ETH zürich

Zurich Information Security and Privacy Center (ZISC)

Annual Review 2020

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

Welcome

The year 2020 was an unprecedented one, as the the Covid-19 pandemic disrupted many aspects of our lives. Also the faculty and the researchers of the ZISC center, and its industry partners, needed to create new ways to keep on working together. Despite such significant challenges, the ZISC center continued delivering 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 partners, we strengthened our collaboration with the industry by starting several new projects. To mention one example, the ZISC researchers are exploring how hardware-assisted security mechanisms paired together with privacy-preserving machine learning techniques can help to solve data collection and processing challenges that are major issues in today's cyber insurance industry. The industry partner in this work is Zurich Insurance. As another example, the ZISC researchers are exploring how augmented reality and virtual reality techniques can address various challenges ranging from corporate training to increased public awareness regarding the company's activities. The industry partner in this research is Die Post. More about these applied research projects, and other similar collaborations, will be explained in the following pages of this report.

In addition to such applied research, the ZISC researchers kept on delivering excellent results on several research projects that address fundamental questions and challenges in information security and privacy. To pick one such example, the project lead by Prof. David Basin developed novel techniques for verification of security properties of software implementations. Such ground-breaking results pave the way towards follow-up applied research projects where similar techniques can be applied to software systems of high importance such as electronic voting.

To mention another example, the research project lead by Prof. Ueli Maurer is developing new ways to model players with quantum computing abilities in cryptographic systems. Such research is at the same time both long-ranging and very topical, as the prospect of practical quantum computers becomes more relevant year by year. More about such basic research activities will be explained later in this report.

During 2020, the members of the ZISC center made also significant contributions to other projects of high societal importance. To highlight one such example, the ZISC faculty members Prof. David Basin, Prof. Srdjan Capkun and Prof. Kenny Paterson were an integral part of an international team of researchers who designed a privacy-preserving contact tracing mechanism known as the DP-3T. This design was adopted by both Google and Apple as the foundation for their exposure notification API that is used by the SwissCovid App and numerous other national contact tracing applications. The ZISC center is extremely proud of this work.

Finally, the ZISC center keeps on growing. During 2020, two new faculty members, Prof. Dennis Hofheinz and Prof. Shweta Shinde, joined the center. With these two new professors the ZISC center increases further its competence and expertise in research areas such as system security, program analysis and cryptography. During this year, the ZISC center also adjusted its leadership structure, with Prof. Srdjan Capkun becoming the Chair of the center and Dr. Kari Kostiainen taking over the role of the Director of the center.

The ZISC center wishes all its partners and collaborators a relaxing holiday season and we are looking forward to working with you again in 2021!

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

Associate Partner

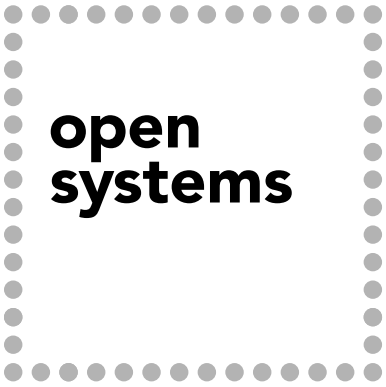

Research Highlights 2020

The EMV Standard: Break, Fix, Verify

EMV, named after its founders Europay, Mastercard, and Visa, is the worldwide standard for smartcard payment. As of December 2019, EMV is being used in over 9 billion cards worldwide. Banks have a strong incentive to adopt EMV due to the liability shift, which relieves banks using the standard from any liability from payment disputes.

EMV: 20 Years of Vulnerabilities

Besides the liability shift, EMV's global acceptance is also attributed to its advertised security. However, EMV's security has been challenged numerous times. Card cloning, downgrade attacks, relay attacks, and card skimming are all examples of successful exploits of the standard's shortcomings. Some of the security issues identified result from flawed implementations of the standard. Others stem from logical flaws that are hard to spot in EMV's lengthy and complex specification, running over 2,000 pages. This raises the question of how we can systematically explore all possible executions and improve the standard to avoid another twenty years of attacks.

Approach Taken: Break, Fix, Verify

In our work we focus on weakness of and improvements to the EMV protocol design. We present a formal, comprehensive model of the EMV standard that accounts for the three card authentication modes, the five cardholder verification methods, the two types of transaction authorization (offline and online), and the two major types of contactless transactions (Visa and Mastercard).

The model is specified in Tamarin, a state-of-the-art verification tool that has been used to analyze real-world protocols such as TLS 1.3 and 5G. An unprecedented feature of our model is that it considers the three roles present in an EMV session: the bank, the terminal, and the card. Previous symbolic models merge the terminal and the bank into a single agent, resulting in the unrealistic assumption that the terminal can check all card-produced cryptographic proofs that the bank can.

Using our model, we have identified various security issues, including the lack of protection against modification of the data that encodes the cardholder verification method used in Visa contactless transactions. We developed a proof-of-concept Android application that exploits this to bypass PIN verification.

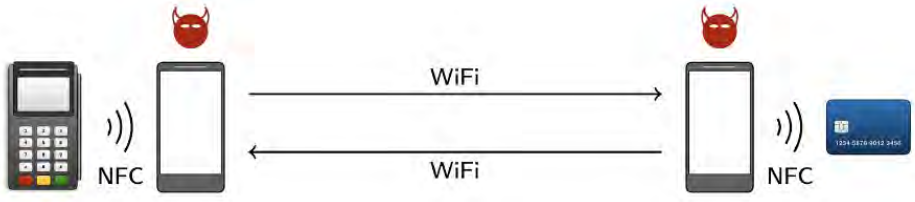

Our app implements the attack on top of a relay attack architecture, thus employing two phones: one that is held near the payment terminal and another that is held near the victim's card (see picture p.7).

Concretely, our man-in-the-middle attack modifies a card-produced message in a way that instructs the terminal that the PIN is not required because the cardholder was verified on the consumer's smartphone. This enables criminals to use a stolen Visa contactless card to pay for expensive goods without knowing the card's PIN. In other words, **the PIN is useless in Visa contactless cards**!

We successfully tested our PIN bypass attack on real-world payment terminals with various Visa-branded cards, including Visa Credit, Visa Electron and V Pay cards. Screenshots of the app are given below. In particular, the screenshot on the left is what the criminal shows at check-out. Notice that it looks just like legitimate payment apps such as Apple Pay or Google Pay.

Finally, we have proposed and machine-checked fixes to the security problems found. The proposed improvements can be deployed with software updates on the payment terminals and therefore do not affect the cards in circulation.

Further information
Demo video of the attack and other information are available at https://emvrace.github.io/.

