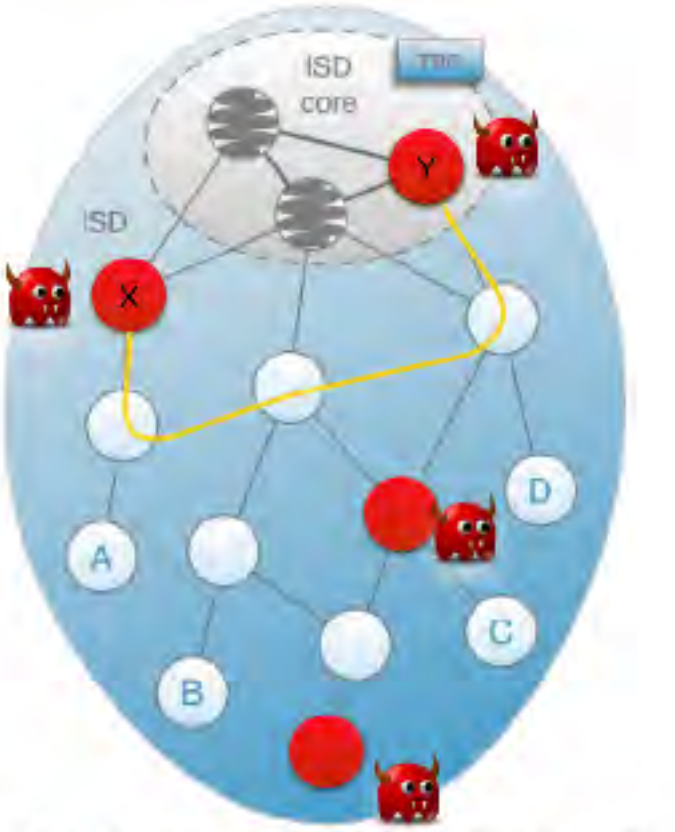

Attacker: controls subset of network (red)
Property: packets traverse authorized path. Rules out ineconomical paths, e.g., valley routing (yellow)

Research Highlights 2021

SmashEx Attack

A few years ago, Intel, the world's leading supplier of PC microprocessors, introduced an improvement that promised greater data security: Software Guard Extensions (SGX). These are hardware-based control mechanisms that ensure that data is secure even if a computer's operating system is malicious or under attack. As operating systems have to perform a huge number of functions, and are highly complex, it makes sense, therefore, to shield applications with sensitive data from the operating system. Software Guard Extensions allow this to be done by means of "enclaves", with certain areas serving to protect the program code of applications that should not be accessed by the operating system.

A newly disclosed vulnerability affecting Intel processors could be abused by an adversary to gain access to sensitive information stored within enclaves and even run arbitrary code on vulnerable systems. The vulnerability was used to stage a confidential data disclosure attack called «SmashEx» that can corrupt private data housed in the enclave and break its integrity.

Introduced with Intel›s Skylake processors, SGX (short for Software Guard eXtensions) allows developers to run selected application modules in a completely isolated secure compartment of memory, called an enclave or a Trusted Execution Environment (TEE), which is designed to be protected from processes running at higher privilege levels like the operating system. SGX ensures that data is secure even if a computer›s operating system has been tampered with or is under attack.

The vulnerability is rated by Intel itself with a CVSS (Common Vulnerability Scoring System) score of 8.2 out of 10. This scoring system indicates the severity of vulnerabilities based on a range of indicators. In this case, one of the reasons for the high score is the fact that the problem affected new hardware and a potentially large number of corporate and private customers – Intel processors with the relevant Software Guard Extensions are very widespread. Thus, among other things, Google products were also affected. Intel SGX enclaves are also often used when IT infrastructure is shared between different parties, or when sensitive data is involved – in the banking or healthcare sectors, for example. The fact that the vulnerability affected a technology designed specifically for sensitive data gives us pause, but it's not a reason to panic. The problem has been solved for the time being using software patches, but hardware adaptation for future processor generations are recommended to make them more secure in the long term. Intel has since released software updates to mitigate this vulnerability with SGX SDK versions 2.13 and 2.14 for Windows and Linux respectively.

Website
https://jasonyu1996.github.io/SmashEx/

Further information
Publication:
Cui J, Zhijingcheng Yu J, Shinde S, Saxena P, Cai Z,
«SmashEx: Smashing SGX Enclaves Using Exceptions»,
Proceedings of the ACM Conference on Computer and Communications Security (CCS). November 2021.

Researchers
Prof. Shweta Shinde

Research Highlights 2021

Security Analysis of Telegram

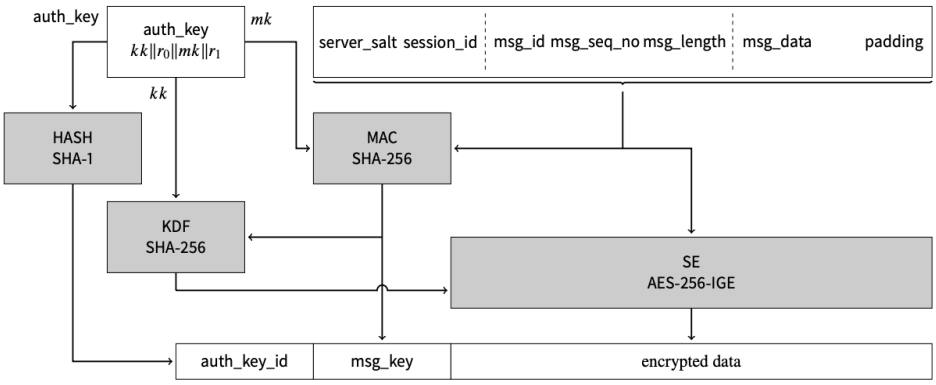

Telegram is a popular messaging platform that, as of January 2021, reportedly has 500M monthly users. We performed a detailed security analysis of the encryption offered by Telegram. As a result of our analysis, we found several cryptographic weaknesses in the protocol, from technically trivial and easy to exploit to more advanced and of theoretical interest.

