Horizon 2020 Program (2014-2020)

Cybersecurity, Trustworthy ICT Research & Innovation Actions
Security-by-design for end-to-end security
H2020-SU-ICT-03-2018

Cyber security cOmpeteNCe fOr Research anD InnovAtion [†]

Work package 1: European Secure, Resilient and Trusted Ecosystem (ESRTE)

Deliverable D1.1: 1st Year Report on Designing and Developing an European Secure, Resilient and Trusted Ecosystem (ESRTE)

This deliverable describes the research activities undertaken by work package WP1 of the CONCORDIA Horizon 2020 project during the first year.

Contractual Date of Delivery	2019-12-31
Actual Date of Delivery	2019-12-28
Deliverable Dissemination Level	Public
Editors	Sotirios P. Chatzis (CUT) Abhilash Hota (JUB) Mattijs Jonker (UT) Jean-Yves Marion (UL) Aiko Pras (UT) Vasileios Prevelakis (TUBS) Nikos Salamanos (CUT) Jürgen Schönwälder (JUB) Michael Sirivianos (CUT) Neeraj Suri (TUD)
Contributors	All partners involved in WP1
Quality Assurance	Pavel Minařík (Flowmon) Tanja Pavleska (IJS)

[†] The research leading to these results has received funding from the European Union Horizon 2020 Program (2014-2020) under grant agreement n° 830927.

The CONCORDIA Consortium

CODE	Research Institute CODE (Coordinator)	Germany
FORTH	Foundation for Research and Technology - Hellas	Greece
UT	University of Twente	Netherlands
SnT	University of Luxembourg	Luxembourg
UL	University of Lorraine	France
UM	University of Maribor	Slovenia
UZH	University of Zurich	Switzerland
JUB	Jacobs University Bremen	Germany
UI	University of Insubria	Italy
CUT	Cyprus University of Technology	Cyprus
UP	University of Patras	Greece
TUBS	Technical University of Braunschweig	Germany
TUD	Technical University of Darmstadt	Germany
MUNI	Masaryk University	Czech Republic
BGU	Ben-Gurion University	Israel
OsloMET	Oslo Metropolitan University	Norway
ICL	Imperial College London	UK
UMIL	University of Milan	Italy
BADW-LRZ	Leibniz Supercomputing Centre	Germany
EIT DIGITAL	EIT DIGITAL	Belgium
TELENOR	Telenor	Norway
ACS	Airbus Cybersecurity	Germany
SECT	secunet Security Networks	Germany
IFAG	Infineon	Germany
SIDN	SIDN	Netherlands
SNET	SurfNet	Netherlands
CYD	Cyber Detect	France
TID	Telefonica I+D	Spain
RD	RUAG Defence	Switzerland
BD	Bitdefender	Romania
ATOS	Atos Spain S.A.	Spain
SAG	Siemens	Germany
Flowmon	Flowmon Networks	Czech Republic
TÜV TRUST IT	TUV TRUST IT GmbH	Germany
TI	Telecom Italia	Italy
EFA	EFACEC	Portugal
ALBV	Arthur's Legal B.V.	Netherlands
EI	eesy innovation	Germany
DFN-CERT	DFN-CERT	Germany
CAIXA	CaixaBank	Spain
BMW	BMW	Germany
GSDP	Ministry of Digital Policy, Telecommunications and Media	Greece
RISE	RISE Research Institutes of Sweden AB	Sweden
Ericsson	Ericsson AB	Sweden
SBA	SBA Research gemeinnutzige GmbH	Austria
IJS	Institut Jozef Stefan	Slovenia
UiO	University of Oslo	Norway
ULANC	University of Lancaster	UK
ISI	Athena ISI	Greece
UNI PASSAU	University Passau	Germany
RUB	Ruhr University Bochum	Germany

Document Revisions and Quality Assurance

Internal Reviewers

1. Pavel Minařík (Flowmon)
2. Tanja Pavleska (IJS)

Revisions

Ver.	Date	By	Overview
0.1	2019-09-05	Jürgen Schönwälder (JUB)	Initial template
0.2	2019-12-05	Jürgen Schönwälder (JUB)	First internal review version
0.3	2019-12-16	Jürgen Schönwälder (JUB)	Second internal review version
1.0	2019-12-28	Jürgen Schönwälder (JUB)	Final version

Contents

Executive Summary

The research within CONCORDIA aims at building a European secure, resilient and trusted ecosystem for innovation in the area of cybersecurity. It brings together partners from academia and industry to stimulate collaboration and increase the impact of European research. The objectives for WP1 are: 1) to perform excellent academic research, 2) to organize scientific events, 3) to play a leading role in the organization of the scientific community and 4) to contribute to standardization, open data and code.

CONCORDIA's excellent academic research has resulted in more than 50 papers at workshops, conferences and journals, which exceeds the objective that is stated in the Description of Action. The fact that CONCORDIA research is inspired by real security problems is reflected by the fact that our top publications are presented at venues where security is *applied* within a specific context (venues like ACM-IMC, ACM-TCPS, IEEE-T-SUSC and IEEE-TNSM). Since the project has started only recently and given that publishing results of top research generally takes a long time, we expect that that the number of high quality papers will increase further in the next years.

Together with the three other EU pilot projects (ECHO, SPARTA and Cybersec-4Europe), CONCORDIA organizes the 4th International NeCS PhD Winter School (Trento) and the IFIP Summer School on Privacy and Identity Management (Brno). In 2019 CONCORDIA researchers contributed to the organization of nearly 25 scientific conferences, including ACM-IMC, and acted as Technical Program Committee member for more than 40 conferences. An interesting observation is that the most active persons in the organization of events within CONCORDIA seem to be women.

CONCORDIA researchers play also a leading role in the organization of the scientific community by acting as chairs of the IFIP Technical Committee 6 (Communication Systems) and IFIP Working Group 6.6 (Management of Network and Distributed Systems). They serve as editorial board members for 13 journals and transactions from ACM, IEEE, Springer and Wiley, as well as Steering Committee members for conferences such as ACM SACMAT, IFIP/IEEE IM and IFIP TMA.

In 2019, researchers from WP1 contributed to three Internet RFCs and one Internet draft. In addition, WP1 researchers are responsible for three open network data repositories: OpenINTEL, Trace-Share and Ethereum datasets. The biggest, OpenINTEL, captures daily snapshots of the state of large parts of the global Domain Name System (DNS). Every day 227 million domains are measured, resulting in 2.4 billion data points per day. Since the start of OpenINTEL in 2015, 3.6 trillion data points have been collected; such data allows us to detect various types of security issues, including the creation of SPAM and phishing domains, or the adaption of DDoS protection services. Note that the complete overview of all CONCORDIA

activities related to standardization and open data is provided in the deliverables of WP5 and WP6.

In general we may conclude that WP1 has reached, or with respect to the number of publications, even exceeded the objectives as described in the Description of Action. Collaboration has emerged between research in WP1 and the pilots in WP2 and WP3. Whereas the first year was used to explore collaboration opportunities for the whole spectrum of research activities, 2020 will emphasize those collaborations that have a high potential for impact. To increase focus, partners with PMs spread over multiple WP1 tasks will also be stimulated to concentrate these PMs in a single or smaller number of tasks.

1 Introduction

This document reports the research activities undertaken by the work package WP1 of the CONCORDIA project during the first year of the project. The goal of WP1 is to organize and coordinate the scientific research within the CONCORDIA project. WP1 has the following objectives:

- Excellent academic research to build an European Secured, Resilient and Trusted Networked Ecosystem, papers for scientific journals, conferences and workshops.
- Organization of scientific events in the area of cybersecurity, including a dedicated annual European cybersecurity conference.
- Leading role in the organization of the scientific community, outreach to different target audiences, including public media and the general public.
- Contributions to standardization, open research data and code, shared via systems like GitHub.

The main objective is to stimulate the publication of scientific results in key journals, conferences, and workshops in the broad field of cybersecurity. The SMART objective is to publish at least 100 of such papers during the project's lifetime.

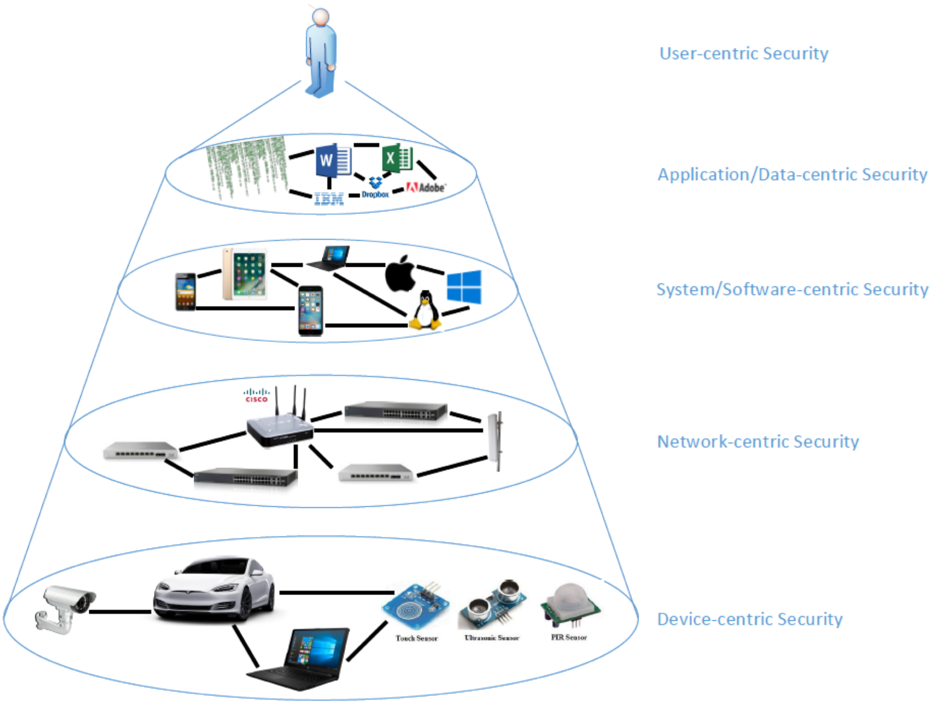

Figure 1: WP1 tasks

Research in WP1 is organized into five tasks, each one focusing on one particular aspect of cybersecurity (see Figure 1):

- T1.1: Device security aspects
- T1.2: Network security aspects
- T1.3: Software and system security aspects
- T1.4: Data and application security aspects
- T1.5: User security and privacy aspects

This deliverable is structured according to the objectives and the tasks. Section 2 summarizes some key achievements of the first year before the research activities of the different tasks are discussed in more detail. Section 3 provides an overview over the research related to device security and Section 4 reviews the research related to network security. Section 5 discusses research on software and system-centric security aspects while Section 6 focuses on data and application specific security aspects. Section 7, finally, discusses research related to user security and privacy. The organization of scientific events and the scientific community is summarized in Section 8. Section 9 discusses the contributions to standards and open research data, and this deliverable concludes in Section 10.

2 Key Achievements

The Description of Action defines the following objectives for WP1:

- Excellent academic research to build an European Secured, Resilient and Trusted Networked Ecosystem, papers for scientific journals, conferences and workshops.
- Organization of scientific events in the area of cybersecurity, including a dedicated annual European cybersecurity conference.
- Leading role in the organization of the scientific community, outreach to different target audiences, including public media and the general public.
- Contributions to standardization, open research data and code, shared via systems like GitHub.

Excellent Academic Research

An important objective of WP1 is to stimulate the publication of scientific results into key journals, conferences and workshops in the broad field of cybersecurity. The SMART objective is to publish at least 100 of such papers during the project's lifetime. As demonstrated in Appendix A, in CONCORDIA's first year researchers have already published around 50 papers in workshops, conferences and journals. Therefore we may conclude that the objective to create at least 100 paper after four years will most likely be exceeded.

An interesting observation is that the top CONCORDIA publications seem to focus on venues where security is *applied* within a specific context. For example, multiple papers have been published at the ACM Internet Measurement Conference (IMC), which is a venue that publishes excellent measurement research regarding (the security of) the Internet. Other CONCORDIA top publications appeared in the *ACM Transactions on Cyber-Physical Systems*, *IEEE Transactions on Sustainable Computing* and *IEEE Transactions on Network and Service Management*. This focus on *applied* or *usable security* is a strong indicator that CONCORDIA research is inspired by, and wants to have impact on, real-world security problems.

Since the project has started only recently and given that publishing results of top research generally takes a long time, we expect that that the number of papers published at top venues will further increase in the next years.

Organization of Scientific Events

Together with the three other EU pilot projects (ECHO, SPARTA and Cybersec-4Europe), CONCORDIA is co-organizing two PhD schools: the 4th International NeCS PhD Winter School and the IFIP Summer School on Privacy and Identity Management. The PhD Winter School takes place in Trento from 20-24 of January

2020; the IFIP Summer School takes place between August 17-21, 2020. Section 8.1 provides the details.

CONCORDIA researchers have contributed to the organization of nearly 25 scientific conferences (for details, see Section 8.2). One of the key events sponsored and organized by CONCORDIA was the ACM Internet Measurements Conference (IMC), which took place from 21-23 October 2019 in Amsterdam, the Netherlands. Although IMC has a broad scope on measurements, the most prominent category of research focuses on (Internet) security. An interesting observation is that the 28 chair positions are equally divided between men and women (14 each).

CONCORDIA researchers have also acted as Technical Program Committee (TPC) member for more than 40 conferences. Also here it is interesting to note that the most active persons within CONCORDIA seem to be women. For details, see Section 8.2.

Organization of the Scientific Community

Professional organizations play a key role in organizing the scientific community. CONCORDIA researchers therefore chair the IFIP Technical Committee 6, which focuses on Communication Systems and which is the largest TC within IFIP, as well as Working Group 6.6, which focuses on management of network and distributed systems. Other important activities to organize the scientific community include membership of editorial boards and steering committees. CONCORDIA researchers serve as editorial board members for 13 transactions and journals by publishers such as ACM, IEEE, Springer and Wiley, as well as steering committee members for conferences such as ACM SACMAT, IFIP/IEEE IM and IFIP TMA. For details, see Section 8.3.