The EMV Standard: Break, Fix, Verify
David Basin, Ralf Sasse, and Jorge Toro-Pozo
In Proceedings of the 42nd IEEE Symposium on Security and Privacy (S&P 2021), San Francisco, CA, USA, May 23-27, 2021. To appear. Available on arXiv at https://arxiv.org/abs/2006.08249

Reseachers
Prof. David Basin, Dr. Ralf Sasse and Dr. Jorge Toro-Pozo

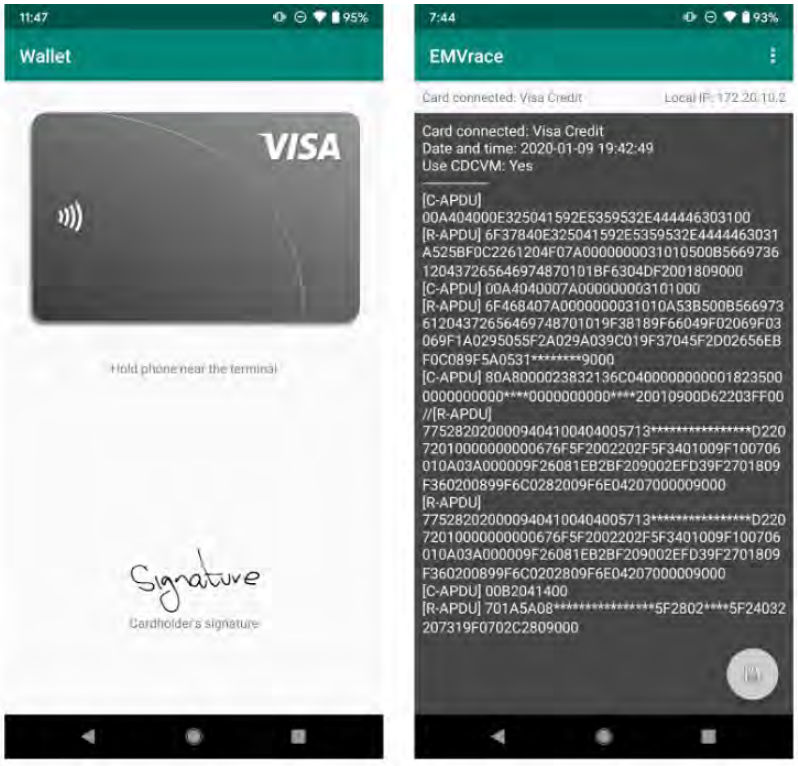

Research Highlights 2020

RDMA Security Vulnerabilities, Attacks, and Mitigations

In recent years, state-of-the-art systems started to leverage remote direct memory access (RDMA) technologies as a high performance communication mechanism and deployments of RDMA in public clouds, such as Microsoft Azure and IBM Cloud, became available. Hence, the security of RDMA architectures is crucial, yet the design of RDMA architectures is mainly focused on performance rather than security and potential security implications and dangers of using RDMA communication in upper layer protocols remain largely unstudied. During the past year, the Network Security Group started a collaboration with the Scalable Parallel Computing Lab (SPCL) that focused on research of RDMA security mechanisms and developing techniques to improve the security of RDMA technologies.

ReDMArk shows that current security mechanisms of InfiniBand (IB) based architectures are insufficient against both in-network attackers and attackers located on end hosts, thus affecting not only secrecy, but also integrity of RDMA applications. We demonstrate multiple vulnerabilities in the design of IB-based architectures and implementations of RDMA-capable network interface cards (RNICs) and exploit those vulnerabilities to enable powerful attacks such as packet injection using impersonation, unauthorized memory access, and Denial-of-Service (DoS) attacks.

One crucial mitigation to secure communication for systems relying on RDMA is provided by our work sRDMA. sRDMA provides strong authenticity and secrecy, prevents several forms of DoS attacks, and employs network interface cards to perform cryptographic operations. Thus, safety- and security-critical applications that rely on RDMA can use sRDMA to prevent attacks by malicious entities within the same network.

In summary, our work shows the negligence of security in RDMA architecture in favor of focusing on high performance and demonstrates implications of these security flaws. Furthermore, it provides a guideline for the development of upcoming versions of RDMA protocols by suggesting different mitigations, such as strong encryption and authentication. Finally, our work increases the awareness of developers of RDMA-enabled systems regarding security threats introduced by RDMA networking.

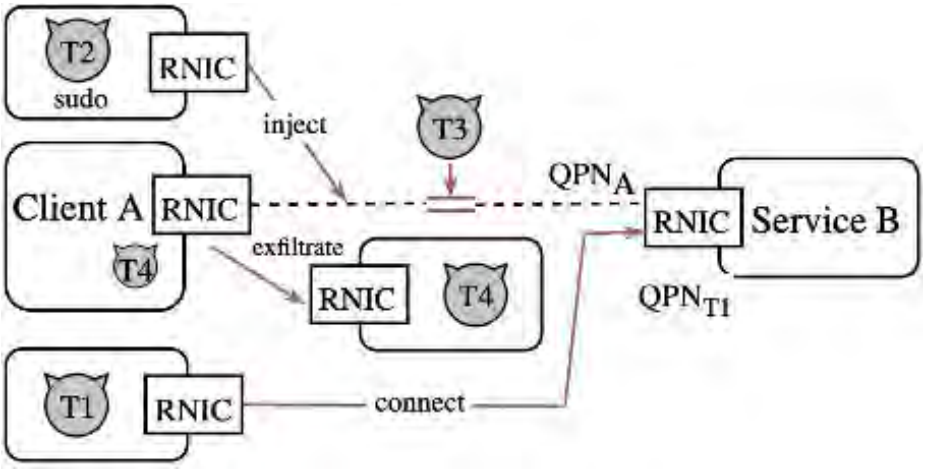

ReDMArk – Bypassing RDMA Security Mechanisms

Current RDMA technologies include multiple plaintext access tokens to enforce isolation and prevent unauthorized access to system memory. As these tokens are transmitted in plaintext, any entity that obtains or guesses them can read and write memory locations that have been exposed by using RDMA on any machine in the network, compromising not only secrecy but also integrity of applications. For example, when memory pages with code are changed remotely, altering packet contents enables remote code injection. To avoid compromise and tampering of these access tokens, RDMA architectures rely on isolation and the assumption that the underlying network is a well-protected resource. Otherwise, an in-network attacker that is located on the path between two communicating parties (e.g., bugged wire or malicious switch) can eavesdrop on or modify access tokens of bypassing packets, with potentially drastic consequences.

ReDMArk shows that current security mechanisms of IB-based architectures are insufficient against both in-network attackers and attackers located on end hosts, thus affecting not only secrecy, but also integrity of RDMA applications. We demonstrate multiple vulnerabilities in the design of IB-based architectures and implementations of RDMA-capable network interface cards (RNICs) and exploit those vulnerabilities to enable powerful attacks such as packet injection using impersonation, unauthorized memory access, and Denial-of-Service (DoS) attacks. To thwart the discovered attacks in ReDMark we propose multiple mitigation mechanisms that are deployable in current RDMA networks.

sRDMA – Efficient NIC-based Authentication and Encryption for Remote Direct Memory Access

To defend against in-network attackers, RDMA networks would require cryptographic authentication and encryption, which is not part of current RDMA specifications. While IPsec transport recently became available for RoCE traffic, the IPsec standard does not support InfiniBand traffic. Furthermore, application-level encryption (e.g., based on TLS) is not possible, since RDMA read and write operations can operate as purely one-sided communication routines and thus can be handled without involvement of the CPU of the remote system. As application-level encryption protocols cannot support purely one-sided communication routines, the applications would need to store packets in a temporal buffer before decryption, which would cause high overhead and completely negate RDMA's performance advantages.