For most users, the immediate risk was low, but these vulnerabilities highlight that Telegram fell short of the cryptographic guarantees enjoyed by other widely deployed cryptographic protocols such as TLS. We made several suggestions to the Telegram developers that enable providing formal assurances that rule out a large class of cryptographic attacks, similarly to other, more established, cryptographic protocols.

By default, Telegram uses its bespoke MTProto protocol to secure communication between clients and its servers as a replacement for the industry-standard Transport Layer Security (TLS) protocol. While Telegram is often referred to as an "encrypted messenger", this level of protection is the only protection offered by default: MTProto-based end-to-end encryption, which would protect communication from Telegram employees or anyone breaking into Telegram's servers, is only optional and not available for group chats. So we focused our efforts on analysing whether Telegram's MTProto offers comparable privacy to surfing the web with HTTPS.

We disclosed four vulnerabilities to the Telegram development team on April 2021 and agreed with them on a disclosure on July 2021. To briefly mention one of the vulnerabilities, an attacker on the network can reorder messages coming from a client to the server. This allows, for example, to alter the order of "pizza" and "crime" in the sequence of messages: "I say yes to", "all the pizzas", "I say no to", "all the crimes". This attack is trivial to carry out. Telegram confirmed the behaviour we observed and addressed this issue in an update of their software.

The central result of our investigation is that Telegram's MTProto can provide a confidential and integrity-protected channel when the changes we suggested are adopted by the Telegram developers. While we prove security of MTProto at a protocol level, communication via Telegram must trust the Telegram servers by default, i.e. end-to-end encryption is optional and not available for group chats. We thus recommend to avoid referring to Telegram as an "encrypted messenger".

Website

https://mtpsym.github.io/

Further information

Publication:
"Four Attacks and a Proof for Telegram", IEEE Symposium on Security & Privacy, to appear, 2022.

Researchers

Prof. Kenny Paterson, Dr. Igors Stepanovs (ETH), Prof. Martin Albrecht, Lenka Mareková, (Royal Holloway, University of London).

Research Highlights 2021

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, ADEM›s 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.

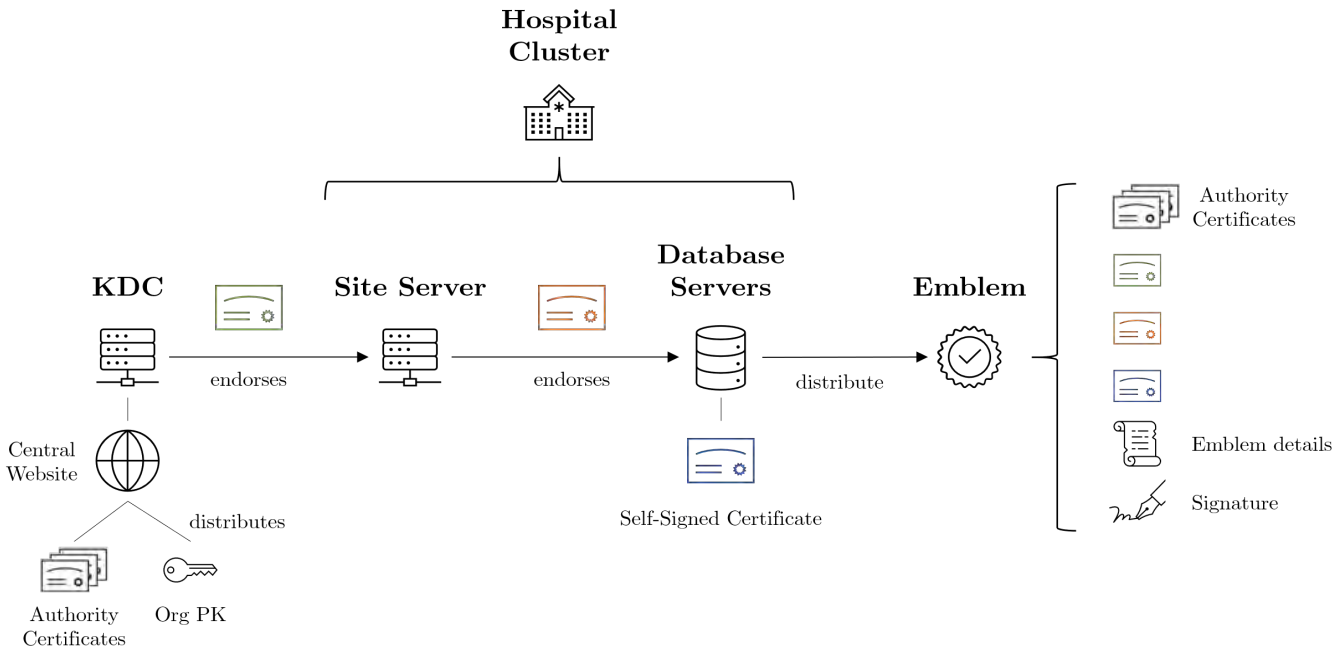

Education 2021

Security in School education

Center of Computer Science Education (ABZ) of ETH Zurich was founded with the goal to introduce computer science as a subject into school education. The main activities of ABZ include developing textbooks and online platforms for teaching computer science on all levels of schools and testing them in school, training teachers more than 4000 in the last three years, popularization of computer science in the whole society, and supporting pupils for different CS competitions like Olympiad in Informatics, Informatics Beaver, ACM Programming Contests.