Contributions to Standardization and Open Research Data

One of the objectives of WP1 is to contribute to standardization, open research data and code, shared via systems like GitHub. Note that not only *researchers* but also *industries* are active in these areas. Therefore the complete overview of all CONCORDIA activities related to standardization and open data is provided in the deliverables of WP5 and WP6.

In 2019, CONCORDIA researchers have contributed to three Internet RFCs and one Internet draft. In addition, WP1 researchers are responsible for the creation and maintenance of three open network data repositories: OpenINTEL, Trace-Share and Ethereum datasets. OpenINTEL captures daily snapshots of the state of large parts of the global Domain Name System (DNS). Because the DNS plays a key role in almost all Internet services, by analyzing DNS information we are able to detect various types of security issues, such as the creation of SPAM and phishing domains, or the adaption of DDoS protection services. By capturing DNS data and recording it over longer periods of time, OpenINTEL allows us to track the evolution of the DNS system. Details are provided in Section 9.

Contributions to the CONCORDIA Pilots

An important achievement in the first year of CONCORDIA is the collaboration that emerged between the research work package (WP1) and the pilots in WP2 and WP3. To bootstrap such collaboration, several meetings have been organized between researchers and the coordinators and members of the various pilots. The first scheduled discussions between researchers and pilots took place during the CONCORDIA kick-off meeting in January 2019. In May 2019 a dedicated meeting was organized in Bremen to exploit collaboration possibilities in detail. After that various forms on interactions emerged between the five WP1 tasks and the seven CONCORDIA pilots. Details of collaborations are described in the various task sections later in this document.

3 Device-Centric Security (T1.1)

Task 1.1 (T1.1) of the CONCORDIA project is concerned with device-centric security, with a special focus on IoT devices with limited resources. The task aims at

- developing techniques for detecting misbehaving IoT devices,
- analyzing automated software update mechanisms of IoT devices,
- investigating hardware components for post-quantum cryptography, and
- researching code analysis techniques to detect advanced persistent threats.

In the following sections, we report on the research activities performed within T1.1. Section 3.1 discusses how devices can be classified according to how easy it is to detect attack on them. A designated device taxonomy has been developed in order to identify which features impact attack detectability. Section 3.2 discusses device assurance, namely the process for the evaluation of the (security) state and trustworthiness of IoT devices. Device assurance is based on the collection of trustworthy evidence on device behavior. Blockchain technology is used to provide a verifiable and permanent record of collected evidence. Blockchains are also used as enabling technology in the research described in Section 3.3, where new techniques are developed to enforce individual privacy preferences in a distributed way directly on IoT devices. Section 3.4 investigates how blockchain technology can be integrated with IoT systems, with a special focus on low range low power wide area networks. The research investigates how light-weight cryptographic mechanisms can help secure access to devices with very limited computational resources. Section 3.5 describes work on assessing the security of software running in a trusted execution environment, such as ARM's TrustZone Technology for Cortex-M processors. Section 3.6 summarizes research on hardware security mechanisms that prevent side-channel attacks or provide light-weight cryptographic algorithms suitable for resource-constrained embedded devices. Section 3.7 summarizes work towards a trust establishment and computation system for vehicular mobile ad-hoc networks. Section 3.8 describes how interoperability between IoT systems can be achieved using so called semantic mediators, that achieve interoperability at the semantic level. The work described in Section 3.9 extends this work by linking semantic mediators into a layered architecture for semantic interoperability mechanisms.

Figure 2 indicates how research done in T1.1 relates to the CONCORDIA pilots.

1. The research by BGU on applying machine learning and artificial intelligence techniques for the efficient prevention and detection of malware and the identification of IoT devices behind network address translators is highly relevant for Ericsson in the context of 5G networks and their work in the Telecom pilot (T2.1).

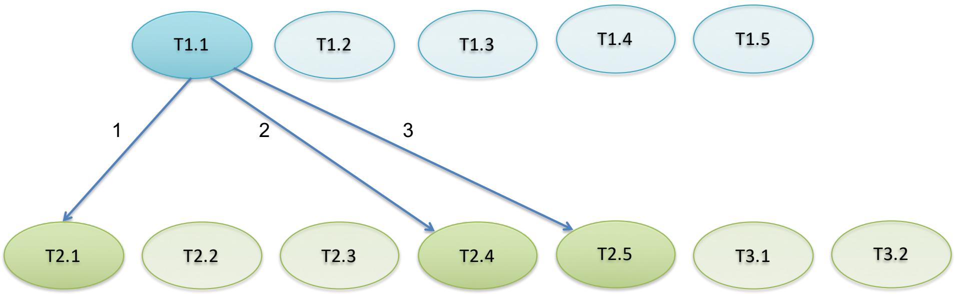

Figure 2: Links between task T1.1 and the CONCORDIA pilots

2. Several partners in T1.1 research topics related to trusted execution environments and more specifically ARM's TrustZone technology. This research is relevant for the eHealth pilot (T2.4). The partner EI is acting as a bridge between T1.1 and T2.4 and planning to integrate some T1.1 research results into the T2.4 prototypes.

3. The unmanned aerial systems pilot (T2.5) is looking for efficient, secure, and lightweight technologies that can be used to identify unmanned aerial systems. The post-quantum crypto work in T1.1 is in particular of relevance here. There are concrete talks between the ACS (leading the T2.5 pilot) and UP and ISI how their technology can be integrated into the pilot.

3.1 Quantification of IoT Attack Detectability

Contact: Yair Meidan (BGU)

The Internet of Things (IoT) is a rapidly evolving paradigm in wireless communications, where various things are being connected to the Internet. Internet enabled things can also cooperate with one another to reach common goals. For the home or office user, such things may include smart webcams, smoke detectors, lightbulbs, refrigerators, etc. Despite their potential benefit in many aspects of modern life, IoT devices are unfortunately also trending as "easy targets" for various cyber attacks, including DDoS, ransomware, data ex-filtration and cryptocurrency mining.

Many studies have explored the cybersecurity aspects of IoT devices. In some of them, network traffic data of real-world IoT devices has been collected to train machine learning classifiers for traffic monitoring and identification of exploitation attempts and attack execution. For instance, the researchers in [82] deployed nine IoT devices from a variety of device makes and types (smart cameras, doorbells, a thermostat and a baby monitor), and collected their benign data for approximately a week. Then they compromised these IoT devices using Mirai and BASHLITE botnets, such that they would send spam and execute network scanning and flooding attacks. Having collected and labeled the resultant malicious traffic data they

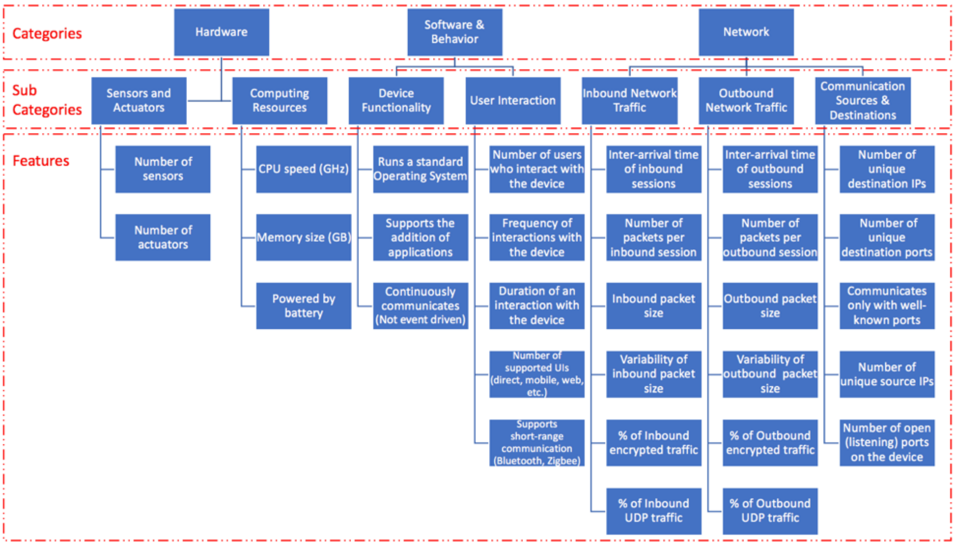
Categories
Hardware
Software & Behavior
Network
Sub Categories
Sensors and Actuators
Computing Resources
Device Functionality
User Interaction
Inbound Network Traffic
Outbound Network Traffic
Communication Sources & Destinations
Features
Number of sensors
Number of actuators
CPU speed (GHz)
Memory size (GB)
Powered by battery
Runs a standard Operating System
Supports the addition of applications
Continuously communicates (Not event driven)
Number of users who interact with the device
Frequency of interactions with the device
Duration of an interaction with the device
Number of supported UIs (direct, mobile, web, etc.)
Supports short-range communication (Bluetooth, Zigbee)
Inter-arrival time of inbound sessions
Number of packets per inbound session
Inbound packet size
Variability of inbound packet size
% of Inbound encrypted traffic
% of Inbound UDP traffic
Inter-arrival time of outbound sessions
Number of packets per outbound session
Outbound packet size
Variability of outbound packet size
% of Outbound encrypted traffic
% of Outbound UDP traffic
Number of unique destination IPs
Number of unique destination ports
Communicates only with well-known ports
Number of unique source IPs
Number of open (listening) ports on the device

Figure 3: Taxonomy of IoT device characteristics organized into categories, subcategories, and features

employed four anomaly detection algorithms in order to detect the attacks. A key finding in this research was that although the False Positive Rate (FPR) was low for most algorithms and IoT devices, some variability in FPR still prevailed, as well as in the detection time.

In the current research, the focus is set on modeling the variability in IoT attack detection metrics. That is, given (a) an attack scenario and (b) the characteristics of an IoT device model to be compromised by this attack, we wish to estimate: How easy would it be to detect this attack on this device model using network traffic analysis? In the first phase of the research we take an expert-based approach, where the characteristics of IoT devices are being ranked by cybersecurity experts based on their potential predictive power (i.e., which feature is more informative?). This ranking will help us in two ways:

1. We will be able to identify the most predictive / informative features for IoT attack detection, thus enable dimensionality reduction for Network Intrusion Detection Systems (NIDSs).

2. We will be able to calculate a detectability score for each combination of attack and device, thereby enabling network security officials to devise detectability-based organizational policies. For example, an IoT device whose typical network traffic behavior is likely to "hide" certain attack scenarios will not be allowed to connect to the organizational network.

Several achievements have been accomplished in this research during 2019:

- We have decomposed various IoT characteristics into a novel hierarchical taxonomy consisting of three categories, seven sub-categories and 30 features (see Fig. 3).
- We designed and implemented an online questionnaire[1] to facilitate the collection of inputs from multiple cybersecurity experts at BGU and from members of the CONCORDIA project.
- We designed a procedure based on the analytic hierarchy process (AHP) to analyze the data and meta-data of the inputs.
- We started writing a research manuscript, to be finalized once enough inputs have been collected and analyzed. This manuscript will be submitted for review before the end of 2019.

This research activity benefits from the genuine collaboration among European cybersecurity experts involved in the project. Knowledge from numerous academic and industrial organizations, spanning over various sub-domains of cybersecurity, is being collected and processed as part of an expert-based IoT security assessment method.

3.2 Security Assurance for IoT Systems

Contact: Claudio A. Ardagna (UMIL)

The existence of billions of cheap and resource-constrained devices connected to the Internet introduces fundamental risks that can threaten users' life and personal sphere. A wealth of services in different domains, such as smart vehicles, smart buildings, e-health, are distributed on the basis of data collected by devices. In this context, assurance evaluation is fundamental to guarantee the correct behavior of the whole system and its devices. The term assurance means, in a wider sense, the technical judgment that a service, process, or device satisfies some properties. The implicit assumption is that data have a sufficient level of trustworthiness to create information, and in turn knowledge and wisdom. This assumption is however not sound when a plethora of devices are used to collect data, and might lead to scenarios where wrong evidence results in wrong decisions and, in turn, untrusted services/applications. It is likely that without an open, protocol-neutral baseline solution for IoT assurance, fundamental risks will create further exploitation opportunities. Research on IoT-based systems assurance is, however, at an early stage, and mainly focused on defining new assurance architectures for IoT. Traditional assurance techniques are affected by important limitations when targeting complex IoT systems as follows:

[1] https://dscore.limequery.com/915153

1. Assurance techniques do not target hybrid systems, where cloud systems at the center are connected via edge networks to smart devices at the periphery and no clear perimeters exist.
2. Hybrid systems rely on data continuously collected by a multitude of devices, which are intrinsically unreliable and under the control of many untrusted providers.
3. Traditional assurance is often driven by untrusted/unverified evidence that is accepted on the basis of the provider reputation.

The need of collecting trustworthy evidence clearly emerges in the above challenges. This need has initially been dealt with in the context of forensics science by defining a systematic and reliable methodology for evidence collection and analysis. Some solutions based on blockchain have been also proposed to guarantee availability, integrity, and verifiability of collected evidence. The research is currently focusing on the need of providing trustworthy evidence collection as the basis for implementing a trustworthy IoT environment, where trustworthy processes and decisions are employed. Data collected from each smart device must be first evaluated and then put into service only if a minimum amount of assurance requirements are addressed.

Our work in the first year of the project focused on the need of trustworthy data at the basis of provable automation and adaptation processes [12]. The idea is to provide an approach that complements existing trustworthy decision processes with a methodology for collecting trustworthy data. The intuition behind our approach is that the more trustworthy data, the higher the decision accuracy. We then provided a novel, service-based methodology for trustworthy evidence collection on the basis of a trustworthy assurance evaluation of IoT processes and systems. Our methodology is implemented with varying granularity, depending on specific performance requirements, from simple trustworthy evidence collection, to trustworthy evidence aggregation, and provable evidence-based automation. It is based on blockchain and smart contracts to guarantee collection of reliable evidence whose integrity is proven over time. Differently from existing solutions, our methodology links the evidence to the way in which it is collected and aggregated, and then a decision based on it is taken. The methodology balances between the given trustworthiness level trustworthiness and its performance, and it is experimentally evaluated using Hyperledger Fabric blockchain.