In our work, we propose sRDMA, a protocol that provides efficient authentication and encryption for RDMA to prevent information leakage and message tampering. sRDMA uses symmetric cryptography and employs network interface cards to perform cryptographic operations. Additionally, we provide an implementation for sRDMA using programmable network adapters.

Further information

ReDMArk: Bypassing RDMA Security Mechanisms.
Benjamin Rothenberger, Konstantin Taranov, Adrian Perrig, and Torsten Hoefler.
In USENIX Security Symposium (USENIX Security 21) 2021.

sRDMA: Efficient NIC-based Authentication and Encryption for Remote Direct Memory Access.
Konstantin Taranov, Benjamin Rothenberger, Adrian Perrig, and Torsten Hoefler.
In Proceedings of the USENIX Annual Technical Conference (USENIX ATC) 2020.

Reseachers

Benjamin Rothenberger, Konstantin Taranov, Prof. Adrian Perrig and Prof. Torsten Hoefler.

Research Highlights 2020

SWiSSSE: System-Wide Secure Searchable Symmetric Encryption

Cloud computing allows clients (individuals and organizations) to store, process and query large databases using third party data centres. Ensuring the privacy of these databases and the privacy of the queries/operations executed on them is a major concern for clients. While traditional encryption techniques allow protecting data at rest and in transit, they do not allow computing directly over encrypted data. In theory, techniques such as fully homomorphic encryption (FHE) and oblivious random-access memory (ORAM) can be used to solve this problem. In practice, such solutions incur prohibitively large costs and do not scale efficiently to large databases with billions of records/files.

ETH researchers from the Applied Cryptography group (Prof. Kenny Paterson, Dr. Sikhar Patranabis) have been working in collaboration with researchers from the University of Bristol (Zichen Gui, Prof. Bogdan Warinschi) on Searchable Symmetric Encryption (SSE). This is a special form of encryption that aims to allow directly querying large encrypted databases (both structured and unstructured) in a secure yet efficient manner.

Over the past year, the research team has identified and addressed several limitations in state-of-the-art proposals for SSE with respect to security. The team identified a new class of leakages in existing SSE schemes called «system-wide leakages» and demonstrated how these can be exploited to completely break the data and/or query privacy guarantees of these schemes. Their findings highlight that there is a fundamental mismatch between the models and schemes for SSE being developed in the scientific community and the real-world threat models that such schemes should be addressing.

The team addressed these shortcomings by providing comprehensive security models and notions for SSE that encompass leakages from the whole SSE system. They also proposed and prototyped SWiSSSE (System-Wide Secure SSE) – the first family of SSE schemes that support fast updates and queries over encrypted databases while achieving security against system-wide leakages. SWiSSSE comes in two flavours – static and dynamic. Static SWiSSSE allows searches for keywords over a fixed set of encrypted documents whilst minimising leakage. Dynamic SWiSSSE augments the static approach by enabling documents (and their keywords) to be added to and deleted from the encrypted database. A particularly novel aspect of dynamic SWiSSSE is that it has oblivious operations. This means that the server cannot distinguish what kind of operation (search, add or delete) is being carried out at any given moment. This makes it significantly easier to analyse the security of dynamic SWiSSSE and pinpoint the limited leakage that it has.

The team also went on to apply and extend state-of-the-art cryptanalysis techniques to try to understand the potential security impact of the residual leakage that

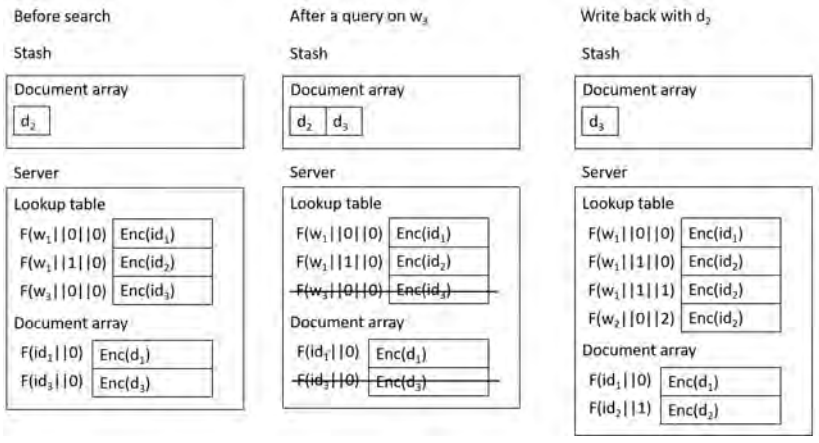

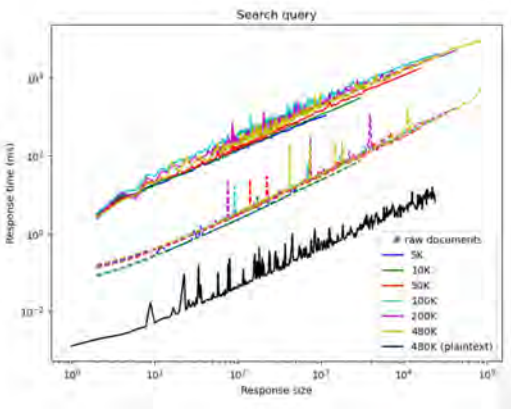

SWiSSSE has. This helped the team to identify how to select practical parameters for SWiSSSE. The team also implemented a prototype of SWiSSSE and experimented with it on a document repository of significant size (the Enron emails, containing over. 500K documents and over 30M keyword-document pairs). The team's benchmarking shows that dynamic SWiSSSE incurs a roughly factor of 10 performance overhead when compared to an unencrypted system. This overhead could be reduced significantly by employing a lower-level programming language and more careful cryptographic optimisations.

In its future work, the team plans to further optimise the performance of SWiSSSE and to increase its functionality, by finding ways to extend it to support more complex search queries over encrypted data. The ultimate goal is to build a fully encrypted database system that supports SQL-like functionality whilst minimising system-level leakage.

In summary, the Applied Cryptography group and its collaborators are making significant progress towards transforming SSE into a scalable and widely-deployable technology that equips outsourced databases with strong, theory-backed security guarantees.