The main achievement is establishing "informatics" as a mandatory subject in Lehrplan 21 for obligatory schools as a result of long-term projects in more than 500 schools and more than 400 appearances in the media.

The contribution to teaching "Security" include: Textbook "Einführung in die Kryptologie", several school projects on this topic and chapters in the new textbook series «Einfach Informatik» containing 15 titles for children of all age groups several courses for teachers for teaching cryptology.

The ABZ also supports the further education of gifted pupils in computer science.

The workshops are held both locally at interested schools throughout Switzerland or directly at ETH Zurich. The content is broad, ranging from programming with LOGO or Python to topics from Computer Science Unplugged.

The ZISC center is proud to support this project!

Grundlagen der Informatik für Schweizer Maturitätsschulen

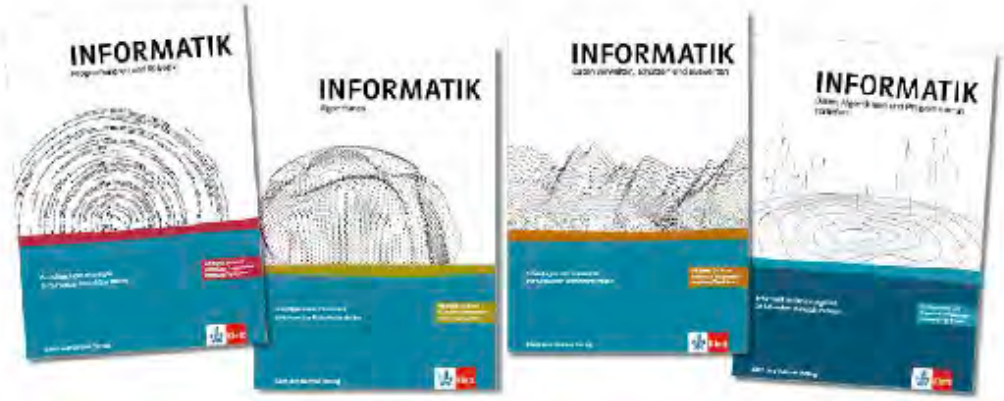

Main Research Areas

Sovereign Smartphone

Prof. Shweta Shinde

Most 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. Users, developers, and governments are left with little to no choice, they need to entrust their security and privacy to OS vendors. Given the wide use of Android and iOS, immediately leaving these ecosystems is not practical, except in niche application areas.
In our work, we draw attention to the magnitude of this problem and why it is an undesirable situation. As an alternative, we advocate the development of a new smartphone architecture that securely transfers the control back to the users while maintaining compatibility with the rich existing smartphone ecosystems. We propose and analyze one such design based on advances in trusted execution environments for ARM and RISC-V.

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 subplanes, 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.

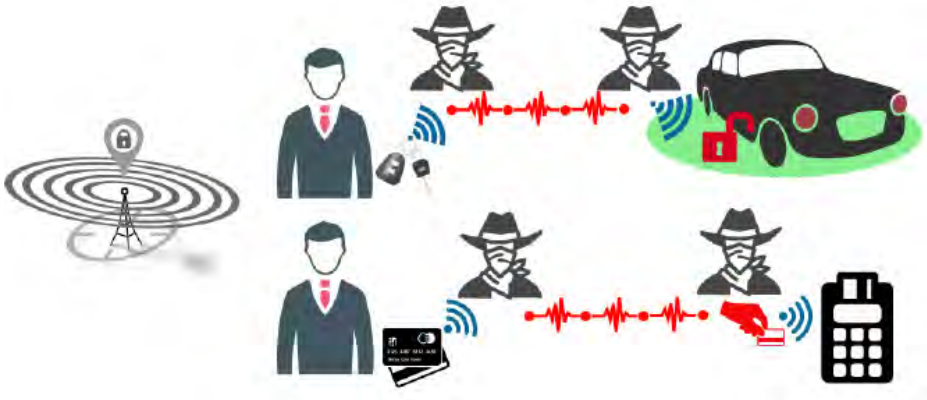

Main Research Areas

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.