A methodology and an approach based on blockchain and smart contracts for trustworthy evidence collection are fundamental for all scenarios and domains where automated and adaptive processes are employed. The decisions taken based on untrustworthy data may have detrimental consequences on both systems, users, and even people. Safety and security of systems and their users can only be maintained when decisions are taken on trustworthy and reliable data on their behavior. These activities, carried out mainly in Task 1.1, are also relevant for activities in Task 1.5

aiming to protect final users. Furthermore, they are important for all CONCORDIA pilots, especially, the pilot "e-Health Sector: Privacy and Data Protection", where trustworthy data collection needs to balance privacy, data traceability and data loss, and the pilot "Security of Unmanned Aerial Systems (UAS)", where the security of unmanned aerial systems must be protected and passes from the quality of collected data used, for instance, for drone authentication.

3.3 Decentralized Enforcement of Individual Privacy Preferences in IoT Environments via Blockchains

Contact: Barbara Carminati (UI)

Internet of Things (IoT) technologies are revolutionizing our daily lives, building around us a pervasive environment of smart objects able, not only to sense data but also to interact with other objects and to aggregate data sensed through different sensors. Smart objects are now able to create new knowledge locally, that could be used to make decisions, such as quickly trigger actions on environments, if needed. Such a scenario enacts the transition from the Internet of Things to the Internet of Everything, a new definition of IoT seen as a loosely coupled, decentralized system of cooperating smart objects, which leverages on alternative architectural patterns with respect to the centralized cloud-based one, such as fog computing. Such a trend towards decentralization reduces the amount of data that is transferred to the cloud for processing and analysis and can also be instrumental in improving data security and individual privacy, two significant concern in the IoT scenario. However, decentralization, if not adequately governed, might also imply loss of control over the data, with consequences on individual privacy. To cope with this problem, as a first step, in [26] we have addressed the problem of specifying and enforcing individual's privacy preferences in the IoT scenario. At this aim, we have defined a privacy model according to which an individual can state a set of conditions on the usage of his sensed data by service providers (aka, conditions on the purpose for which the data can be sensed/used, the maximum retention time of collected data, how/whether data sensed by different devices can be aggregate together etc.). Moreover, in [26] we have also addressed the problem of associating new privacy preferences to newly derived data, that is, data generated as results of processing sensed data. For the enforcement mechanism, in [26] we have considered a centralized architecture, that is, a scenario where devices have only the capability to sense data and send them to a data center for further analysis. As such, the proposed enforcement monitor statically analyzes every consumer query and decides if the privacy policy of the service provider satisfies the privacy preferences specified by owners of devices generating the data.

As a second step, in [120], we focused on the a decentralized privacy enforcement mechanism, where the compliance check of user individual privacy preferences is performed directly by smart objects, rather than by a central entity. To address these challenges, in [120], we extend the privacy preference model proposed in [26], by designing a set of privacy meta-data that are used by smart objects for locally checking privacy preferences and for locally enforcing user privacy preferences at smart object level. Smart objects are thus able to derive privacy meta-data for newly created data items, keep track of the operations performed over data items, denoted as history, in order to ease privacy preference enforcement, and, finally, check compliance of the privacy policy of the data consumer with the privacy preferences of data items.

Embedding the enforcement mechanism into smart objects implies to rely on trustworthiness of these devices. In this first year of the CONCORDIA project, we have investigated how to relax this assumption. The key idea is exploiting a blockchain platform for the enforcement of owners privacy preferences. In general, a blockchain is a distributed data structure, replicated and shared among members of a network, acting as a distributed ledger, used to keep track of every exchange of resources or assets between participants of a network. These changes are recorded into transactions, batched into time-stamped linked blocks, forming the so-called chain of blocks. Transactions are inserted into blocks only if they are considered valid by the network participants. Transaction validation is reached through a distributed consensus protocol that, in general, is considered secure if the majority of network participants are honest. An important aspect, when leveraging on a blockchain, is that the computation involved in transaction validation can be encoded into predefined programs, called smart contracts. As such, blockchain can be seen as a distributed ledger storing results (i.e., transactions) of (smart) contracts whose correct evaluation have been validated by network participants.

In light of this, the scenario we are considering is the one of a smart environment, that is, a set of connected IoT devices, owned by a given user, able to sense data, eventually locally elaborate them, and then send them to the provider servers. We further assume that the IoT network is connected to provider servers via a limited set of special IoT devices, called gateways. In this view, we are investigating a framework where users' gateways are registered in the blockchain. This implies that, in addition to be instructed to interact with their IoT networks, gateways act also as Blockchain nodes. Then, in order to exploit the blockchain for the enforcement of owner privacy preferences, we are investigating how to encode the privacy compliance checks into smart contracts. This brings the nice benefit of not requiring to trust the single smart object enforcing the privacy compliance check.

3.4 Scalable Transport Mechanisms for Blockchain IoT Applications

Contact: Eryk Schiller, Sina Rafati, Burkhard Stiller (UZH)

Internet of Things (IoT)-integrated use cases have raised a high attention in the past decade [45], as supply chain monitoring, environmental monitoring, smart cities, smart industries, and health-care focus on data immutability and require IoT systems for measurements, data collection, and active control. Thus, the integration

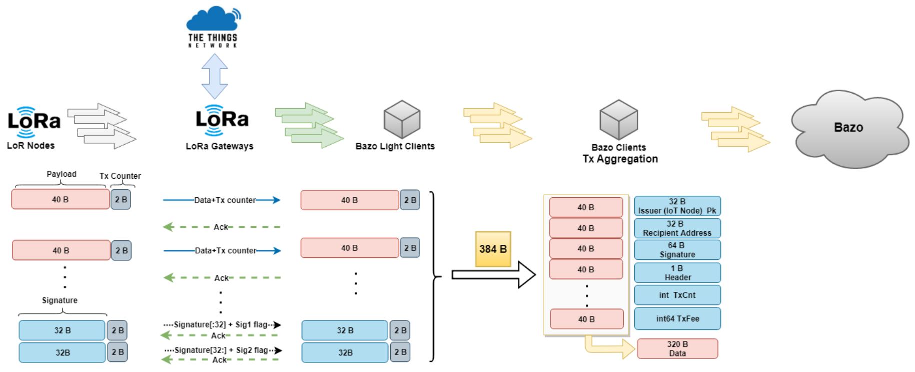

Figure 4: Transmission scheme used by the prototype system. IoT devices (i.e., LoRa nodes) use LoRa adapters to communicate with the TTN network. The TTN provides a mean for the e2e connectivity through LoRa Gateways securing a connection between IoT devices and Bazo Clients, which in turn submit blockchain transactions to the Bazo blockchain. First, several data samples are reported. Second, previously delivered data packets are signed using one Ed25519 ECDSA signature of 64 bytes. A reliable communication scheme in the LoRa network is materialized through sequential tx counters, Ack messages, as well as the ARQ scheme retransmitting lost packets (i.e., data chunks and signatures).

of Blockchains (BCs) and IoT into Blockchain-IoT (BIoT)-supportive applications responds to demands of persistent storage of strongly secured data, where automated data collection becomes a key for offering transparency and reliability. A highly demanding role of BIoT applications requires an elaboration and analysis of underlying IoT protocols, which form the communication basis for IoT systems. Thus, the studies on the range of communication, data rates, maximum transmission units (MTUs), reliability of communication protocol, and the energy efficiency are required to appropriately support IoT deployments [45]. To understand the requirements of BIoT systems, prototype BIoT applications were developed and analyzed in the context of Long Range (LoRa) Low Power Wide Area Networks (LPWAN), which is a highly recognized technology used in IoT data streaming.

BIoT systems require research attention due to high technical potential in various application domains such as micro payments, smart-cities, or smart-grids. This work studies various methods that improve the performance of Blockchain systems integrated with the Internet of Things (BIoT) using the LoRaWAN access method. Duty Cycle Enforcement (DCE) and Listen Before Talk (LBT) mechanisms as the LoRa channel access methods, Automatic Repeat reQuest (ARQ) on the LoRa Transport Layer, and transaction aggregation on the Application Layer were evaluated. The main focus was put on the system performance studying the maximal number of transactions submitted, reliability of transport schemes, and the energy efficiency of the BIoT system.

In a first step, the characteristics of IoT traffic in The Things Network (TTN) were evaluated, where regular traffic with certain periodicities ranging between 0-300 hours was discovered in the majority of devices connected to TTN. Second, the performance of Edwards-curve Digital Signature Algorithm (EdDSA), i.e., Ed25519 [16] on Class 1 [24] devices was evaluated and established at an acceptable level. The Arduino Device AT2560 with an 8-bit AVR RISC-based micro-controller having 8 kB of SRAM using the Arduino Cryptographic Library (ACL) [1] needs around 6 seconds to either (i) generate a public/private key-pair or (ii) sign a variable data length transaction (TX) for payloads between 0 and 2000 bytes (see Fig. 4). Third, the work performance proves that the combination of LBT-based LoRa MAC, the ARQ-enabled LoRa Transport Layer, and transaction aggregation at the Application Layer provides a good trade-off between the submitted transaction count, packet loss, and energy efficiency. The proposed scheme complies to the data integrity demands of BIoT applications by specifying a reliable end-to-end (e2e) data transmission scheme from IoT end-devices to the BC. The elimination of the heavy EdDSA processing upon every data packet in *multi-packet transactions* provided much better performance in terms of the capacity of submitted data streams and energy efficiency for *multi-packet transaction* schemes. The paper was presented at IEEE LCN 2019 [125], and a demo was given at IEEE ICBC 2019 [112].

BC based systems can improve data transparency and reliability of IoT-BC e2e applications providing indeniability and immutability of the information stored in the BC. Moreover, this work studies different cryptographic primitives that can be deployed directly on IoT devices to improve the security of the overall system.

This work relates to T1.2 (Network-Centric Security) through the analysis of data traffic in the TTN network, where the periodic traffic was discovered among the majority devices. This can help to build traffic classifier and discover traffic anomalies in the TTN network. Finally, the link to T1.4 (Data/Application-Centric Security) was established by originated BC TXs directly from IoT devices, which requires the deployment of security algorithms (e.g., Ed25519) directly on IoT devices. Furthermore, BIoT system can allow for BC micro-payments in the case T2.3 (TransportE-Mobility Sector: Security for the e-Charging Infrastructure) materializing vehicle charging without any interaction from the driver side, i.e., the IoT infrastructure on the vehicle will perform all the tasks required to complete vehicle charging. BIoT systems can also have a deep impact on T2.4 (e-Health Sector: Privacy and Data Protection) providing data reliability, indeniability, and immutability through the BC-enabled storage of TX originated at IoT devices. Finally, the research on BIoT systems links to T2.5 (Security of Unmanned Aerial Systems (UAS)), while BC-enabled network can enhance auditing in defense usecases.

3.5 Security Assessment of TrustZone-M based Software

Contact: Antonio Ken Iannillo (SnT)

Trusted hardware technologies are commonly used as anti-tamper technologies to make software more resistant against attack and protect critical program elements. It is generally more difficult to successfully attack trusted hardware than a software-only protection scheme. With the advent of the Internet of Things (IoT), computing and networking capabilities are extending to devices that are not considered as computers, enabling them to interact with the physical world or other software entities with minimal or no human input. These devices are powered by embedded computers: small hardware (micro-controllers) equipped with specialized sensors and actuators that run a constrained software to handle data and external communication. Micro-controllers processors have much more limitations than application processors. Indeed, the main requirements for micro-controller applications are low power consumption, real-time processing, deterministic behavior, and low interrupt latency. Thus, hardware security extensions for application processors cannot be directly applied, because they have been developed for more relaxed use cases. However, we strongly require IoT devices to be secure. We are inclined to exploit the machines because they can be ubiquitous and do not have the fallibility that humans do, but we can't sufficiently trust them because they are vulnerable, and failures in cyber-physical systems can cause actual harm. As extreme examples, changing the mix ratio of disinfectants at a water treatment plant or stopping the cooling system at a nuclear power plant could potentially place a whole city in immediate danger.

Recently, ARM Holding, which already owns the largest share of mobile and embedded markets (60%), has further extended TrustZone-support for the tiniest low-end devices. ARM designed a hardware security extension from the ground up, instead of reusing it from application processors, for micro-controllers with the name of TrustZone Technology for Cortex-M profile or TrustZone-M. There are several security applications that can be enabled by TrustZone-M, for example: IoT device makers can use it to store intellectual property in secure memory while still allowing non-secure application to access it via APIs; secure storage of critical information; a root of trust implementation provides a secure foundation for over-the-air (OTA) firmware updates and mutual authentication between devices in a system. The main research idea is to create tools and frameworks that can help practitioners to assess the security properties of the software based on this technology.

A methodology has been proposed in 2019 for the security assessment of software based on TrustZone-M [61], the ARM hardware security extension for micro-controllers. The methodology consists of the exploitation of a verification and validation framework to automatically test TrustZone-M based software.

Assessing the security of IoT devices is of paramount importance for all the project related research that provides new solutions based on IoT devices (T1.1 research). In particular, there is a direct link with the activities carried in 5.6 where this research can be directly applied since it is based on the same technology (ARM TrustZone-M).

Furthermore, this research is very relevant for the e-health pilot and the UAV pilot. In the first case, the e-health pilot, and in particular eesy-innovation, is using ARM processor for the connected medical devices. This pilot will eventually integrate TrustZone-M enabled devices and this research may guarantee their security. In the second case, Airbus is leading the implementation of a use case for UAVs that consists also of authentication and authorizations between devices. The security of such communications can be guaranteed by ARM TrustZone-M devices and assessed by the outputs of this research.