To read more about the group, please visit: https://appliedcrypto.ethz.ch/

Further information

SWiSSSE: System-Wide Security for Searchable Symmetric Encryption
Zichen Gui and Kenneth G. Paterson and Sikhar Patranabis and Bogdan Warinschi
In: Cryptology ePrint Archive, 2020

Reseachers

Zichen Gui, Prof. Kenneth G. Paterson, Dr. Sikhar Patranabis and Bogdan Warinschi

Dr. Sinisa Matetic, Varun Maram, Mia Filic, Dr. Patrick Towa, Dr. Anu Unnikrishnan
Matilda Backendal, Dr. Felix Günther, Prof. Dr. Kenny Paterson, Dr. Benjamin Dowling, Lukas Burkhalter, Jan Gilcher
Dr. Sikhar Patranabis, Barbara von Allmen Wilson, Dr. Igors Stepanov, Alexander Viand

Research Highlights 2020

Message Time of Arrival Codes: A Fundamental Primitive for Secure Distance Measurement

Navigation systems, contactless payment, and radio-based access systems are some of the many applications that rely on a notion of distance or proximity between two devices. A transmitter-receiver pair can estimate physical distance by measuring the time a signal takes to travel from one device to another. The integrity of the signal›s time of arrival is critical to the security of the applications that build upon it:

If cars miscalculate their distance while driving autonomously, they may cause a collision. If an adversary can manipulate an electronic key›s distance estimation, the underlying access control system can be subverted. Attacks on electronic car-keys have gained notoriety over the years as major media outlets reported multiple car thefts. These attacks typically involved radio devices that relay the signal between car and key.

With programmable, versatile signal generation equipment becoming cheaper and easier to use, researchers expect such attacks to become even more widespread and more elaborate in the future.

In the wake of this development, ETH researchers from the System Security and Applied Cryptography Groups have conducted an effort to formally define the security properties of signals that resist all distance-modifying attacks. The resulting MTAC-primitive serves as a design reference for any radio signal used for secure time of arrival measurement.

While a significant research effort has already gone into the higher-level protocols that allow two entities to establish their proximity securely, this work extends the security formalism to the physical layer, which is responsible for generating and receiving signals.

This step is crucial since weak physical-layer security can break a system irrespective of higher-level protocols and no matter how strong the cryptographic algorithms are.

Physical-layer design for secure distance estimation is subject to challenging design constraints: Signals allowing precise time-of-arrival measurement typically are pulsed and, in consequence, wide in spectrum. On the other hand, regulators heavily limit the power levels of such ultra-wideband signals. Sending a pulsed signal over a meaningful distance requires a level of time-redundancy, which can be at odds with its security against adversarial on-the-fly manipulation.

The researchers propose a procedure according to the formalism that satisfies these real-world design constraints. They show that the performance-security tension can be managed for system parameters that are in line with the recently updated IEEE 802.15.4 standard for ultra-wideband communication.

In the future, the researchers aim to apply the primitive to wireless technologies that play a role in upcoming positioning systems: With the ongoing deployment of 5G, there is a trend to wider frequency bands in the cellular domain. While creating a potential for precise positioning within the licensed spectrum, this also gives rise to new physical-layer security challenges that can be analyzed with the proposed formalism.

Further information

Message Time of Arrival Codes: A Fundamental Primitive for Secure Distance Measurement
Patrick Leu, Mridula Singh, Marc Roeschlin, Kenneth G. Paterson, Srdjan Capkun
in IEEE Symposium on Security and Privacy (S&P), 2020

Reseachers
Patrick Leu, Mridula Singh,
Dr. Marc Roeschlin,
Prof. Kenneth G. Paterson,
Prof. Srdjan Capkun

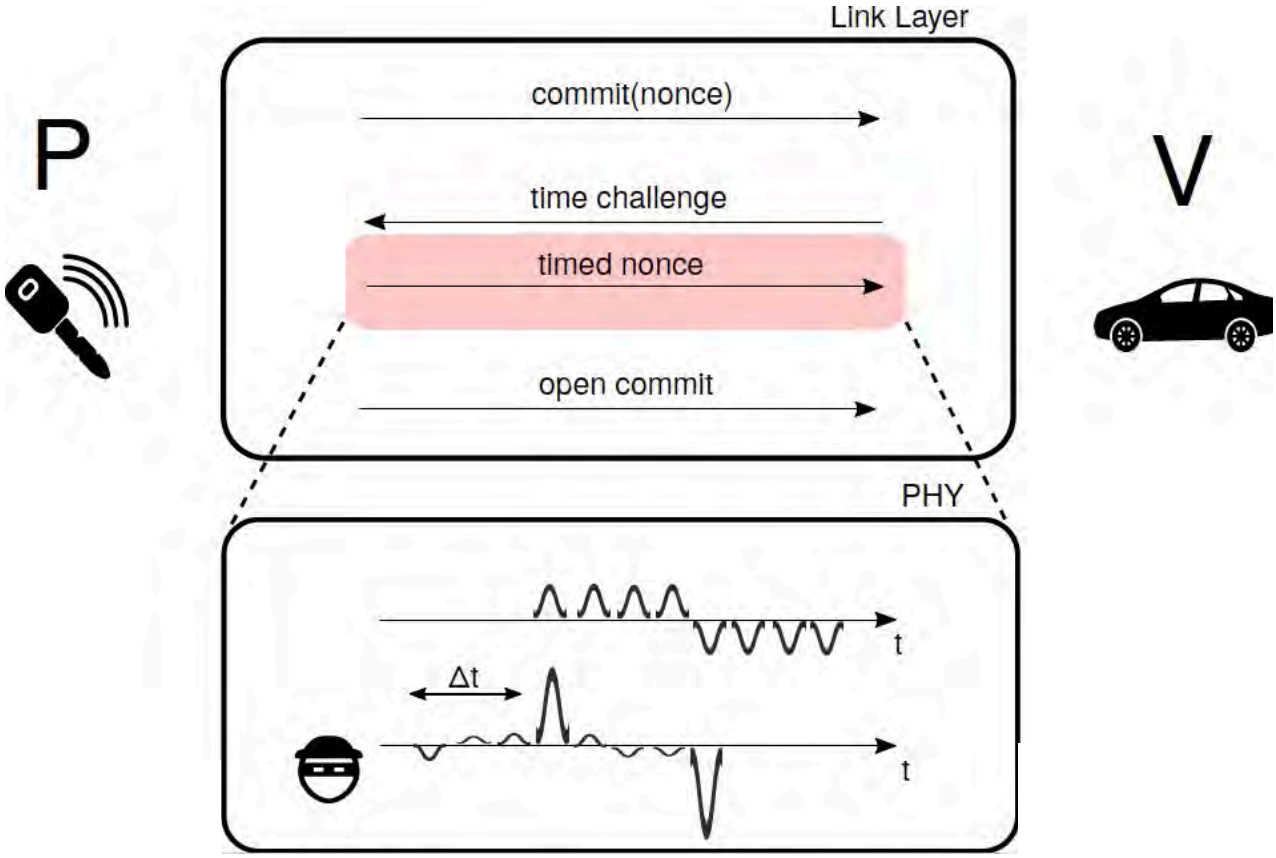

New Faculty Members

We are happy to welcome Prof. D. Hofheinz and Prof. S. Shinde to our team

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

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

Education 2020

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, 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 testbook series «Einfach Informatik» 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!

Main Research Areas

Smartphone Security

In this project, we focus on smartphone security. In particular, we look at how smartphones can enhance the security of our daily activities as well as how secure is data stored by users on smartphones.