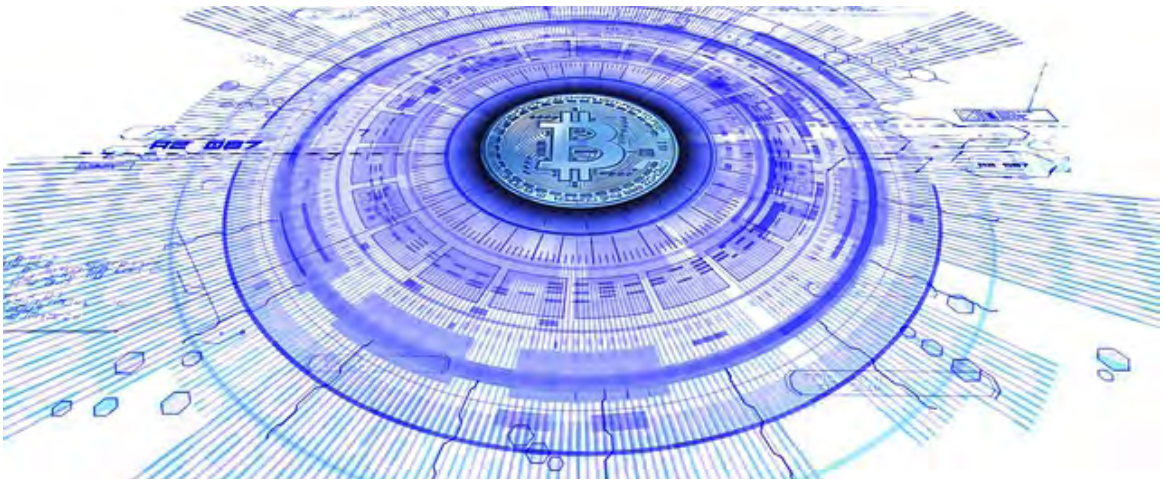

Research Projects

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

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.

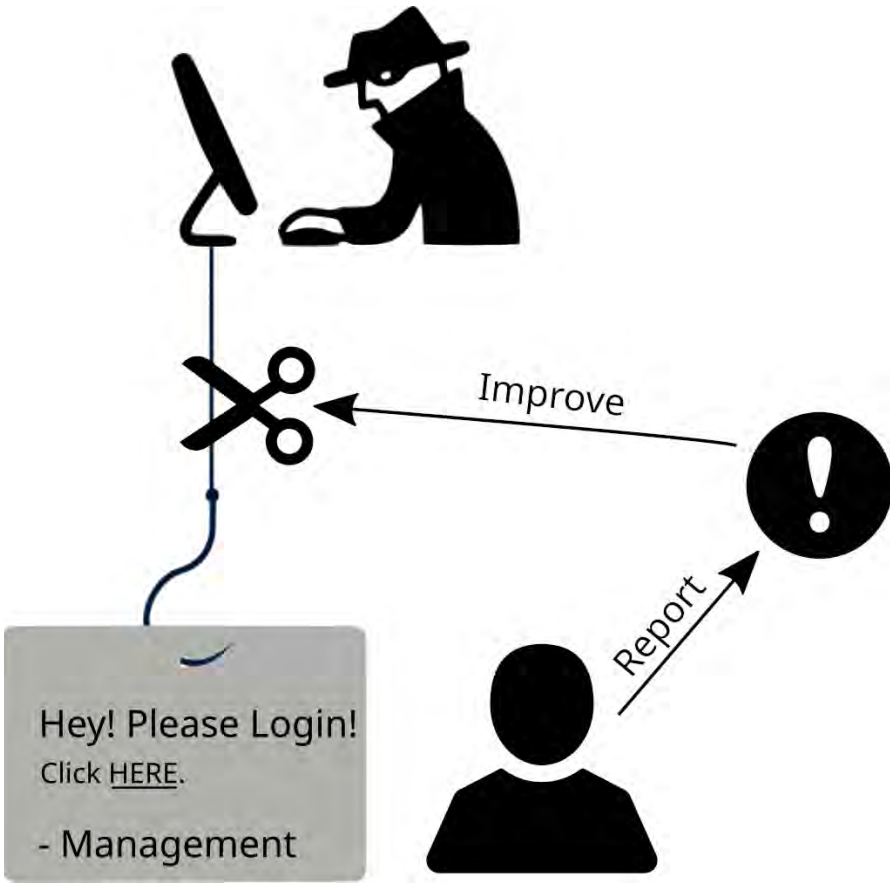

In this project, in partnership with the Swiss Post, we aim to conduct a large-scale study on phishing prevention, detection, and education. Users will be involved in phishing detection, by having the ability to report suspicious emails and to get feedback by automated analyses and human analysts after 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 at addressing various issues in cloud and blockchain security in an aim to improve their security and scalability.
In the area of blockchain technology, our project focuses on the security and privacy of different blockchain technologies and on the development of new protocols and systems to enhance functionality.

We propose 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 design and implement a system called BITE where enclaves on full nodes serve privacy-preserving requests from lightweight clients. As we will show, naive serving of client requests from within SGX enclaves still leaks user information. BITE therefore integrates several privacy preservation measures that address external leakage as well as SGX side-channels. We show that the resulting solution provides strong privacy protection and at the same time improves the performance of current lightweight clients.

Further, we 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 multisignature transactions signed by a quorum of them. Service providers in Bitcontracts are stateless and do not need to communicate with the Blockchain's peer-to-peer network.

Lastly, we investigated problems with mining centralization and analyzed approaches that try to solve these issues with decentralization of mining pools. We 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 propose a solution using trusted execution environments.

In the area of cloud security, our project investigated secure data deduplication and novel access control paradigms in the cloud. Deduplication allows storage reduction and makes cloud storage financially attractive to customers, but also generates numerous privacy and security challenges. Moreover, although the cloud encourages data sharing, existing access control paradigms do not fit all the new requirements arising from shared storage between partially-trusted partners. Therefore, in this project, we devised novel access control paradigms that allow data sharing according to users' needs.

Researchers
Karl Wüst (ETH)
Kari Kostiainen (ETH)

Industry partner

NEC

Towards Provably Secure Internet Communication

Nowadays, the wide-spread access to the Internet enables quick communication, unrestrained by physical location. However, this comes at a cost of new security risks, since now private messages become available to adversarial entities, located anywhere around the world. Hence, cryptographic protocols that add security to the communication become essential.