3.6 Hardware Cybersecurity for Embedded Systems and IoT Devices

Contact: Odysseas Koufopavlou (UP), Apostolos P. Fournaris (ISI)

The University of Patras (UP) and the Industrial Systems Institute (ISI) of the Athena Research Center are working on providing hardware assisted/protected cybersecurity solutions to enhance embedded systems and IoT device security. Embedded systems and devices are gradually gaining a considerable market and technology share of computational system devices in various domains. Some of those domains, like Critical Infrastructure (CI) environments, have a need for fast responsiveness thus requiring real-time embedded system but in parallel, they also have a high demand for strong security. Cybersecurity attackers can find fertile ground in the embedded system domain since such systems' original structure was not meant to provide security but rather safety and in several occasions real time response.

Embedded systems, as most computing units, cannot be considered trusted due to many known vulnerabilities. Typically, software tools (e.g., antivirus, antimalware, firewalls) can protect a system from attackers taking advantage of those vulnerabilities to inject some malicious software. However, there exist security gaps and vulnerabilities that cannot be traced by the above-mentioned software tools since they are not stationed on software applications, drivers or the operating system but rather on the computer architecture and hardware structure itself. Creating software to exploit these vulnerabilities remains hidden from most software cybersecurity tools and thus constitute a serious security risk for all devices using these commodity, vulnerable, computer structures.

One of the most effective means of enhancing computer security without compromises in quality of service (i.e., real-time responsiveness, computation and network performance, power consumption etc.) is to migrate security functions in hardware. Thus, a hardware/software co-designed security design approach has been promoted in applications where strong security is needed. Under this framework, a security-by-design approach should be enforced in security enabled devices that promotes a notion of trust in all the device's architectural layers from hardware to firmware, to operating system kernel and to software applications. For this reason, there exist many hardware security elements along with appropriate software libraries that advertise security and trust.

However, security and trust on hardware cybersecurity approaches is strongly associated with the protection against implementation attacks since any real-time embedded system device can leak sensitive information to a knowledgeable attacker. These, side channel attacks (SCAs) aim at collecting leakage data during processing and through statistical computations extract sensitive information. SCAs can be very potent to embedded system devices left unwatched for long periods of time so that an attacker can have physical access or proximity to them. Characteristic examples of such unattended devices are critical infrastructure systems' end nodes, e.g., embedded system devices like sensors, in-field actuators, programmable logic controller (PLC), or other monitoring and control devices. However, even if physical access to devices is not possible, there exist SCA types, like micro-architectural/cache SCAs, that can be mounted remotely.

Under the light of the above dangers, focused on a device's architecture and hardware structure, embedded system devices that are traditionally designed and built mainly for high QoS (reliability, high response time, high performance) and real-time responsiveness are left unprotected.

In CONCORDIA, considering the above described research background, we perform applied research on dedicated security devices, like hardware security modules, for the embedded system world in order to instill a security and trust-by-design approach. Our goal is to propose self-contained security and trust tokens that will be able to support next generation cryptography services like postquantum cryptography primitives for achieving confidentiality, authentication, availability and integrity while in parallel can be highly protected/trusted against side channel and micro-architectural attacks. Our research is focused on the following research directives:

- Hardware Cybersecurity: In the field of number theory arithmetic and cryptographic engineering, our expertise constitutes of the design and implementation of Hardware/Software Cryptographic components, focusing on HW/SW co-design solutions and with a future scope at the emerging necessity that is Post-Quantum Cryptography. Goal is the creation of Hardware assisted Cybersecurity components, like hardware security modules.

- Secure implementation: Focus is also being given to the secure implementation of these cryptographic components with the use of appropriate countermeasures, in order to instill trust against side channel analysis (differential power analysis or electromagnetic emission attacks), and/or fault injection analysis attacks. Additionally, software vulnerabilities on modern embedded systems are assessed, through the form of micro-architectural attacks (remote cache attacks) and remote fault injection (rowhammer based) attacks.

- Embedded Systems IoT Security: Real-world use of end-to-end security (TLS/DTLS design) for a wide range of IoT Systems, like wireless sensor networks, industrial systems node security etc.

During 2019, UP and ISI have been collaborating on research on fault injection attack resistance against vulnerability that exists in embedded system DRAM cells, on providing lightweight cryptography hardware based cryptography cores for resource constrained environments, and on designing embedded system hardware security sensors (in the form of hardware security tokens) for secure, authenticated collection of intrusion detection logs and events.

- As the cell density of DRAM modules keeps increasing to guarantee the growing demand for memory capacity in modern computer and embedded system, the electromagnetic interference between memory cells increase significantly, leading to security vulnerabilities that can be exploited. In particular, the widespread Rowhammer vulnerability has been proven to exist in the most common DRAM modules manufactured after 2012. In our paper [48] we are revisiting the exploitation approaches using the Rowhammer bug and we are providing an overall study of their security implications on Embedded Systems such as mobiles or tablets based on ARM architecture. Furthermore, we present our implementation approach on the Phys Feng Shui methodological attack for Android on a LG Nexus 5 device and we highlight the practical issues that arise when trying to trigger the Rowhammer attack on a real system.

- More and more devices, from the most high-end smartphones, to the most modest environmental sensors, already exhibit or acquire the ability to communicate with each other and be parts of networks. Connectivity is mainly accomplished through wireless communications and privacy is the way to go, regarding security over the air. A common approach dictates the incorporation of a crypto-core inside the transceiver. However, a privacy implementation must take into account several restrictions, imposed by the constrained resources of a small, power efficient device. For this purpose, an incorporation of a hardware implemented lightweight Simeck32/64 block cipher for IEEE 802.15.4 transceiver was proposed. In our work [77], we examine the FPGA implementation cost of such an approach, and we propose newly designed and efficient implementations of the proposed Simeck32/64 Crypto-Core. Our implementations achieve great improvement in throughput values, with a reasonable increase in hardware area. Furthermore, modifications of these approaches, exploiting the special functional blocks of a Xilinx FPGA, namely BRAM and DSP blocks, are implemented and the achieved results are evaluated. Additionally, the effects of applying the DSP multi-pumping methodology on the aforementioned implementations are also examined.

- Anomaly detection systems (ADS), as part of a security information and event management (SIEM) system, are cybersecurity tools for identifying potential threats inside an information technology system. They are widely used in critical infrastructure (CI) systems for protection against attacks that can cause severe problems to public security and welfare. ADS collect information from various kinds of sources and correlate them to identify anomaly events. Such sources can be devices and software sensors which inside a CI context (factories, power plants, remote locations) are placed in open areas and left unattended. These devices are vulnerable to tampering and malicious manipulation which may then lead an ADS or SIEM system to ignore or falsely alert of possible cybersecurity problems. In our work, we describe strategies to mitigate the above problem using hardware means in order to enhance trust on ADS sensors. Furthermore we propose a hardware/software based approach for legacy CI devices that can act as an ADS sensor or a tool for ensuring software ADS sensor data are not tampered.

The above described activities are a close match of the activities of T1.1 Device-centric Security. In the task description it is explicitly mentioned that ”we will also investigate the hardware cybersecurity structures of embedded IoT nodes with a focus on the latest (e.g., authenticated encryption, lightweight encryption) and upcoming (e.g., post-quantum) cryptography implementation solutions.”. Thus, in year one, the bases of the hardware cybersecurity research effort as this is dictated by T1.1 description are set in order to progress with more advanced hardware security approaches in the remaining years of the project (e.g., secure postquantum design, side channel attack assessment using machine learning algorithms etc).

The first year’s research outcomes (in the form of publications) or research activities can be an important tool in securing exposed or even legacy devices used in the CONCORDIA pilots, and especially those not offering extensive security configuration capabilities. In particular, pilots that use embedded systems in their core (e.g., T2.5 pilot on unmanned aerial vehicles) can benefit from the research effort on hardware cybersecurity solutions since the presented research activities can create a secure environment for embedded system operations. For example, authentication can be provided in embedded systems with an added feature of security against side channel analysis attacks on cryptographic algorithms. In the remaining years of the project, our focus will be shifted to next generation cryptography protocols (including post-quantum solutions), staying ahead of the curve on an even faster approaching post-quantum reality.

3.7 MobileTrust: Secure Knowledge Integration in VANETs

Contact: Sotiris Ioannidis (FORTH)

Vehicular Ad hoc NETworks (VANET) are becoming popular due to the emergence of the Internet of Things and ambient intelligence applications. In such networks, secure resource sharing functionality is accomplished by incorporating trust schemes. Current solutions adopt peer-to-peer technologies that can cover the large operational area. However, these systems fail to capture some inherent properties of VANETs, such as fast and ephemeral interaction, making robust trust evaluation of crowd-sourcing challenging. In this work, we propose MobileTrust – a hybrid trust-based system for secure resource sharing in VANETs. The proposal is a breakthrough in centralized trust computing that utilizes cloud and upcoming 5G technologies in order to provide robust trust establishment with global scalability. The ad hoc communication is energy-efficient and protects the system against threats that are not countered by the current settings. To evaluate its performance and effectiveness, MobileTrust is modeled in the SUMO simulator and tested on the traffic features of the small-size German city of Eichstatt. Similar schemes are implemented in the same platform in order to provide a fair comparison. Moreover, MobileTrust is deployed on a typical embedded system platform and applied on a real smart car installation for monitoring traffic and road-state parameters of an urban application.

The proposed system is developed in such a way to provide security, privacy, and trust in an intelligent and energy-aware transportation scenario, bringing closer the vision of sustainable circular economy.

During 2019 we focused on the design and implementation of MobileTrust – a hybrid trust-based system for secure resource sharing in VANETs and published our work on ACM Transactions on Cyber-Physical Systems [58]. Specifically, we study the MobileTrust design as well as the knowledge integration and we implement the centralized trust management component. Next the trust frameworks need to be implemented. For every user, MobileTrust evaluates their trustworthiness based on past interactions. All new users start with neutral trust and reputation values, which are continuously updated as new pieces of knowledge are integrated. Next, in order to maintain the reputation records and operate in a trustworthy manner, the users that contribute information to MobileTrust must be authenticated, so the authentication and privacy scheme is implemented.

The vehicle registration step enforces a procedure to examine that the user is driving a real vehicle. Otherwise, a malicious entity could create several user accounts for virtual vehicles, which would later launch coordinated sophisticated attacks on-line by sending fake positioning information along with erroneous road events.

The next step is to deploy the theoretical analysis of the proposed approach and its effectiveness in countering the attacker models that are detailed in the simulation study. This analysis uses mainstream cryptographic primitives (i.e., TLS, RSA, AES, and SHA512) that need to be implemented and configured properly. Consequently, their usage is considered theoretically secure. At first, we examine the security aspects that are provided by the deployed communication protocols. Then, we give proofs regarding the capabilities offering MobileTrust in supporting the legitimate functionality in presence of attackers. This step initially includes the protocol analysis, where we ensure that the analysis of the communication links between the involved entities and the MobileTrust component is modeled in the verification tool ProVerif [23]. It is a widely-used automatic symbolic protocol verifier that proves the security properties of the examined protocol, like authentication, secrecy, and adversary equivalence aspects.

This work focuses on the centralized trust computing using cloud and upcoming 5G technologies in order to provide robust trust establishment with global scalability. It presents a trust-based resource sharing system for Vehicular Ad hoc NETworks that builds upon a main security functionality (i.e., IEEE 802.11p, ETSI ITS G5) and the goal is to evaluate the entities' behavior once the cryptographically secure communication is achieved. Consequently, this study is related to CONCORDIA pilot Task T2.3 of transport and e-mobility sector.

3.8 Secure Semantic Interoperability for IoT Applications with Linked Data

Contact: Sotiris Ioannidis (FORTH)

Interoperability stands for the capacity of a system to interact with the units of another entity. Although it is quite easy to accomplish this within the products of the same brand, it is not facile to provide compatibility for the whole spectrum of the Internet-of-Things (IoT) and the Linked Data (LD) world. Currently, the different applications and devices operate in their own cloud/platform, without supporting sufficient interaction with different vendor-products. As it concerns the meaning of data, which is the main focus of this study, semantics can settle commonly agreed information models and ontologies for the used terms. However, as there are several ontologies for describing each distinct 'Thing', we need Semantic Mediators (SMs) in order to perform common data mapping across the various utilized formats (i.e., XML or JSON) and ontology alignment (e.g., resolve conflicts).Our goal is to enable end-to-end vertical compatibility and horizontal cooperation at all levels (field/network/backend). Moreover, the implication of security must be taken into consideration as the unsafe adoption of semantic technologies exposes the linking data and the user's privacy, issues that are neglected by the majority of the semantic-web studies. A motivating example of smart sensing is described along with a preliminary implementation on real heterogeneous devices. Two different IoT platforms are integrating in the case study, detailing the main SM features. The proposed setting is secure, scalable, and the overall overhead is sufficient for runtime operation, while providing significant advances over state-of-the-art solutions.

Specifically, semantic interoperability is the designed property where various systems can interact with each other and exchange data with unambiguous, shared meaning. This enables knowledge discovery, machine computable reasoning, and federation of different information systems. Interoperability is materialized by including information regarding the data (metadata) and linking each element to a commonly shared vocabulary. Thus, the meaning of the data is exchanged along the data itself in a self-describing information package. The shared vocabulary and the associations to an ontology enable machine interoperation, logic, and inference. Ontology is the explicit specification of a conceptualization and includes a formal representation of the properties and relations between the entities, concepts and data of a specific application domain.

In general, technologies from the Semantic Web are adapted in order to capture the inherited properties of an IoT ecosystem. They are mainly eXtensible Markup Language (XML) schemes, such as the Resource Description Framework (RDF), RDF Scheme (RDFS), and Web Ontology Language (OWL) for ontologies, and for services the Web Services Description Language (WSDL). These primitives provide common definitions of data or services, describe things with the underlying properties, and accommodate the semantic annotations, discovery of resources, inference of knowledge, and access control, in an interoperable and machine-readable fashion.