Throughout our work we highlight the interaction of security with usability and deployability — two key components that cannot be ignored when designing and analyzing a secure system. We will see how in some cases decreasing or removing the user interaction requirements from a system render it more secure. In other cases, in contrast, it is the user interaction and attentiveness that play an important role in safekeeping. the data stored on a user's smartphone.

Monitoring

It is a growing concern for companies, administrations, and end users alike whether IT systems comply with policies regulating the usage of sensitive data. Checking compliance is particularly acute as our modern infrastructures (communication, entertainment, finance, etc.) collect, process, and share data.

A prominent approach to compliance checking is runtime monitoring. Here, system actions are observed and automatically checked for compliance against a given policy. We develop efficient and scalable monitoring algorithms for expressive policy specification languages, e.g., metric first-order temporal logic. We are also interested in policy enforcement, that is, preventing policy violations instead of only detecting them.

Future Internet Architecture SCION

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

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

Secure Positioning and Localization

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.

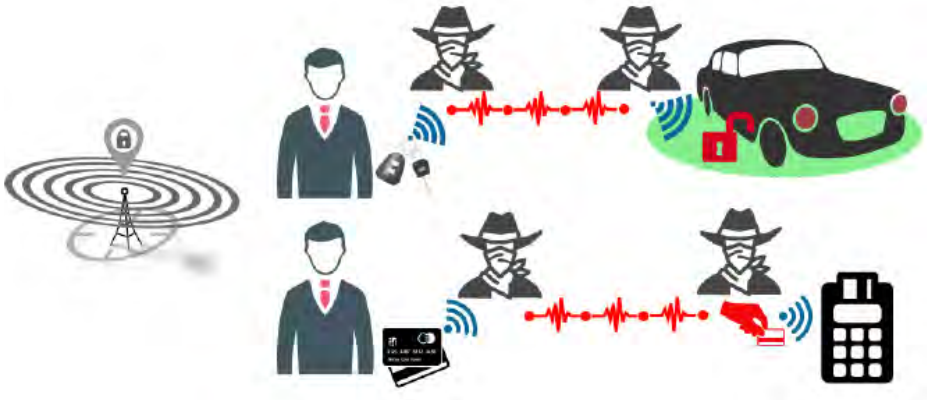

Main Research Areas

Access control

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

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

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

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

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.

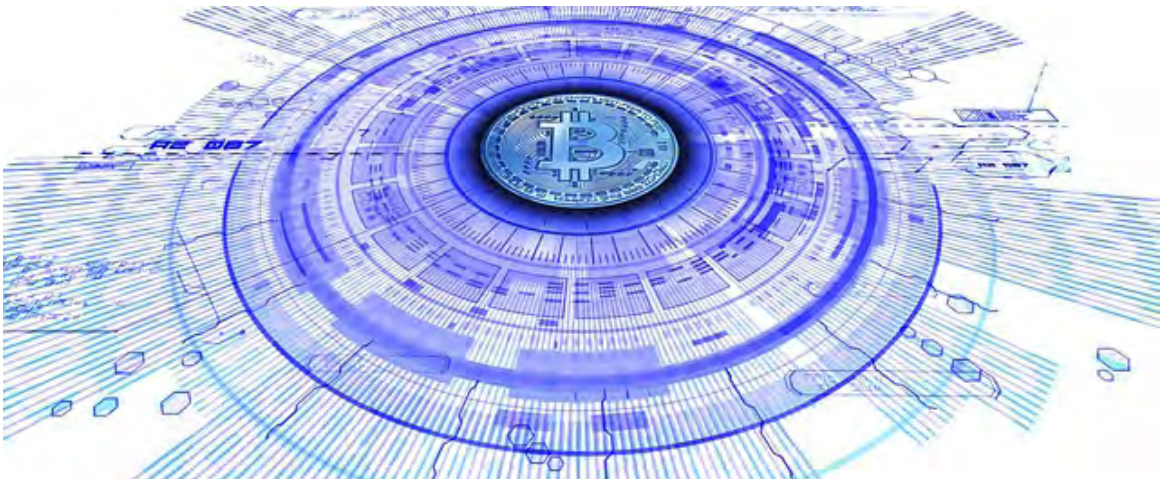

Research Projects

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.

Markus Legner, Tobias Klenze, Marc Wyss, Christoph Sprenger, and Adrian Perrig.
EPIC: Every Packet Is Checked in the Data Plane of a Path-Aware Internet.
In Proceedings of the USENIX Security Symposium 2020.

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.

Reseachers

Various members of the Network Security Group.

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.

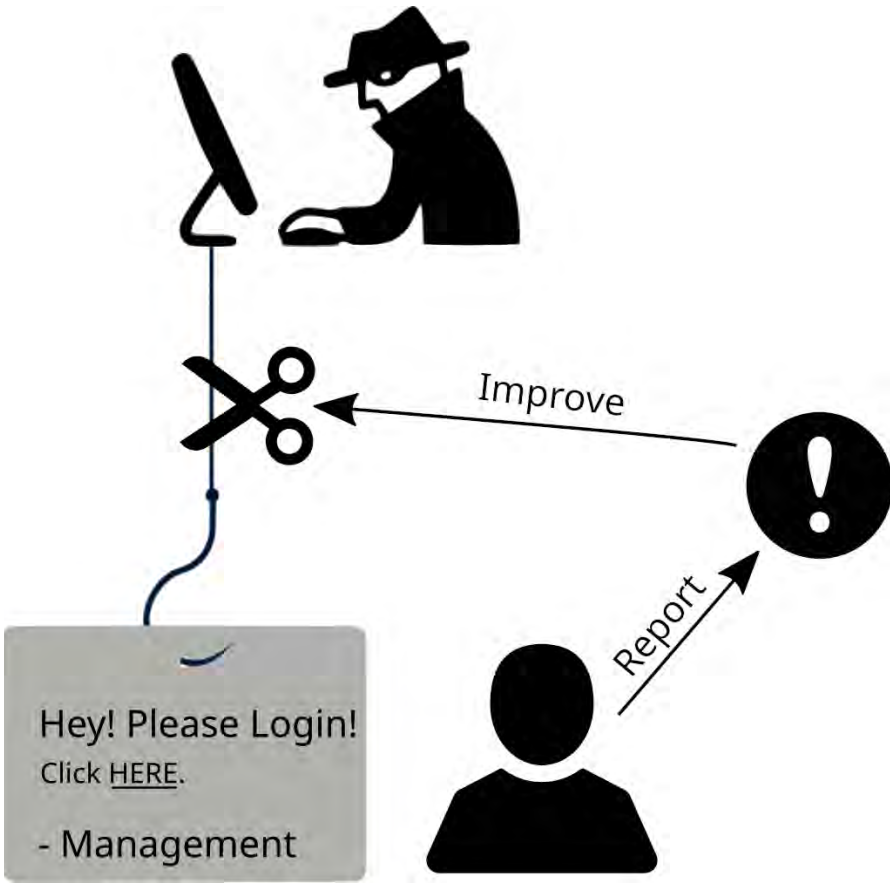

In this project, in partnership with the Swiss Post, we aim to conduct a large-scale study on phishing prevention, detection, and education. Users will be involved in phishing detection, by having the ability to report suspicious emails and to get feedback by automated analyses and human analysts after 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.