Since different situations have different functional and security requirements, the number of secure-communication protocols with different security-functionality-efficiency trade-offs is rapidly growing. For example, we have various session-establishment protocols (such as TCP-based TLS, or Google's QUIC based on faster but less reliable UDP), various secure-messaging protocols (such as Signal's double ratchet, or the group messaging protocol currently being standardized by the MLS working group), and many more. The large number of use cases, trade-offs and accompanying protocols (often designed in an ad-hoc fashion and without clearly specified security guarantees) motivates the goal of this project, which is to explore from the cryptographic perspective the space of secure-communication protocols.

More specifically, for various functionality requirements, we specify different security guarantees, where usually stronger guarantees require less efficient protocols. This is done with the help of cryptographic modeling tools, such as the (standard) game-based security analysis and the constructive cryptography framework (which, in particular, allows to express the strong guarantee of composability, i.e. a protocol is secure even if arbitrary other protocols are executed simultaneously).

This allows to, first, express the exact guarantees of existing protocols (and either verify that they meet their intuitive goals, or discover a gap between the intuition and reality) and, second, provide new protocols offering previously unexplored trade-offs.

Further information

A Unified and Composable Take on Ratcheting

Daniel Jost and Ueli Maurer and Marta Mularczyk

Theory of Cryptography Conference, TCC 2019

Continuous Group Key Agreement with Active Security

Joël Alwen, Sandro Coretti, Daniel Jost, and Marta Mularczyk

Theory of Cryptography Conference, TCC 2020

On the Insider Security of MLS

Joël Alwen and Daniel Jost and Marta Mularczyk

Cryptology ePrint Archive, 2020

Researchers

Marta Mularczyk (ETH)

Ueli Maurer (ETH)

Research Projects

Topology-Hiding Computation

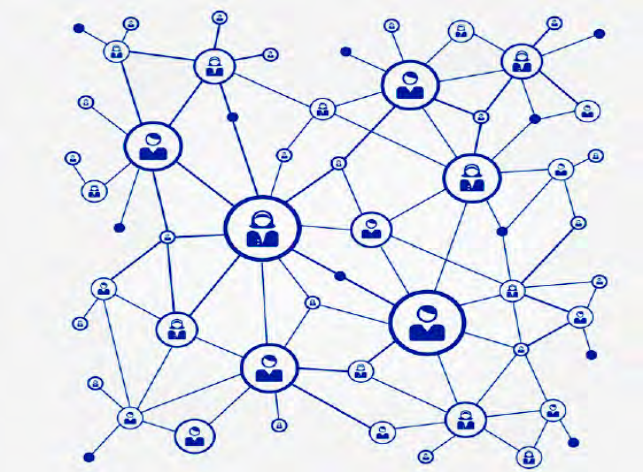

Secure communication over an insecure network is one of the fundamental goals of cryptography. The security goal can be to hide different aspects of the communication, ranging from the content (secrecy), the participants' identity (anonymity), the existence of communication (steganography), to hiding the topology of the underlying network in case it is not complete.

Incomplete networks arise in many contexts, such as social networks, the Internet of Things (IoT) or ad-hoc vehicular networks. Hiding the topology can, for example, be important because the position of a node within the network depends on the node's location. This could in turn leak information about the node's identity or other confidential parameters.

Incomplete networks have been studied in the context of communication security, referred to as secure message transmission, where the goal is to enable communication between any pair of entities, despite an incomplete communication graph. Also, anonymous communication has been studied extensively. Unfortunately, none of these approaches can be used to hide the network topology. In fact, secure message transmission protocols assume (for their execution) that the network graph is public knowledge.

The goal of this project is to design topology-hiding communication protocols, which allow a set of parties connected by an incomplete network with unknown communication graph, where each party only knows its neighbors, to communicate in such a way that the network topology remains hidden even from a powerful adversary who can corrupt parties. These communication protocols can then be used to perform arbitrary tasks, for example secure multi-party computation, in a topology-hiding manner. In the formal analysis, we consider different degrees of network hiding. For example, a network may be completely hidden, or some partial knowledge about it may leak to the adversary. Recent results show that we can hide the topology up to leaking 1 bit of information about it with probability p.

Further information

Topology-Hiding Computation Beyond Semi-Honest Adversaries
Rio Lavigne and Chen-Da Liu-Zhang and Ueli Maurer and Tal Moran and Marta Mularczyk and Daniel Tschudi
Theory of Cryptography Conference, TCC, 2018.

Topology-Hiding Computation for Networks with Unknown Delays
Rio LaVigne, Chen-Da Liu Zhang, Ueli Maurer, Tal Moran, Marta Mularczyk, and Daniel Tschudi
Public-Key Cryptography — PKC 2020, LNCS, Springer, vol. 12111, pp. 215–245, Apr 2020.

Researchers
Martin Hirt (ETH)
Ueli Maurer (ETH Zurich)
Marta Mularczyk (ETH)

Enhanced 5G Security

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

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

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

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

Further information

Jonghoon Kwon, Taeho Lee, Claude Hähni, and Adrian Perrig.
SVLAN: Secure & Scalable Network Virtualization.
In Proceedings of the Symposium on Network and Distributed System Security (NDSS) 2020.