The common format and meaning of semantics in a universally accepted ontology, as suggested above, would be fruitful. Yet, this is not the current status [54]. While various systems could employ standardized or popular ontologies, eventually they extend them and settle own interfaces and semantics. Thereby, the direct interaction of such systems is infeasible. A smart watch for example, which is developed in IOS could not interwork with smart bulbs without a relevant proprietary gated application from the same brand. Therefore, islands of IoT functionality are established, leading towards a vertical 'Intranet-of-Things' instead of the actual vision of an 'Internet-of-Things'. To presume upon the full potential of the IoT setting, we require standards for accomplishing the desired horizontal and vertical operation, communication, and programming across platforms/devices, independent of their vendor and/or model.

Thus, the deployment of Semantic Mediators (SMs) is recommended in this study in order to correlate the required information and materialize cross-domain interaction with interoperability between systems of different semantics. The SMs transform data in the same format and resolve potential conflicts between the different thing descriptions. Security countermeasures are also deployed, protecting the data both in transit and at rest. The main contributions of the proposed SMs include:

- cooperation with legacy, XML, and JavaScript Object Notation (JSON) formats,
- support of World Wide Web Consortium (W3C) initiatives for IoT and LD, i.e., standardized ontologies, iot.schema.org, JSON for Linking Data (JSON-LD),
- direct processing of JSON-LD data by the inference and reasoning modules (with SPARQL-LD),
- secure transmission of data (Transport Layer Security (TLS) at all communications) defending the system against data in transit attacks,
- validation of the data legitimacy prior their usage (i.e., JSON Web Signature, JWS Javascript Object Signing and Encryption (JOSE) framework), protecting against the data in rest semantic attacks,
- distributed functionality across the edge, network and backend systems,
- efficient and scalable operation,
- advancements over state-of-the-art solutions.

During 2019 we completed the implementation of the described idea above and a paper has been accepted in the IEEE Global Communications Conference [57].

This study focuses on the secure, semantic interoperability for IoT applications. This activity can be related to any pilots incorporating in their infrastructure IoT environments.

3.9 End-to-End Semantic Interoperability Mechanisms for IoT

Contact: Sotiris Ioannidis (FORTH)

Semantic interoperability is the designed property where various systems can interact with each other and exchange data with unambiguous, shared meaning. This enables knowledge discovery, machine computable reasoning and federation of different information systems. Traditionally, technologies from the Semantic Web are adapted in order to capture the inherited properties of the Internet of Things (IoT) domain. Such technologies provide common description and representation of data and services; they characterize things and their capabilities, deal with the semantic annotation, resource discovery, access management, knowledge extraction in a machine-readable or interoperable manner. Thus, the common interpretation of semantic information in a globally shared ontology could be quite useful. However, several local systems may utilize popular or standardized ontologies, eventually they extend them and establish their own semantics and interfaces. As a result, the direct interaction between these systems is not feasible. With this in mind, our goal is to propose the use of semantic interoperability mechanisms, which correlates the

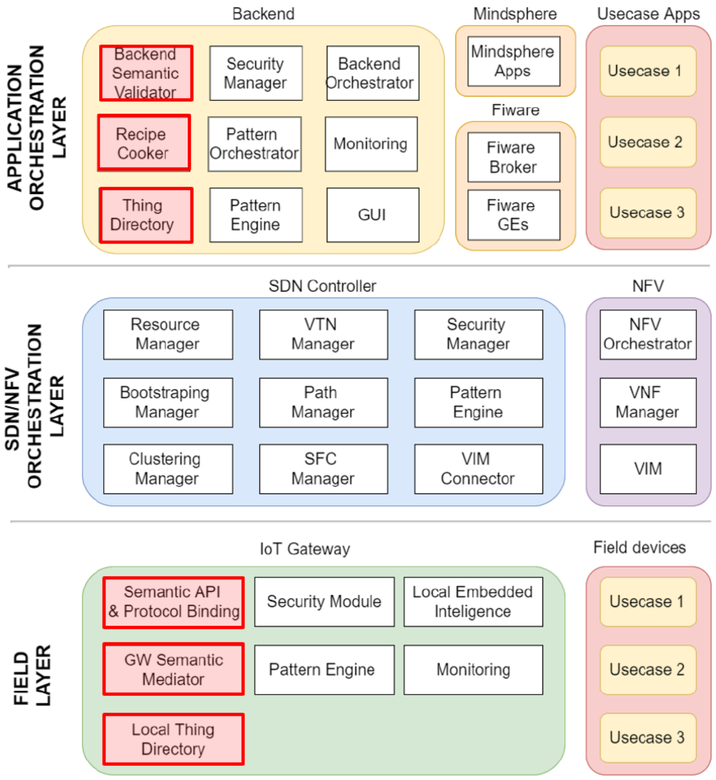

Figure 5: Layered architecture for semantic interoperability mechanisms

required information and enables the interoperability of systems with different semantics or cross-domain interaction. A motivating example of smart sensing is analyzed along with the implementation of the proposed approach.

Specifically we propose a framework that enables end-to-end compatibility and cooperation at all layers. Thus, semantic interoperability mechanisms across all levels must be deployed in order to resolve semantics between the field and backend layers. This approach is based on the architecture in [106], which consists of three layers: Field Layer, SDN/NFV Orchestration Layer and Application Orchestration Layer, as shown in Figure 5.

During 2019 we completed the implementation of our idea and we published a paper in the 24th IEEE International Workshop on Computer Aided Modeling and Design of Communication Links and Networks (CAMAD 2019) [74].

This study demonstrates the development of data transformation techniques and validation mechanisms to ensure end-to-end semantic interoperability. The proposed approach includes mappings between datatypes to ensure that data flow is possible between smart objects (Things). In addition, semantic transformation methods are defined with the purpose of resolving, if possible, conflicts among the semantic annotations. The representative motivated scenario is briefly described involving smart temperature sensing; the functional components of the architecture in [106], which are involved and are responsible for logical and structural data transformation, are illustrated in detail. Consequently it is related to CONCORDIA's pilot 2.1 Telecom Sector: Threat Intelligence for the Telco Sector and specifically with "Handling Privacy an Anonymity with Machine Learning".

4 Network-Centric Security (T1.2)

The focus of task 1.2 (T1.2) of the CONCORDIA project is network-centric security. T1.2 has particular emphases on three topics within this context: protection against (Distributed) Denial-of-Service ([D]DoS), analysis of encrypted network traffic, and usage of Software-Defined Networking (SDN) to provide resilient network services. More specifically, the intention of T1.2 is to

- investigate proactive, coordinated and distributed strategies to defend against DDoS attacks;
- collect, share and analyze data on attacks to the end of attack mitigation and attribution;
- develop techniques to monitor encrypted traffic for security purposes; and
- investigate SDN as a means to form a trusted and resilient Internet for Europe.

In the following sections, we report ongoing research efforts within the context of T1.2. In addition, we will outline how these efforts relate to the pilot projects of the CONCORDIA project.

Statistics show that over half of Internet traffic is nowadays encrypted, using encryption protocols such as TLS and IPsec. To the end of network security and threat detection, being able to monitor (meta-level) properties of encrypted communications is important. Section 4.1 outlines our research efforts to enrich encrypted network traffic traces with metadata that can be used, e.g., by monitoring systems, to the end of threat detection. In Section 4.2, we investigate the utilization of hardware accelerators such as graphics processing units (GPUs) to process encrypted network traffic. This paves the way for scalable, relatively inexpensive solutions to analyze traffic. Section 4.3 details advances that we made when it comes to the creation of encrypted network traffic datasets. Such datasets are generally lacking publicly, yet key when it comes to testing, e.g., threat detection solutions.

Within the context of network security, threat intelligence, situational awareness and decision-making are pivotal to deal with security incidents and to bolster resilience and security. This self-evidently goes beyond just encrypted traffic. In Section 4.4, we introduce a novel approach for supplementing information security metrics with time series analysis to enable anomaly detection and to support decision making when faced with security incidents. Section 4.5 proposes *MENTOR*, a recommender system for protection services that can support decision-making for cybersecurity management. We also investigated the effectiveness of Internet-wide DDoS attack mitigation efforts. Specifically, in Section 4.6, we investigate the effects of an FBI takedown campaign on DDoS-for-hire websites (i.e., Booters). The Domain Name System Security Extensions (DNSSEC) enable securing information in the DNS against forgery. In Section 4.7, we do a comprehensive analysis of the first ever DNS root KSK (key signing key) rollover. The Resource Public Key Infrastructure (RPKI) can be used to secure the Border Gateway Protocol (read: Internet routing) against certain types of hijacks. So shifting from DNS security to routing security, we perform a longitudinal study of RPKI in Section 4.8.

The blockchain technology has the potential to redefine collaborative processes in many domains, including cybersecurity. To integrate this technology in security solutions, it is paramount to first understand and evaluate the security and resilience of the blockchain technology itself. In Section 4.9 we investigate various novel, theoretical attacks on blockchain. Specifically, attacks that target blockchain on the application-level. We also outline ongoing efforts to investigate attacks on the underlying (peer-to-peer) networks. Section 4.10 presents the Blockchain Signaling System in which we introduce a service model for a purpose-built mitigation-as-a-service to combat DDoS attacks on networks. Such a system can be used to coordinate and distribute attack defense. Multi-disciplinary collaboration and threat intelligence exchange involves diverse data. A major challenge is to visualize this data to create situational awareness and support mitigation decision-making processes. Section 4.11 discusses our efforts in this area. Section 4.12 investigates the fault tolerance of blockchain networks, specifically focusing on popular, open-source frameworks that are readily available to be integrated into cybersecurity solutions.

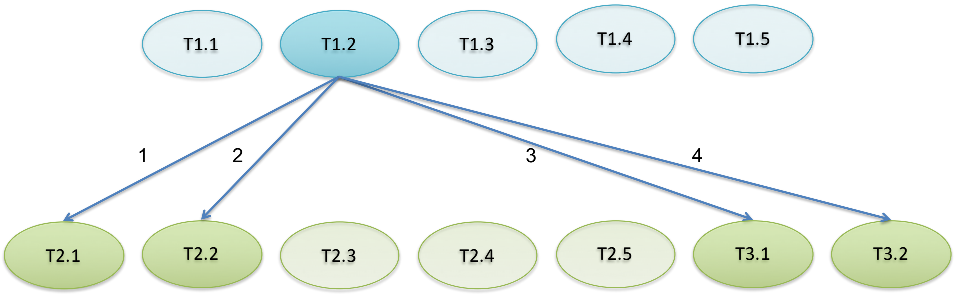

Figure 6: Links between task T1.2 and the CONCORDIA pilots

Figure 6 shows how research done in T1.1 relates to the CONCORDIA pilots.

1. The diverse T1.2 research efforts on encrypted network communication annotation and line-rate traffic inspection and processing ties into all CONCORDIA task in which threats may be learned from encrypted network traffic and traces (e.g., T2.1, T2.2 and T3.2).

2. The T1.2 work on test dataset creation (Trace-Share) can service as a building block for tasks in which CONCORDIA partners produce or consume (baseline) datasets (e.g., to test threat intelligence systems).

3. The work on recommendation systems for DDoS mitigation ties into, among others, the Threat Intelligence for Europe and DDoS Clearing House pilots (tasks T3.1 and T3.2, respectively).
4. The investigation of blockchain security serves as a precursor for CONCORDIA tasks that are envisioned to leverage blockchain technology. This, for example, includes Threat Intelligence for the Telco Sector (T2.1) and Piloting a DDoS Clearing House for Europe (T3.2).
5. The investigation of blockchain as a means for a collaborative DDoS attack mitigation system (BloSS) ties into CONCORDIA tasks that can build upon this work for information sharing about attacks, malware, and threats. This includes the Finance Sector Threat Intelligence (T2.2), among others.

4.1 Extracting Metadata from Encrypted Traffic

Contact: Pavel Minarik (Flowmon)

Statistics show that more than 60% of the Internet traffic is now encrypted, while this percentage is constantly increasing. The majority of communications are secured using common encryption protocols such as TLS and IPsec in order to ensure security and protect the privacy of Internet users.

Transport Layer Security (TLS) is a data encryption protocol that aims to secure transferred data from unauthorized access. In the case an attacker gains access to the transferred data, the attacker can only read random values instead of transferred application data. The protocol is implemented on top of the TCP connections, and all the TCP protocols can use it to protect their data. Probably the most common protocol which uses TLS is HTTP, in which case the protocol is called HTTPS.

Even though the transferred application data are encrypted, there is still some information useful for traffic analysis. This information is called: communication metadata. An example of such information is a list of supported cipher suites or server name. All these metadata are transferred during the initial communication stage called TLS handshake. From the perspective of network traffic analysis, it's important to be able to recognize and analyze these data.

Currently, the most used version of the TLS protocol is 1.2 (defined in RFC 5246). The most recent version is 1.3 (defined in RFC 8446), which tries to solve several possible security vulnerabilities and to speed up the process of connecting two devices by the TLS protocol. From the point of traffic analysis, the new version of TLS protocols reduces information transferred during the TLS handshake, which is available for analysis (are unencrypted).

In 2019, we designed and implemented a module for our Flowmon Probe that extracts specific metadata about network traffic encrypted with TLSv1.3 and export those metadata as part of traffic statistics in IPFIX format. By processing and analyzing TLSv1.3 metadata the following use cases can be accomplished:

- *cryptographic compliance* – ensures that only up to date and secure enough encryption algorithms and protocols are being used. Communication is protected based on valid certificates.
- *threat detection* – server name indication enables to distinguish real domain corresponding to user request. Client fingerprint exposed to the network indicates infection with malicious software.