Industry partner

Swiss Post

Reseachers

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

Research Projects

Formal Methods for Federated Identity Management

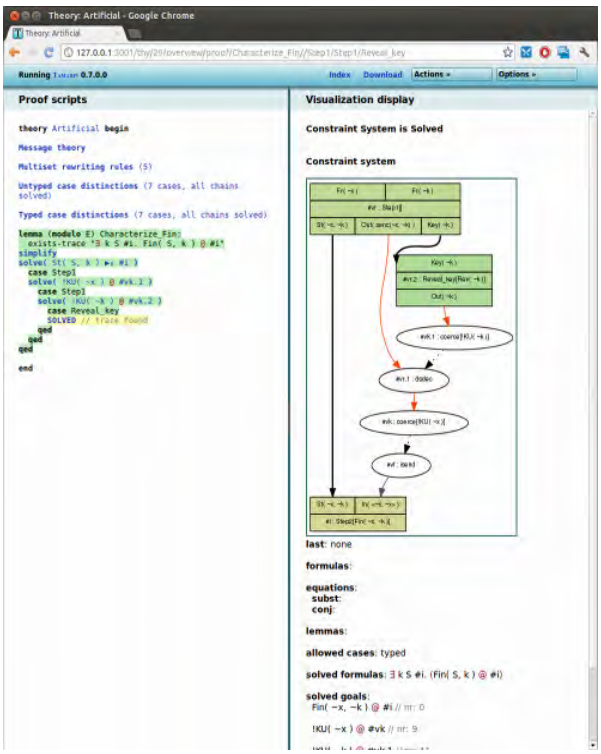

Single sign-on based on the OpenID Connect protocol are widespread in the Internet today, allowing users to use their account with an identity provider (IdP) to log in to other services, called relying parties (RPs). This complex web of IdPs, RPs, and users, is referred to as «federated identity management». While the security of OpenID Connect has been proven, there have been unresolved privacy issues with respect to IdPs. In particular, the IdP learns to which RPs its users log in, and the exact time of these logins. This information may reveal sensitive behavioral information about the user. Thus, users must trust the IdP to handle this information confidentially and protect it against potential attackers. We are happy to report that since last year›s report, two research papers related to this project have been published.

First, we have developed a variant of OpenID Connect, called «Privacy-Preserving OpenID Connect» (POIDC), which solves the mentioned privacy issue.

In particular, it prevents the IdP from learning to which RPs the user logs in; the IdP only learns that a login happened. However, it cannot even distinguish between repeated logins to the same RP and logins to different RPs.

Furthermore, we have proven that POIDC fulfills the same desired security properties as standard OpenID Connect. For this, we used the state-of-the-art protocol verification tool Tamarin. This work has been published this year in the Proceedings of the 15th ACM Asia Conference on Computer and Communications Security, ASIACCS 2020.

Second, a closely related research effort is that of «User Account Access Graphs»

We provide a methodology to systematically analyze a user›s entire account setup, which may include accounts at different IdPs, connecting to many RPs.

We also consider more traditional accounts using passwords and email-based account recovery options. We graphically model a user›s setup of accounts, credentials, devices, and their connections. This work has been published last year in the Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security, CCS 2019.

Industry partner

ZKB

Reseachers

Prof. David Basin (ETH)
Sven Hammann (ETH)
Ralf Sasse (ETH)

Security of Avionics Communication Systems

Next-generation avionics communication systems were — in the majority — standardised several decades ago, when only the most apt attackers were able to even receive aircraft signals. Therefore, security was not on the radar of the standardisation bodies who mainly focused on the safety impact of these protocols. This mismatch between safety and security led to systems that would keep the aircraft operating in a safe manner, even if individual systems failed.

Additionally, military and commercial (e.g. Amazon, Swiss Post) drones will in the intermediate future coexist with civilian aircraft in public airspace. This coexistence requires unmanned aerial vehicles (UAVs) to handle civilian aircraft communication in addition to the communication channels that control the drone.

This forest of communication system is a possible target for manipulation, eavesdropping and more advanced attacks. Any manipulation of the individual systems might lead to unforeseen consequences, as human intervention might be inhibited by an attacker.

In this project, we investigate the security of many different (mostly wireless) communication protocols. We will examine novel attacks against GPS technology and also take a look at the Traffic Alert and Collision Avoidance System (TCAS) employed by civilian aircraft and future drone systems. Further, we want to review satellite communication which serves as a backbone for medium and large sized drone systems as well as a multitude of commercial aircraft.

The target of this project is to research the security and privacy of today's aircraft and UAV communication systems and, where possible and applicable, propose changes to ensure the safety and security of tomorrow's flight operations.

Industry partner

Armasuisse

Reseachers

Prof. Dr. Srdjan Capkun (ETH)
Daniel Moser (Armasuisse)

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.

Industry partner

NEC

Reseachers

Karl Wüst (ETH)
Kari Kostiainen (ETH)

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

Reseachers
Marta Mularczyk (ETH)
Ueli Maurer (ETH)
Daniel Jost (ETH)

Research Projects

Topology-Hiding Computation

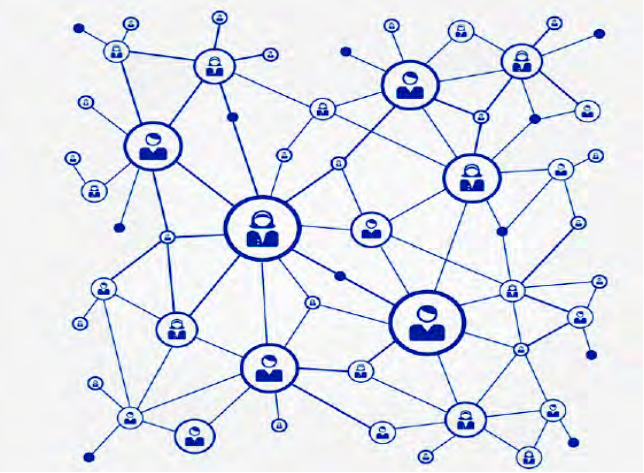

Secure communication over an insecure network is one of the fundamental goals of cryptography. The security goal can be to hide different aspects of the communication, ranging from the content (secrecy), the participants' identity (anonymity), the existence of communication (steganography), to hiding 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
Cryptology ePrint Archive – 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.

Reseachers
Martin Hirt (ETH)
Chen-Da Liu Zhang (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.