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

Researchers

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

Industry partner

NEC

Research Projects

Improving Network Security Through Programmability

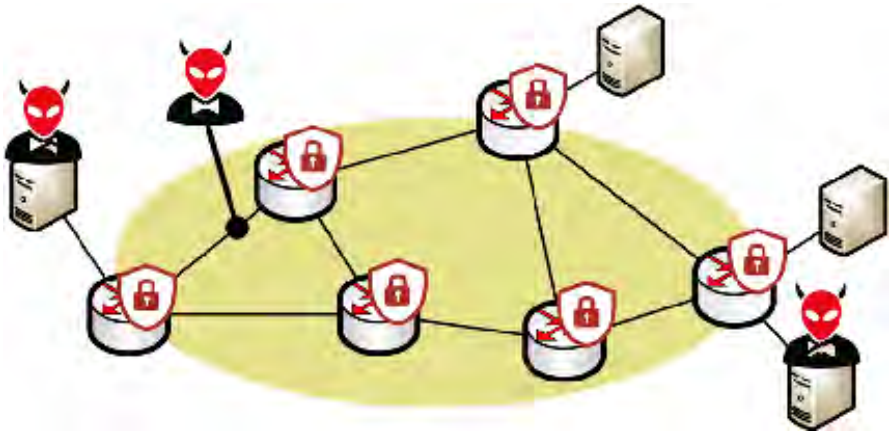

In this project, we argue that the network itself should be able to detect and mitigate attacks instead of relying purely on perimeter-based protection provided by dedicated appliances. To do so, we plan to leverage recent advances in network programmability which enable both the control plane and the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

Roland Meier, Petar Tsankov, Vincent Lenders, Laurent Vanbever, Martin Vechev
NetHide: Secure and Practical Network Topology Obfuscation
USENIX Security 2018. Baltimore, MD, USA (August 2018).
For more details, see: https://nethide.ethz.ch

Roland Meier, Thomas Holterbach, Stephan Keck, Matthias Stähli, Vincent Lenders, Ankit Singla, Laurent Vanbever
(Self) Driving Under the Influence: Intoxicating Adversarial Network Inputs
ACM HotNets 2019. Princeton, NJ, USA (November 2019).
For more details, see: https://nsg.ee.ethz.ch

Researchers

Roland Meier (ETH)
Laurent Vanbever (ETH)

Industry partner

armasuisse

Self-securing Networks

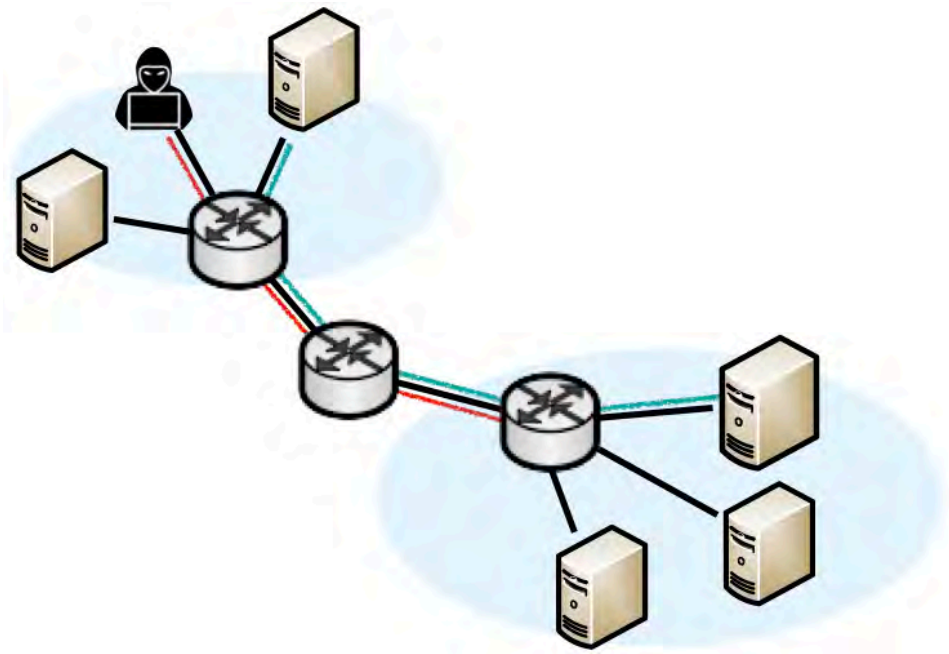

The goal of this project is to build data-driven network infrastructures that can autonomously protect, detect and defend themselves against attacks. We intend to develop network-specific learning and inference algorithms that can run directly in the data plane, in real-time, to perform tasks that are 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.

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

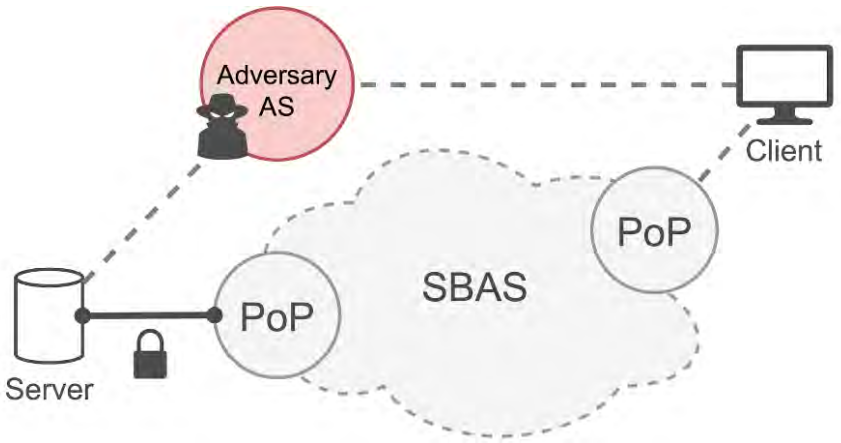

The recent Facebook outage went on record as one of the largest outages for a major application provider. With the root cause for Facebook, Instagram, and WhatsApp going offline being the BGP routing protocol, there is more awareness than ever that more reliable approaches are required to route Internet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

Robotic stacking of parcels in containers and roll cages

Problem statement: The employees of Swiss Post do a great amount of manual work to load and unload parcels in containers. The unloading process can be automated mechanically (e.g. by tilting the roll cages), but an acceptable level of filling by the loading process cannot be achieved through a mechanical solution: parcels must be intelligently stacked in order to optimize the volume transported.
A robot could do this work, but the current technologies are based on a process in which every item to stack is known in advance. In the postal industry, it is impossible to know which item is coming next, although some properties of each item are known (e.g. size and weight). This project will focus on the development of an intelligent robotic system capable of loading containers and roll cages using low to no buffer.