We have also examined the encrypted traffic patterns of the TLS protocol and its changes introduced in version 1.3. We trained a machine learning model on TLS handshake parameters to identify the operating system of the transmitting device and compare its results to well-known identification methods. Our results were summarized in the article *Using TLS Fingerprints for OS Identification in Encrypted Traffic*, which was accepted for publication in the experience session of the *IEEE/IFIP Network Operations and Management Symposium (NOMS) 2020.*

Flowmon 10.3, which provides extended visibility into the TLS 1.3 protocol and other reporting on encrypted traffic, has been made generally available.

In 2020 we plan to analyze the new encryption protocol called QUIC (Quick UDP Internet Connections). The QUIC protocol is designed to securely transfer HTTP traffic over UDP connection, which will result in faster communication. The expected result is again module for Flowmon Probe providing additional insights into network traffic encoded in QUIC protocol.

The extensive annotation of encrypted traffic with metadata is projected to enable or improve efforts by consortium members to analyze potential threats in network traffic.

4.2 GPU-Accelerated Encrypted Network Traffic Inspection

Contact: Sotiris Ioannidis (FORTH)

Traditionally, Internet traffic analysis and monitoring is based on techniques such as deep packet inspection (DPI). The core functionality of such DPI implementations is based on pattern matching, that enables searching for specific strings or regular expressions inside the packet contents. Common applications of DPI include, but are not limited to, firewalls, intrusion detection and prevention systems, L7 filtering and packet forwarding. With the widespread adoption of network encryption though, DPI tools that rely on packet content are becoming less effective, demanding the development of more sophisticated techniques in order not to become obsolete. Traditional DPI implementations can only extract very coarse-grained information for the majority of encrypted traffic, even though its analysis is a core operation for many network systems. Self-evidently, network inspection systems need to be improved and adapted to current encryption trends.

An approach to inspect encrypted network traffic is the generation of signatures based on packet metadata, such as the packet timestamp, size and direction. These metadata can be usable even with encrypted traffic, since they can be easily extracted from packet headers. Recent related work has proven that revealing the traffic nature in encrypted communication channels is feasible. For instance, Conti et al. [35] proposed a system to analyze encrypted network traffic to identify user actions on Android devices, such as email exchange, interactions over social network, etc . Their framework leverages information that is available in TCP/IP packets, such as IP addresses and ports, among other features, such as packet size, direction and timing. Using machine learning techniques, they conduct their experiments that show that the system can achieve accuracy and precision higher than 95% for a number of user actions. Papadogiannaki et al. [102] proposed a pattern language to describe packet trains for fine-grained identification of application-level events in encrypted network traffic. They provided an efficient implementation of this language, namely *OTTer*, based on an extended version of the Aho-Corasick algorithm [102]. This approach is tested against real traffic and presents a minor CPU overhead when integrated with a proprietary DPI engine. Current solutions that focus on detecting malicious network traffic include Symantec's Encrypted Traffic Management (ETM) and Cisco's Encrypted Traffic Analytics (ETA) tools. ETM gains visibility into encrypted traffic to stop threats. Yet, this approach could violate user privacy since traffic is decrypted using SSL visibility appliances. ETA uses a more sophisticated technique that combines many different features of traffic. Still, this solution remains proprietary.

In this work, we investigate the utilization of hardware accelerators, such as GPUs, for high performance metadata matching against network traffic. The benefits for such an implementation is the high processing throughput as well as the low cost of powerful commodity high-end GPUs (in contrast to expensive server setups) . Since GPUs offer stream processing, real-time traffic inspection can be achieved. Fast metadata matching can enhance the implementation of numerous applications tailored for encrypted networks, such as traffic monitoring and intrusion detection.

During 2019 our research efforts were primarily concerned with the design and initial implementation of this idea. We also presented our work as a poster in the ACM Celebration of Women in Computing womENcourage 2019 – "Diversity Drives Societal Change" [103].

This work focuses on network traffic inspection through GPUs in order to accelerate the performance of pattern matching. The results are therefore useful, among others, for systems that inspect network traffic for cyber threats. As such, our work can be connected to two CONCORDIA threat intelligence pilots. Specifically, Task 2.1 ("Threat Intelligence for the Telco Sector") and Task 2.2 ("Assessing Cyber Risks, Threat Intelligence for the Finance Sector").

4.3 Generation of Encrypted Network Traffic Datasets

Contact: Martin Drasar (MUNI)

Monitoring of encrypted traffic is a multifaceted problem, of which we research two main domains: creation of encrypted datasets and encrypted traffic analysis. The advances in both domains have impact beyond these domains and contribute to cyber situational awareness, anomaly detection, attacker behavior prediction, and so on.

The availability of publicly shared datasets of sufficient quality is a prerequisite for reproducible research, and as such is highly encouraged [15]. Without such datasets, research results cannot be reliably proven, and common mistakes are repeated [3]. In the area of network anomaly detection, such a situation is less than satisfactory. The creation of acceptable datasets is challenging because both real-world and artificial datasets come with their share of issues. As a result, there is currently no scientific consensus on accepted high-quality datasets [15], and the evaluation of different analytic methods (e.g., for threat detection) is either not possible or limited to obsolete datasets.

One of our research efforts within this task is to produce current, modifiable, extensible, reproducible, and publicly shareable datasets with optionally encrypted parts, which can be used as a verified baseline by the community. To this end we have developed a methodology to create semi-labeled network traffic datasets and developed a platform for their actual creation. The crux of the methodology is a behavior-aware combination of so-called *annotated units* with un-annotated background traffic. The *annotated units* contain traffic of interest, such as network attacks, and are a result of careful extraction from real-world traffic or a product of simulation. These units are normalized, modified, and then combined with provided background traffic. The combination accounts for statistical differences in a number of traffic features to make it blend with the background traffic.

To address the especially problematic lack of encrypted datasets, we have suggested an extension to the aforementioned dataset creation pipeline. It enables replaying of datasets through virtual tunnels and subsequent collection of encrypted data, which can be mixed back into the background traffic.

We developed and released publicly a platform to create datasets. The platform is called Trace-Share (`https://github.com/CSIRT-MU/Trace-Share`). Trace-Share will be the basic building block for future on-demand creation of encrypted datasets, and will enable partners in the CONCORDIA project to produce or consume baseline datasets.

Our research in the following year 2020 will focus on the limits of this process regarding alterations of traffic properties, and on evaluation of dataset quality in synthetic and real-world environments.

We also assert that even with the proposed approach to mixing attacks into a background traffic, the contents would only reflect existing attacks, which were detected in the real networks. To provide wider attacks' variety for better testing and training of detection methods, we propose to create a custom cybersecurity-oriented environment for simulation and evaluation of new variants of cyber attacks. This environment will harness the current advances in reinforcement-learning- and heuristic-based agents. The environment will be closely tied to the aforementioned Trace-Share platform and will leverage it to directly produce datasets from the simulation results.

The aim of this work is to produce quality datasets of both encrypted and unencrypted traffic, together with extensive annotation of their contents. As such, it can be employed by various members of the project consortium, who want to train or benchmark their analytic methods. There are tentative plans to use Trace-Share generated datasets in Task 2.1 ("Threat Intelligence for the Telco Sector") to stress-test the detection methods on data that would be hard to acquire from production telco infrastructure.

4.4 Security Metrics for Quality Control and Situational Awareness

Contact: Jan Kohlrausch, Christian Keil

Recently, technical platforms have been developed to share, collect, and analyze threat intelligence data. Relevant platforms that are actively deployed in the security community as well as the CONCORDIA project include MISP [140] and the "Central Clearing House" (CCH). The deployment of these platforms involves multiple challenges. As demonstrated by recent research, assuring data quality (e.g., Indicators of Compromise) is crucial for threat intelligence applications such as attack detection. Focusing on threat intelligence feeds, Li et al. [76] and Pinto et al. [107] introduce quality metrics to assess and evaluate different dimensions of data quality. We refer to [63] for a widely accepted definition of information security metrics. Other research shows that threat intelligence data can be used for network situational awareness by detecting anomalies in the monitored data. Zhou et al. [151] and Yaacob et al. [145] apply ARIMA time series analysis on network security data for anomaly detection and to forecast network traffic, supporting situational awareness.

Following the success of applying information security metrics and ARIMA time series analysis on specific threat intelligence applications, we started researching a generalized combination of both approaches, facilitating and improving processes for quality assurance and situational awareness. More specifically, we propose to introduce a novel process supplementing well-defined sets of information security metrics with ARIMA time series analysis to support the following aims:

Enabling Anomaly Detection ARIMA time series analysis allows us to forecast future measurements and to estimate confidence intervals of measurements.

Based on this, anomalies and significant changes in the data can be detected, supporting situational awareness processes.

Detecting Time Dependence Time series analysis allows us to detect and assess time dependence of measurements taken by metrics. For example, numbers of alerts or attacks may be constant (only exhibit statistical fluctuations), constantly change over time, or exhibit abrupt changes (change points). Such properties can be revealed by applying ARIMA time series analysis. This enables, for example, to more precisely estimate future resources and costs of defense mechanism and security controls (situational awareness).

Supporting Decision Making The estimated statistical uncertainties and confidence intervals help to improve decision making as required by quality assurance and situational awareness processes. More specifically, the statistical properties allow to compute the likelihood for measurements to be an anomaly. Thus, decision making based on information security metrics can be made more transparent. This is, for example, crucial for defining escalation procedures in quality assurance processes to react to a significant drop in data quality.

This process can be applied to all metrics that either count observations (e.g., number of IDS alerts per day) or that compute the average of measurements (e.g., average costs of incidents). It embraces metric selection, determining a suitable ARIMA model, fitting the parameters of the ARIMA model, and its application for data analysis and anomaly detection.

Based on the “CSIRT Handbook” [144] and the “Computer Security Incident Response Team (CSIRT) Services Framework” [46] we explored how data analysts and security practitioners can incorporate this method into the existing best practices for CSIRT services to support and facilitate quality assurance and situational awareness.

Our early efforts resulted in the paper *ARIMA Supplemented Security Metrics for Quality Assurance and Situational Awareness*, which was submitted for publication at *ACM DTRAP, FIRST special issue, 2019*. After a minor revision the paper is supposed to be accepted.

The new class of ARIMA supplemented metrics will be implemented in the Centralised Clearing House (CCH) to support quality assurance and situational awareness processes. The CCH is part of CONCORDIA’s Threat Intelligence Platform developed in Work Package 3. An important aspect of quality assurance is to measure parameters indicating, for example, technical problems affecting data ingestion. Moreover, the parameters allow to assess the properties of submitted data (e.g., the timeliness of reporting security events) and to detect changes over time. Since a large number of sites contribute data to the CCH, it is reasonable to assume that major incidents on the Internet, such as the spreading of new worms or global botnet activity, will result in a significant change of the quantity of reports submitted to the CCH. We will apply the anomaly detection capabilities as provided by the ARIMA time series analysis to detect such anomalies. These results will supplement the threat intelligence platform as provided by the CONCORDIA project.

4.5 MENTOR: On the Recommendation of Protection Services

Contact: Muriel Franco, Burkhard Stiller (UZH)

Currently, companies invest in protection services (e.g., firewalls and anti-malware tools) and response teams to ensure availability and to protect crucial services and infrastructure. The cybersecurity market is worth billions of dollars [85] and investments are steadily rising. Thus, there are financial incentives for Protection Service Providers (PSP) to enter the market by offering protection services, while end-users can reduce protection costs by leveraging a competitive market for cybersecurity to meet their specific demands. These protections may include the acquisition of physical appliances, software licenses, virtual network functions, and cloud-based protection. Thus, although traditional models will still meet specific demands, a notable amount of next-generation protection services can adapt to flexible business models and provide a different level of protection on-demand.

It is not a trivial task for end-users to select one of them. Decision-making is even more critical when infrastructure is under attack and the decision to mitigate the attack should be provided on the basis of information about the infrastructure, such as its economic aspects, demands, and the characteristics of the attack. In this scenario, it is essential to observe not only how often attacks surpass the on-site infrastructure capacity, but also which off-site services can provide the necessary protection, considering their different service flavors, such as the amount of traffic supported, the capacity to address particularities of a determined attack, and price conditions. In this sense, recommender systems [128] provide a valuable security management tool to support decision-making during the detection and mitigation process.

We propose *MENTOR*, a protection service recommender system as a support tool for cybersecurity management that is able to recommend services for the prevention and mitigation of cyber attacks. This work investigates similarity measure techniques to correlate information, such as budget constraints and the type of service required, from customers with different services available. Based on this, *MENTOR* is able to indicate an adequate service to protect infrastructures according to different demands, such as region, deployment time, and price conditions. In addition, an evaluation and discussion determine the performance and accuracy of each similarity measure technique implemented within *MENTOR*.

The *MENTOR* system assists network operators during the decision-making process concerning measures to protect critical infrastructure. For this, the recommender engine indicates protection services available from different Protection

Service Providers (PSP) to prevent and mitigate threats. *MENTOR* considers different properties from available protection services, the customer profile, and characteristics of the cyber attack to establish a fair recommender system, where one or more services from different PSPs (e.g., both small companies and global players) can be proposed to neutralize a threat efficiently, while minimizing costs and reducing damage.

This work resulted in the paper *MENTOR: The Design and Evaluation of a Protection Services Recommender System*, which was published and presented at the IFIP 15th International Conference on Network and Service Management (CNSM 2019) [49]. Besides that, part of this work was presented during the CONCORDIA Open Door talk "Update on Economic Aspects of Cyber-security", where MENTOR was introduced as a tool to help in cybersecurity economic planning (e.g., costs and risks reduction).

The results of the research efforts related to MENTOR can be integrated in various parts of the CONCORDIA ecosystem. For example, they can be integrated with solutions developed in the context of T4.3, which focus on the economic aspects of cybersecurity (e.g., the SEConomy framework [117]). Concerning pilots, MENTOR has the potential to be integrated with the DDoS Clearing House (T3.2) as well as with the Threat Intelligence Center (T3.1) since data can be used as input to recommend adequate protections. Our work on MENTOR has already enabled it to use data from the DDoSDB [123], which is part of the DDoS Clearing house. Specifically, it uses DDoSDB data to understand the characteristics of real-world attacks and recommend the adequate solution.