MONDRIAN: Comprehensive Inter-domain Network Zoning Architecture, upcoming at NDSS 2021

Industry partner

NEC

Reseachers

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

Research Projects

Improving Network Security Through Programmability

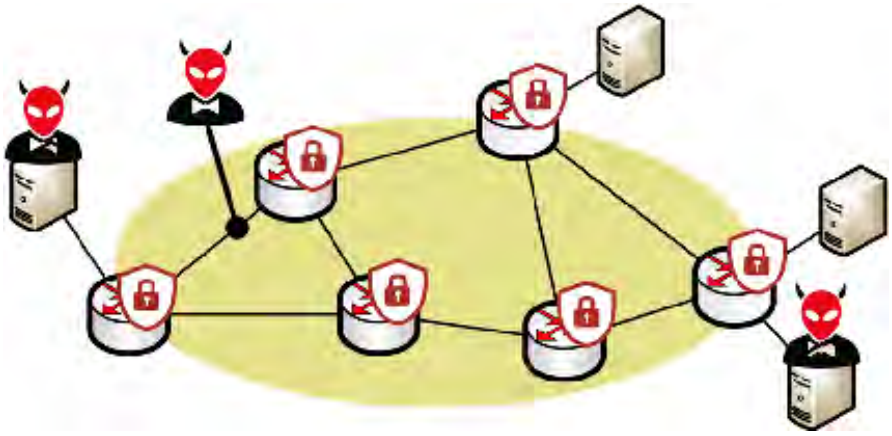

In this project, we argue that the network itself should be able to detect and mitigate attacks instead of relying purely on perimeter-based protection provided by dedicated appliances. To do so, we plan to leverage recent advances in network programmability which enable both the control plane and the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

Reseachers

Roland Meier (ETH)
Laurent Vanbever (ETH)

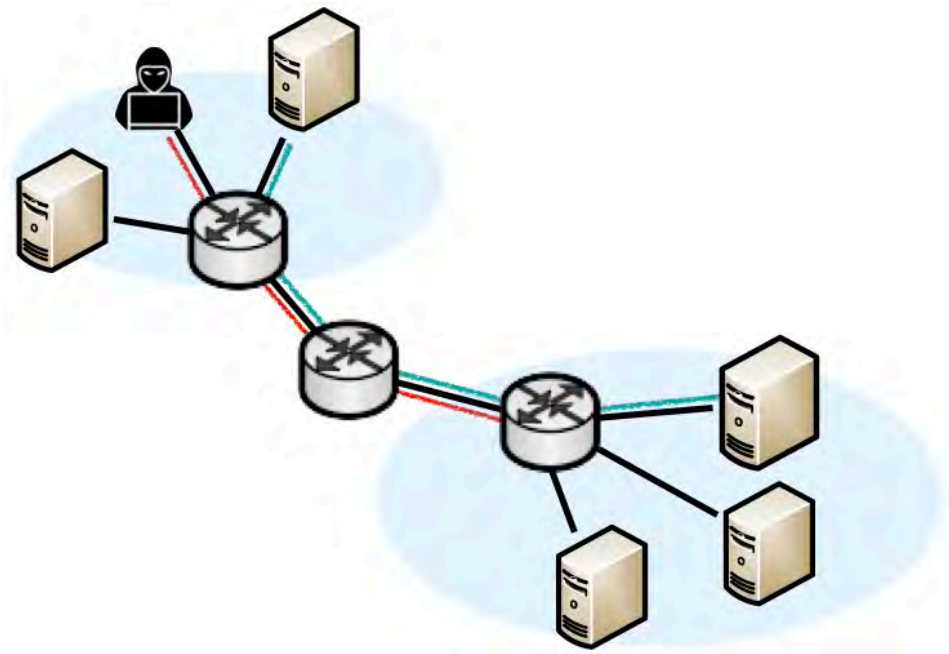

Self-securing Networks

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.

Industry partner

Armasuisse

Reseachers

Albert Gran Alcoz (ETH)
Laurent Vanbever (ETH)

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 project, we examine the SCION protocols in detail and formally verify that they have the desired security properties. We first formalize the protocols and security guarantees, and then use techniques from refinement and interactive theorem proving for their verification. Finally, we extract from the proven assertions a low-level specification of the IO-behavior of SCION components.

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 2020 we published our "Igloo" framework, which realizes the sound link between protocol and code and provides the methodology for full-stack verification. We have also made substantial steps in the modelling of the SCION forwarding protocol and the tools required to verify its implementations.

Further information

Christoph Sprenger, Tobias Klenze, Marco Eilers, Felix A. Wolf, Peter Müller, Martin Clochard, and David Basin
Igloo: Soundly Linking Compositional Refinement and Separation Logic for Distributed System Verification
OOPSLA 2020.

Marco Eilers, Peter Müller, and Samuel Hitz
Modular product programs
TOPLAS 2019

Reseachers

Prof. David Basin (ETH)
Prof. Peter Müller (ETH)
Prof. Adrian Perrig (ETH)

Quantum players in constructive cryptography

Quantum mechanics is one of the most successful physical theories, and has been verified by numerous experiments. But what does this imply for cryptography? On one hand, adversaries may have abilities that are not captured by a "classical" adversary. On the other, the (honest) users may also use quantum technology to increase the security of their protocols. But before being able to formulate the risks and benefits of quantum players, one needs cryptographic models and security definitions that encompass such parties.

The goal of this project is to model quantum players in the constructive cryptography framework of Maurer and Renner. The first part of the project involves modifying the framework itself so that it has the power need to capture such quantum players. For example, quantum mechanics allows a message to be in a superposition of sent and not sent, or a superposition of sent to Alice and sent to Bob, which needs to fit in the underlying communication model used by the framework. Furthermore, one may consider various message scheduling models, e.g., sequential scheduling (the players are activated one after the other), time-based scheduling (the time it takes to send and receive messages is explicitly modeled, and used to determine the order in which messages are processed) and non-determiistic scheduling (one computes all possible orders of messages and looks at the worst case). This projects studies these different scheduling models in the quantum context.

The second part of the project consists in using the framework to model cryptographic security in various applications. For example, we wish to find the best way to model CPA and CCA attacks on schemes that encrypt quantum messages. Another example is to study device-independent cryptography, and model the reuse of devices in a composable framework. It is indeed well-known that current security proofs only hold for devices that are used just once.

Further information

Fabio Banfi, Ueli Maurer, Christopher Portmann, Jiamin Zhu.
Composable and Finite Computational Security of Quantum Message Transmission.
TCC 2019, LNCS, Springer, vol. 11891, pp. 282–311, 2019.

Christopher Portmann
Quantum Authentication with Key Recycling.
Advances in Cryptology – EUROCRYPT 2017 – Proceedings, Part III, pp. 339-368, 2017.