Basic research: Solving the problem stated above requires intelligent robots that know how to dynamically manipulate rigid boxes. This task requires specialized motion planning algorithms for 1) robust grasping and 2) collision-free trajectories to efficiently move boxes from the conveyor belt to their final location in the container.

Both sub-tasks must take into account the workspace of the robot. For example, if reachability is somewhat limited, then the boxes could be tossed gently, or placed down and pushed into their final spot. Such strategies, which are often employed by human workers, require robots to poses a deep understanding of contacts and friction, dynamics, robustness against unanticipated perturbations, dynamic regrasping strategies, etc. The ultimate goal of this PhD thesis is to endow robots with human-level skill when it comes to loading parcels.

Technical foundations that the CRL group will contribute to:

1. Physics-based simulation models that will let robots understand and predict the physical implication of their actions.

2. A differentiable simulator as the technical foundation for trajectory optimization algorithms that will generate dynamic motion plans.

3. Robotic tele-operation as the means to learn complex motion skills from demonstrations.

Researchers

Prof. Stelian Coros (ETH)
Dr. Roi Poranne (ETH)

Industry partner

SWISS POST

Manipulation of non-rigid e-commerce parcels

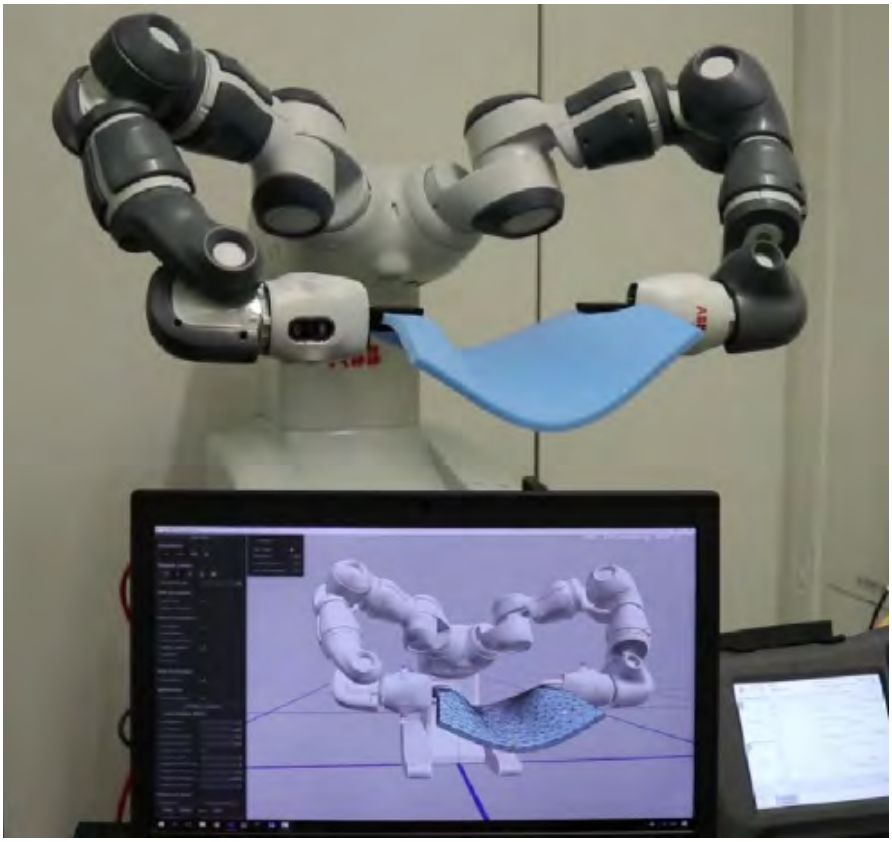

Problem statement: The Swiss Post knows how to process form-stable items (Letters or packages). For these types of items, we have appropriate technologies. With the growing volume of e-commerce items from Asia, however, we have the challenge that the range (material, size, unstable structure, surface pressure, different shapes …) of these items vary massively.

For this spectrum, we do not yet have suitable technologies in the postal industry. Therefore, the decision for the subsequent process is made by a manual (feel with the hand) and visual judgment by the employee. They touch, turn, look, bend the item.

Basic research: Solving the problem stated above requires robots that can dexterously manipulate soft, unstructured parcels and polybags. To this end, we will build on the model-based methodology my research group has recently introduced.

In particular, the goal of this thesis will be to develop technical foundations to allow the robot 1) to build an internal mechanical model of soft/unstructured parcels by feeling/scanning/manipulating the items, and 2), to autonomously understand how to grasp, pick up, and dynamically place the soft parcel on a conveyer belt in a prescribed configuration.

Technical foundations that the CRL group will contribute to:

1. Physics-based simulation models that will let robots understand and predict the physical implication of their actions.