4.6 The Effects of Taking Down DDoS-for-Hire Services

Contact: Jair Santanna, Mattijs Jonker (UT)

Booter services continue to provide popular DDoS-as-a-service platforms and enable anyone irrespective of their technical ability, to execute DDoS attacks with devastating impact. Since booters are a serious threat to Internet operations and can cause significant financial and reputational damage, they also draw the attention of law enforcement agencies and related counter activities.

In the first year, we investigated booter-based DDoS attacks in the wild and the impact of an FBI takedown targeting 15 booter websites in December 2018. We did so through the lens of a major IXP, a tier-1 ISP, and a tier-2 ISP. By purchasing attacks against our own infrastructure from 4 popular booters, we studied booter capabilities. The attack traffic levels generated by cheaper non-VIP services were considerably higher than reported in the literature (avg. 1.4 Gbps). We were the first to report the capabilities of a premium (VIP) booter service that peaks at 20 Gbps while promising 60-80 Gbps. In our data sets, we observed NTP-based DDoS attack traffic to be prevalent at all three vantage points. The attacks observed involve substantial traffic rates of up to 600 Gbps. To study if booter takedowns of law enforcement agencies help to reduce the attack traffic, we analyzed the effect of an FBI-led mass-seizure of 15 booter domains in Dec. 2018 on NTP, DNS, and Memcached-based DDoS attacks. We revealed that the takedown immediately had an effect on the DDoS amplification traffic especially towards reflectors. However, it did not have any significant effect on DDoS traffic hitting victims or on the number of attacks observed. This shows that seizing the seizing the domains of the booter websites is not enough as the underlying infrastructure of reflectors remains online and is utilized by third parties. Moreover, we found at least one booter to become active under a new domain shortly after the seizure, while the number of booter service domains in total increased over the measurement period despite the seizure. Our work aims to inform network operators to better understand the current threat-level, but also law enforcement agencies to recognize the need of additional efforts, e.g., to shut down open reflectors. Since our study is limited to technical parameters, the question arises whether this is sufficient to assess the health of the booter ecosystem. This motivates the need to better study the effects of law enforcement on the booter economy, e.g., on infrastructures, financing, or involved entities.

Our work was published in the renowned ACM Internet Measurement Conference (IMC'19) [71]. The results of this research updated earlier findings on Booter attack characteristics, which tie into the DDoS Clearing House pilot (T3.2).

4.7 A Comprehensive Study of the DNSSEC Root Key Signing Key Rollover

Contact: Roland van Rijswijk-Deij (UT)

The DNS Security Extensions (DNSSEC) add authenticity and integrity to the naming system of the Internet. Resolvers that validate information in the DNS need to know the cryptographic public key used to sign the root zone of the DNS. Eight years after its introduction and one year after the originally scheduled date, this key was replaced by ICANN for the first time in October 2018. ICANN considered this event, called a rollover, "an overwhelming success" and during the rollover they detected "no significant outages."

In our work, we independently follow the process of the rollover starting from the events that led to its postponement in 2017 until the removal of the old key in 2019. We collected data from multiple vantage points in the DNS ecosystem for the entire duration of the rollover process. Using this data, we study key events of the rollover. These events include telemetry signals that led to the rollover being postponed, a near real-time view of the actual rollover in resolvers and a significant increase in queries to the root of the DNS once the old key was revoked. Our analysis contributes significantly to identifying the causes of challenges observed during the rollover. We show that while from an end-user perspective, the roll indeed passed without major problems, there are many opportunities for improvement and important lessons to be learned from events that occurred over the entire duration of the rollover. Based on these lessons, we propose improvements to the process for future rollovers.

Our work was published in the renowned ACM Internet Measurement Conference (IMC'19) [87]. The results of this research speak to the security of the DNS, which does not only create situational awareness about DNS security, but also ties into higher-level applications developed within the context of CONCORDIA that rely on the DNS.

4.8 A Longitudinal Study of RPKI Deployment and Invalid BGP Route Origins

Contact: Roland van Rijswijk-Deij (UT)

Despite its fundamental and critical role in establishing Internet connectivity, the Border Gateway Protocol (BGP) remains highly vulnerable to attacks such as prefix hijacking, where an Autonomous System (AS) announces routes for IP space it does not control. To address this issue, the Resource Public Key Infrastructure (RPKI) was developed starting in 2008, with deployment beginning in 2011. This paper performs the first comprehensive, longitudinal study of the deployment, coverage, and quality of RPKI.

We use a unique dataset containing all RPKI Route Origin Authorizations (ROAs) from the moment RPKI was first deployed, more than 8 years ago. We combine this dataset with BGP announcements from more than 3300 BGP collectors worldwide. Our analysis shows the after a gradual start, RPKI has seen a rapid increase in adoption over the past two years. We also show that although misconfigurations were rampant when RPKI was first deployed (causing many announcements to appear as invalid) they are quite rare today. We develop a taxonomy of invalid RPKI announcements, then quantify their prevalence. We further identify suspicious announcements indicative of prefix hijacking and present case studies of likely hijacks. Overall, we conclude that while misconfigurations still do occur, RPKI is 'ready for the big screen," and routing security can be increased by dropping invalid announcements. To foster reproducibility and further studies, we release all RPKI data and the tools we used to analyze it into the public domain.

Our work was published in the renowned ACM Internet Measurement Conference (IMC'19) [32]. The results of this research speak to the security of the BGP layer, which is the primary routing protocol of the Internet and thus is a vital part under the hood of prospective CONCORDIA infrastructures (e.g., for data sharing and threat intelligence exchange).

4.9 State of the Art of Blockchains' Network Monitoring and Security

Contact: Thibault Cholez, Jean-Philippe Eisenbarth (UL)

The blockchain technology got tremendous attention in the last few years. It is considered as a disruptive technology that is able to redefine the collaborative processes in many activity domains, including cybersecurity. The billions of assets exchanged on a daily basis by public blockchains make them a new and attractive target for attackers. As such, blockchains can be considered in a dual way regarding cybersecurity: first, as a target to protect; and second, as a new tool to build new collaborative defense mechanisms. For instance, solutions have been proposed to create a blockchain-based platform to help detect and fight against distributed denial of service, or provide a secure logging of security incidents. Many other applications are possible in the field of cyber security.

With the rise of collaborative platforms that incorporate blockchain, the question of the performance and the resiliency of the blockchain networking infrastructure arises. Indeed, the P2P infrastructures and protocols supporting the blockchains become critical assets as more and more money and services are made on top of them, but they are largely undocumented and may be prone to severe attacks. With regard to the state of the art on P2P networks security, the fact that the infrastructure is distributed is not sufficient to assess its reliability, as much bias (e.g., if nodes are concentrated in a given geographical location) and attacks (e.g., eclipse attacks, Sybil attacks or partition attacks) are still possible and may severely disturb the network.

Eyal et al. introduced a new attack on the Bitcoin blockchain named the Selfish Mining Attack [43]. In this attack, a selfish miner (or a mining pool) chooses when to release/propagate his newly mined block rather than propagating it immediately to its peers. The main goal of this attack is to increase the profitability of the selfish miner. Attack scenarios involve two parameters: 1. α, the mining power of the selfish miner; and 2. γ, the fraction of honest nodes that mine on the selfish miner fork. Nayak et al. [90] proposed variants to this attack that optimize the profitability. These are called stubborn mining and trailing mining with sub strategies (lead mining, equal fork, n-trail). Nayak et al. argue that the stubborn mining strategies can lead to 25% additional gains compared to selfish mining. In more recent work, Grunspan et al. [51] showed that selfish mining strategies are more profitable than honest mining only after the difficulty adjustment (2 weeks in Bitcoin). Given some values of α and γ, the found the attack to not even be profitable anymore. There are two simulators of the selfish mining strategy that take into account the difficulty adjustment.[2] [3] Natoli et al. proposed an attack called the Balance Attack [89]. This attack is similar to the selfish mining attack. A selfish miner could

[2] https://greywyvern.com/code/javascript/selfishmining

[3] https://armankhosravi.github.io/dirtypool/

revert the state of the blockchain by releasing his private fork, thereby creating a double spend attack.

Overall, the proof-of-work consensus and the size of main public blockchains lead to huge computational costs for attacks. As such, the security offered becomes prohibitive and often superior to the expected gain of attacks.

Rather than only focusing on the application level, an attacker should rather try to disturb the underlying P2P network to weaken the consensus in some specific parts of the blockchain. Apostolaki et al. [11] showed that an adversary taking advantage of the apparent centralization of the bitcoin network could use BGP hijacks to partition the network and delay the propagation of blocks. An attacker could isolate 50% of the mining power by hijacking fewer than 100 IP prefixes and he could also waste up to 63% of a peer's mining power by delaying 50% of its connections. Heilman et al. [59] presented an attack on the bitcoin network that completely isolates a specific node by monopolizing all of its outgoing and incoming connections. Concerning the outgoing connections of a node, it tries to connect to peers it learned about previously (from the `tried` and `new` tables). The authors managed to populate the tables of the victim with its malicious nodes and to force the eviction of honest nodes. They showed that in the worst case for the attacker, there is a 85% probability to successfully perform this attack with a 4600-nodes botnet. Marcus et al. [80] introduced an eclipse attack on the ethereum network. Basically, to perform this attack, they just needed to generate a lot of Ethereum NodeID (executing a lot of clients on the same host) and send unsolicited `PING` message to the target to force him to rewrite its entire db table with the malicious NodeID (pointing to the same IP address). Once the target has created all its outgoing connections, the attacker only needs to fill all the remaining incoming connections of the victim with another host he controls. The developers of `geth` (the official Ethereum client) fix this vulnerability by limiting to 2 per buckets the number of nodes within the same IP prefix (/24 subnetwork) and to 10 for the entire `table`. The partial conclusion is that despite poor initial design choices that made attacks possible on the P2P network, the security mechanisms added in the last versions of the main clients are efficient to mitigate large attacks.

All things considered, it becomes evident that to develop security services on top of blockchains, it is of prime importance to analyze the underlying network to unveil possible limitations and vulnerabilities.

Our first objective in 2019 was to evaluate how reliable public blockchain platforms are based on the literature. We started investigating the state of the art to this end and identified existing attacks and their feasibility. We are currently investigating the feasibility of attacks that target the network (rather than the application) layer. Specifically, given the topology of networks, we are studying if small attacks targeting a few number of nodes can still provide the attackers a significant advantage over the blockchain or not.

This work is relevant to CONCORDIA to several extends. As mentioned before, the blockchain technology is increasingly used and becomes a cyber object that must be well understood and protected. Secondly, several tasks of CONCORDIA are envisioned to leverage the blockchain technology to implement the collaboration between partners and, more generally, stakeholders to build the European competence network. Two prominent examples are the collaborative DDoS clearing house platform (T2.1, T3.2) and the courses certificate framework (T3.4).

4.10 Cooperative Network Defense based on a Blockchain Signaling Platform

Contact: Bruno Rodrigues, Burkhard Stiller (UZH)

DDoS attacks attempt to render target system resources unavailable to legitimate users. Although widely known and studied for years, DDoS remains one of the major causes of concerns for service providers around the world. More recently, nations are also concerned with cyber threats [149, 105]. Defending against large-scale DDoS attacks *on-premises* becomes infeasible due to an increased volume of traffic achieved by those attacks. The main reason that drives the escalation of DDoS attacks is the increasing availability of stationary and portable devices, ranging from small sensors to cameras and home gateways connected to the Internet (i.e., the Internet of Things), that empower attackers to exploit a large number of insecure devices to generate malicious attacks on operational machinery and systems.

A cooperative defense is an alternative way to cope with large-scale DDoS attacks. In this approach, mitigation takes place at the egress points of the attack traffic. Its advantages over traditional/on-premise defenses have been widely recognized in the literature [149, 105, 86]. For example, by combining detection and mitigation capabilities of different domains, a cooperative defense approach reduces mitigation overhead at a single point. Moreover, malicious traffic can be tackled near its source. To date, however, no widespread deployment of a cooperative defense system has succeeded.

With the Blockchain Signaling System (BloSS) [115] we introduce a service model in which costs can be shifted from the autonomous system (AS) operators to potential customers subscribing to a purpose-built Mitigation-as-a-Service (MaaS) offering [79]. This works by directly using the fees paid by the customer as incentives to the operator of the AS. A provider under attack (i.e., the Target) may, if necessary, request mitigation services by submitting transactions to members whose purpose is to offer mitigation services and have infrastructure available to influence attacking traffic. A provider whose purpose is to offer mitigation services (i.e., the Mitigator) may define, in terms of incentives, what is necessary to deliver service expressing these terms in their smart contract. Once a mitigation service is accepted, a deadline to upload a proof of completion is started. The Mitigator can act in a rational way and upload a proof or miss the upload. Thereafter, the Target can rate the service of the Mitigator and, based on this rating, the funds initially locked in the contract are released to the Mitigator.

We demonstrated BloSS at the ACM SIGCOMM 2019 conference, during at the demo session [118]. The 2019 SIGCOMM edition saw around 1200 participants. Our demo experienced a considerable interest of participants working in the field of security and DDoS mitigation.

A blockchain-based system can transparently provide means to ensure that the exchange of information is not only visible to all members of a blockchain network, but also rewarded in order to foster the cooperative behavior. Therefore, a contribution to an information repository (e.g., threat intelligence) or the direct exchange of mitigation services can be rewarded whenever this contribution proves itself useful (e.g., when requested) or a mitigation service is effectively completed. Furthermore, insights gained on BloSS could be applied not only to distribute incentives, but to rate the quality of a potential contribution to a repository. The role of blockchain should be the guarantee of transparency and not of the mass storage of data itself, for which decentralized storage system can be used in combination with blockchain (e.g., IPFS).