Reseachers

Christopher Portmann (ETH Zurich)
Ueli Maurer (ETH Zurich)

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)

Manipulation of non-rigid e-commerce parcels

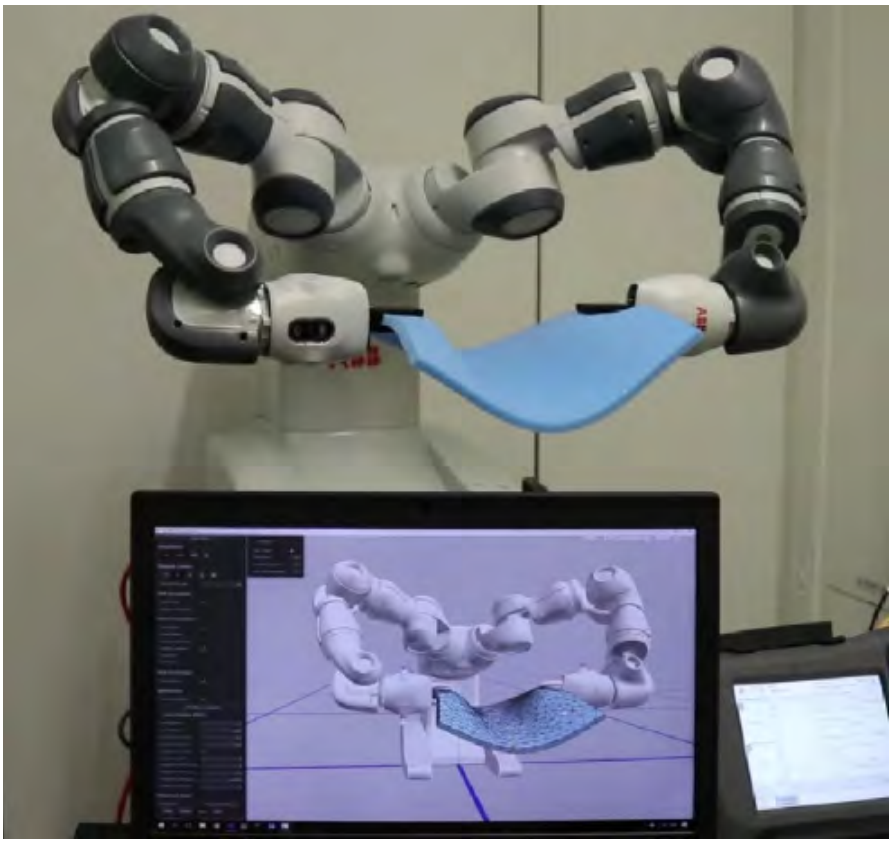

Problem statement: The Swiss Post knows how to process form-stable items (Letters or packages). For these types of items, we have appropriate technologies. With the growing volume of e-commerce items from Asia, however, we have the challenge that the range (material, size, unstable structure, surface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

Reseachers

Prof. Stelian Coros (ETH)
Miguel Mora (ETH)
Dr. David Hahn (ETH)

Research Projects

Secure Governance Schemes for Blockchains

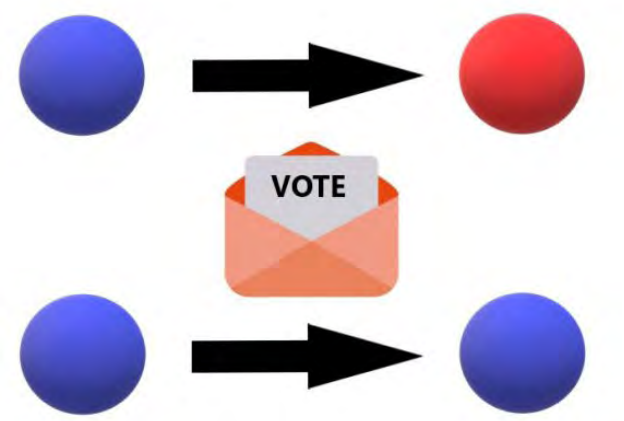

Systems based on blockchain technology are promising, as they can be decentralized and rendered robust against attacks. A blockchain is a (distributed) ledger, in which all transactions are recorded sequentially. Because such systems build on distributed consensus –i.e. they require a large number of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

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

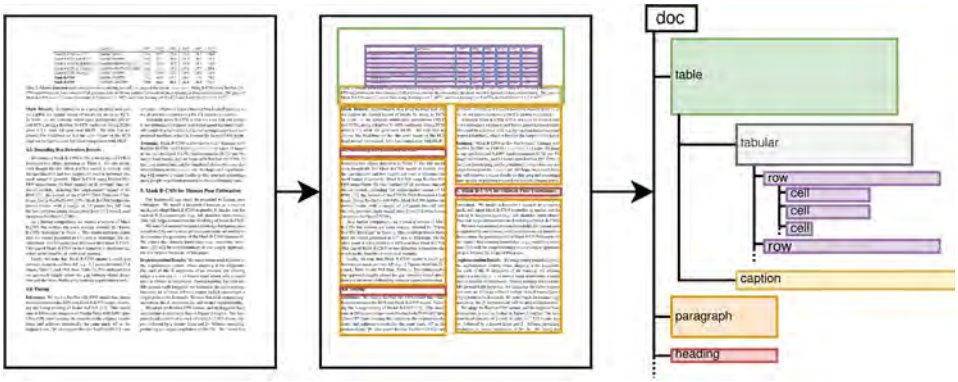

Automatic Visual Document Parsing

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.

Industry partner

Zurich

Reseachers

Ce Zhang (ETH),
Johannes Rausch (ETH)

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.

Industry partner

Zurich

Reseachers

Dr. Kari Kostiainen (ETH
Théo Giovanna (ETH)

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.

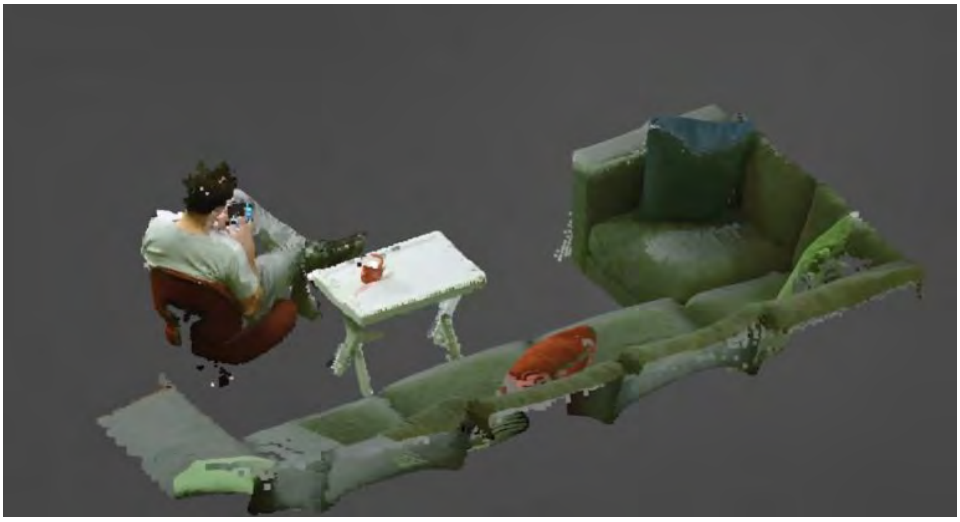

Approach

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.

Industry partner

Swiss Post

Researchers

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

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

People

ETH Faculty in ZISC
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 (ZISC director) 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. Kenneth Paterson leads the Applied Cryptography group. The group's research interests lie in all aspects of Cryptography, especially Applied Cryptography.

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.

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