2. A differentiable simulator as the technical foundation for trajectory optimization algorithms that will generate dynamic motion plans.

3. Robotic teleoperation as the means to learn complex motion skills from demonstrations.

Researchers

Prof. Stelian Coros (ETH)
Miguel Mora (ETH)
Oliver Stark (ETH)

Industry partner

SWISS POST

Research Projects

Secure Governance Schemes for Blockchains

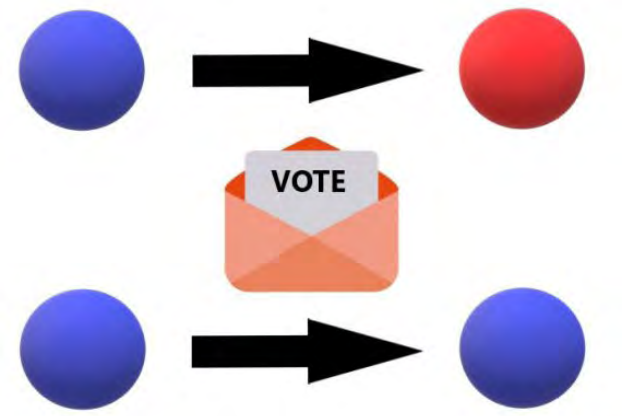

Systems based on blockchain technology are promising, as they can be decentralized and rendered robust against attacks. A blockchain is a (distributed) ledger, in which all transactions are recorded sequentially. Because such systems build on distributed consensus –i.e. they require a large number of 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)

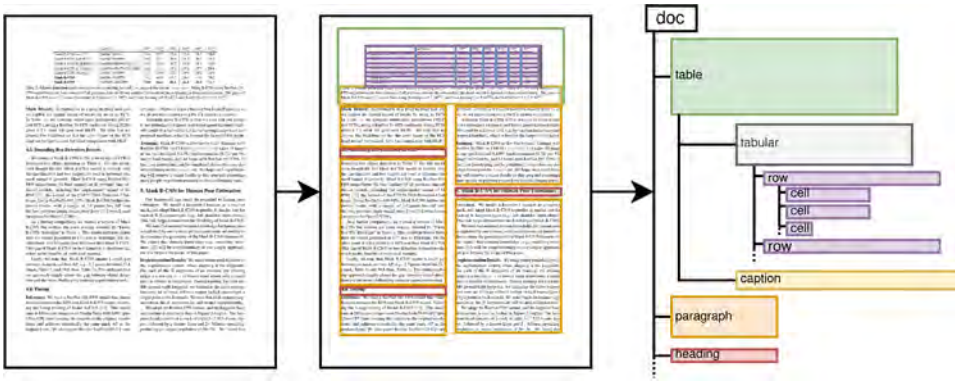

Automatic Visual Document Parsing

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.

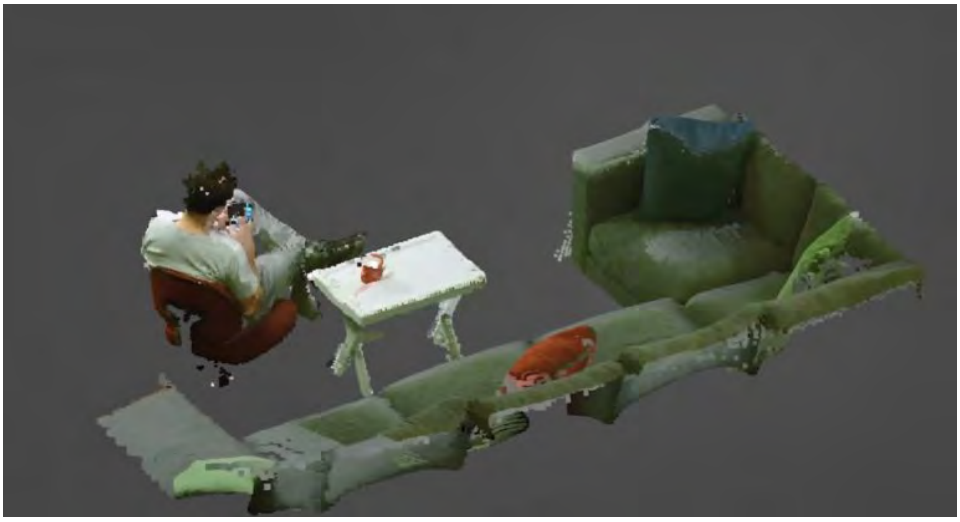

Approach

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.

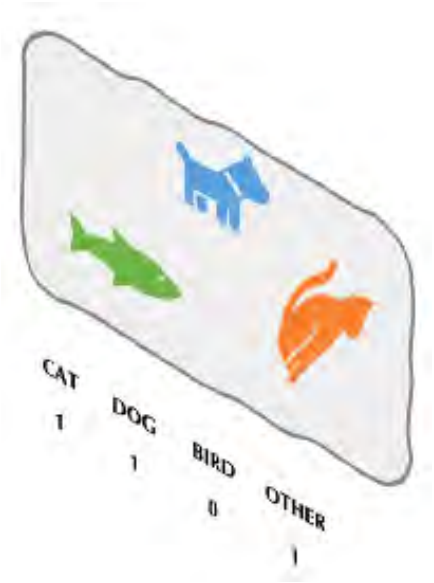

In the second part of the project, we perform a similar experimental study, this time on a real-world dataset provided by our partner, building on top of their system that is currently in use. This email-forwarding task is quite interesting due to the sheer particularities of the task and the dataset --- 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

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