BloSS aims to provide a solution for one of CONCORDIA's major research focuses, a coordinated collaborative effort for DDoS attacks. In this sense, there are research oriented tasks in WP1 such as (T1.2 - Network Centric) as well as the development of a Threat Intelligence Center (T3.1) and a DDoS Clearing House (T3.2) that could build upon BloSS. Therefore, the T2.2 pilot, whose objective is the construction of a Telcos intelligence center, can benefit from the advances of this research. While focusing specifically on DDoS attacks, BloSS can be used for sharing information about malware and threats.

4.11 Security Management and Visualization in a Blockchain-based Collaborative Defense

Contact: Christian Killer, Bruno Rodrigues, Burkhard Stiller (UZH)

Internet access becomes progressively democratic, including many different people using diverse types of stationary and portable devices. This leads to a concern about device security and the awareness of personal security. Recent statistics on security reports show, for example, not only a steady increase in the number of DDoS attacks, but also the number of long-lasting attacks (e.g., the most extensive attack in terms of time lasted longer than 12 days [64]). Centralized defenses can become a bottleneck due to the need to analyze all traffic measurements at a single location. Thus, the distributed nature of DDoS attacks suggests that a distributed and coordinated defense is the best alternative for a successful defense [95].

In this regard, a permissioned blockchain (BC) is a trustworthy, decentralized, and publicly available data storage, effectively supporting all members of the cooperative DDoS defense alliance. BC capabilities can be leveraged to build a platform for signaling attacks, serving as an immutable platform for the exchange of mitigation services defined in smart contracts (SC) of different peers, but also to provide incentives stimulating the cooperative behavior among service providers [116, 115]. Thus, if an attack is highly sophisticated and there are no countermeasures available, it is possible to request for cooperative mitigation for any domain participating in the alliance.

The major challenge, regardless of the underlying technology, is how to visualize information in a clear and objective manner considering the particularities of a collaborative defense. Thus, it is also required to consider that specialists must analyze not only internal threats, but also external mitigation requests. As the first step to threat mitigation is realizing its existence, it is critical for analysts to use a proper tool to structure and categorize data such that visualization "makes sense". A collaborative defense involves multi-disciplinary concepts, and the decision-making process usually requires a low response time from the user, but selecting an appropriate type of graphical representation and flow of interaction is not a straightforward task.

A cooperative defense adds a layer of complexity, in which not only should internal threats be analyzed and classified, but also threats of cooperative autonomous systems. We started investigating visualization in effort to complement the architecture of BloSS. Recall (from Section 4.10) that BloSS is based on a permissioned BC with a Proof-of-Authority consensus, enabling the visualization and management of cooperative defense requests in a dashboard. In more detail, our threat management dashboard provides a simple and objective interface for cybersecurity analysts for managing both incoming mitigation requests and requests for mitigation service, with the ability being possible to follow their progress in near real-time.

This work has been presented at the IEEE ICBC 2019. A poster [65] was discussed interactively with many participants during the coffee breaks.

A blockchain-based system can transparently provide means to ensure that the exchange of information is not only visible to all members of a blockchain network, but also rewarded in order to foster the cooperative behavior. Making this cooperative behavior accessible for cybersecurity analysts is crucial, since human operators need to decide whether or not to accept mitigation requests.

This work complements BloSS, which aims to provide a solution for one of CONCORDIA's major research focuses: a coordinated collaborative effort for DDoS attacks. The efforts to build a "Threat Intelligence for Europe" (T3.1) as well as the DDoS Clearing House pilot (T3.2.) can build upon or BloSS. In addition, T2.2, the objective of which is the construction of a Telco Threat intelligence center, can benefit from the advances of this research. While focusing specifically on DDoS attacks, BloSS can be used for sharing information about malware and threats too.

4.12 Fault Tolerance in Permissioned Blockchain Networks

Contact: Blaž Podgorelec, Muhamed Turkanović (UM)

Hyperledger is an open-source project created to advance cross-industry blockchain technologies. The project is hosted by the Linux Foundation. According to Forbes, the Hyperledger blockchain frameworks are among the most-considered frameworks by companies that are actively exploring the possibilities of blockchain technologies. Hyperledger Enterprise contains various, distinct frameworks (e.g., Hyperledger Fabric, Hyperledger Iroha).

The crash fault tolerance (CFT) property of blockchain systems relates to the requirement that the system operates normally despite node failures. Blockchain networks that are not tolerant when node crashes occur may suffer various security-related problems (e.g., failure of critical infrastructure).

Blockchain as a technology has potential uses in, for example, distributed collaborative defenses against (see: Sections 4.10 and 4.11). This motivated us to study the fault tolerance property of permissioned blockchain networks available in the public domain. Specifically, we focused our efforts on the popular Hyperledger frameworks: Hyperledger Fabric, Hyperledger Iroha. Our study involved a rigorous experiment [108], in which we analyzed the fault tolerance of both aforementioned blockchain platforms. By performing these analyses we established the CFT characteristics of the blockchain networks.

Our results show that anomalies in the blockchain networks may occur, whereby even the CFT property is not always 100% satisfied. During the analysis, we identified some anomalies associated with the CFT properties in the blockchain networks within one of the selected platforms. Following established practices for responsible disclosure, we reported this to the software maintainer, who updated the software to address the problem. Altogether, we reduced the risk of potential damage to enterprises that were actively using the frameworks in a production environment. For the aforementioned reasons, the research showed the importance that permissioned blockchain networks should furthermore be analyzed and tested, considering the slowly expanding adoption of such networks and platforms.

The results of our work have utility for pilots within the CONCORDIA project in which the possible use of a blockchain-based component is considered.

5 Software/System-Centric Security (T1.3)

Task 1.3 (T1.3) of the CONCORDIA project is concerned with software/system-centric security. Specifically, the main research topics addressed by T1.3 are

- malware analysis,
- security by design: adaptive software and OSs,
- detecting service dependencies, and
- system security validation and zero-days.

In this section we report on the research activities performed on these topics within T1.3. We start with shortly introducing each topic and identifying the papers that have been published per topic. In the remainder of this section the abstract for each paper is presented.

Malware Analysis

Malware detection and malware analysis are both crucial aspects of cybersecurity. For this, we develop two kinds of methods. The first is based on machine learning/AI models and has been successfully applied by the company Bitdefender (BD). The second is based on morphological analysis that has been devised at Lorraine University (UL) and which is now co-developed with the start-up Cyber-Detect (CYD). There are different issues to consider. Malware analysis is made by combining static and dynamic analysis. Dynamic malware analysis is traditionally performed by running a sample on a host system, suitably instrumented to detect interesting behaviors, and suitably isolated so as to prevent infection of the user's development/production environment. One of the challenges in such a setup is overcoming anti-research techniques used by malicious samples to avoid detection. Static malware analysis is traditionally performed by identifying malware by pattern analysis. One if the challenge is to go beyond, because malware writers know how to fool pattern analysis. That is why we propose and develop machine learning methods, morphological analysis methods and various other formal methods like concolic execution and model checking.

Within CONCORDIA, three partners perform research in the area of Malware Analysis: Bitdefender (BD), Cyber-Detect (CYD) and the University of Lorraine (UL). The following malware analysis papers have been published in 2019:

- Preventing File-less Attacks with Machine Learning Techniques, A.G. Bucevschi, G. Balan and D.B. Prelipcean, Proceedings of the 21st International Symposium on Symbolic and Numeric Algorithms for Scientific Computing, SYNASC 2019
- Improving Detection of Malicious Office Documents using One-Side Classifiers, S.C. Viţel, G. Balan and D.B. Prelipcean, Proceedings of the 21st

International Symposium on Symbolic and Numeric Algorithms for Scientific Computing, SYNASC 2019
- LockerGoga Quickly Reversed, G. Bonfante, C. Jannier, J.Y. Marion and F. Sabatier, Malcon Conference 2019

As we have a strong software activity in this part, we have added below their descriptions at the same level as the publications.

- GravityZone, developed by Bitdefender, is a business security solution built from ground-up for virtualization and cloud to deliver security services to physical endpoints, mobile devices, virtual machines in private, public cloud and Exchange mail servers.
- Gorille Pro, developed by Cyber-Detect and University of Lorraine, is a business malware analyzer of windows binary files, which is able to detect and analysis new threats.

Security by Design

The holy grail of cybersecurity is comprehensive security-by-design. However, this is also known to be an elusive target given the impossibility to be able to fully characterize all threats and efficiently mitigate them. The CONCORDIA approach is thus to provide dynamic, run-time and adaptive on-demand security by providing security "enhancers" to software and middleware by (a) identifying new threats as they occur, (b) analyzing the attack surface that leads to propagation of the attacks, and (c) providing run-time mitigation security of system/software components. A basic premise behind this approach is to fundamentally understand the ecosystem of technology, to explore and determine the attack surfaces at the system/software level over varied layers of abstraction, across components, interfaces and middleware protocols. We are developing technologies that facilitate trust on system components and data. Within CONCORDIA, two partners perform research in the area of Security by Design: RISE and Lancaster/TUD.

Code obfuscation is a major tool for protecting software intellectual property from attacks such as reverse engineering or code tampering. Yet, recently proposed (automated) attacks based on Dynamic Symbolic Execution (DSE) shows very promising results, hence threatening software integrity. Current defenses are not fully satisfactory, being either not efficient against symbolic reasoning, or affecting runtime performance too much, or being too easy to spot. We present and study a new class of anti-DSE protections coined as path-oriented protections targeting the weakest spot of DSE, namely path exploration. We propose a lightweight, efficient, resistant and analytically proved class of obfuscation algorithms designed to hinder DSE-based attacks. Extensive evaluation demonstrates that these approaches critically counter symbolic deobfuscation while yielding only a very slight overhead.

Within CONCORDIA, three partners perform research in the area of Security by Design: RISE, UL and Lancaster/TUD. The following security by design papers have been published in 2019:

- Assessing the State and Improving the Art of Parallel Testing for C, O. Schwahn, N. Coppik, S. Winter, N. Suri, Proceedings of the 28th ACM SIGSOFT International Symposium on Software Testing and Analysis, ISSTA 2019, 2019
- Inferring Performance Bug Patterns from Developer Commits, Y. Chen , S. Winter, N. Suri, Proceedings of the International Symposium on Software Reliability Engineering (ISSRE), 2019.
- How to Kill Symbolic Deobfuscation for Free (or: Unleashing the Potential of Path-oriented Protections), Mathilde Ollivier and Sébastien Bardin and Richard Bonichon and Jean-Yves Marion, Proceedings of the 35th Annual Computer Security Applications Conference, ACSAC 2019, pp 177–189.

Detecting Service Dependencies

The 2016 attack on DynDNS impacted Amazon, Netflix, Reddit, Spotify, Tumblr, and Twitter because of service dependencies. As of today, we have no good means and tools to discover and analyze service dependencies, which we believe are far wider spread that we usually believe. The analysis of service dependencies has already been started in the nineties, however the automated detection of service dependencies was not solved appropriately. It is necessary to understand how attackers can prolong the duration of disruptions. Furthermore, it is necessary to look at procedures used to "repair" attacked services and to think about how to effectively fail or slow down any repair actions. To achieve effective security assurance management, it behooves to obtain insight into the service and functional dependencies that exist in information systems. The dependency describes the effect of an object's attribute change (e.g., data value variance, program status alterations, or control flow redirection) that is caused by another object under a given condition. For example, the data dependency enables the unexpected error propagations that make it difficult to locate the real source of the errors. Besides, the combination of multiple dependencies triggers the disaster of dependency explosion that exacerbates the situation for achieving effective security assurance management. Therefore, the existence of the dependencies introduces substantial challenge for improving the security performance of information systems. To address the problem, an effective dependency analysis methodology is desired. Within CONCORDIA, the Technical University of Darmstadt (Lancaster/TUD) performs research in this area.

The following papers have been published in the area of service dependency detection in 2019:

- Gyro: A Modular Scale-Out Layer for Single-Server DBMS, H. Saissi, M. Serafini, N. Suri, Proceedings of the Symposium on Reliable Distributed Systems (SRDS), 2019
- Analyzing and Improving Customer-side Cloud Security Certifiability, S. Zhao, Y. Chen, S. Winter and N. Suri, Proceedings of the IEEE International Workshop on Software Certification (WoSoCer), 2019

- Extracting Safe Thread Schedules from Incomplete Model Checking Results, P. Metzler, N. Suri, G. Weissenbacher, International Symposium on Model Checking of Software (SPIN), 2019

System Validation and Zero Days

Security claims for systems and middleware are elusive. While protocols and crypto primitives can typically be specified via a rigorous set of assumptions with the correctness assurance provided either analytically or formally, there are no such techniques available at the systems level. As security properties are realized (on either hardware or software), the implicit limitations inherent in these implementations can often entail the inadvertent violation of the assumptions resulting in a compromise of the security claims. Consequently, the need exists for systematic testing and validation of assumptions over their (software) realizations. The task will explore validation methodologies for defining and executing tests and perturbations (i.e., stress conditions) in the implementation as attack scenarios Robots are cyber-physical systems, designed to perform specific tasks and ease human work. Current applications include, but are not limited to, military, industrial, agricultural, and domestic robots. With robots, the physical world is highly coupled with the cyberspace. Firstly, sensors perceive the physical environment; then, the control software chooses for actions, potentially in collaboration with other agents of the cyberspace (for example, other robots or the cloud); finally, actuators perform those actions on the physical environment. Regardless of their capabilities and size, robots take their actions based on what they sense. Thus, robots are targeted by sensor spoofing attacks that can force an incorrect behavior in the robotic system and undermine the success and safety of critical operations. Spoofing is the action of disguising a communication from an unknown source as being from a known one. Within CONCORDIA, the University of Luxembourg (SnT) performs research in the area system validation and zero days.

On the topic of system validation and zero days, the following paper has been published in 2019:

- Auto-encoding Robot State against Sensor Spoofing Attacks, S. Rivera, S. Lagraa, A.K. Iannillo and R. State, IEEE International Symposium on Software Reliability Engineering Workshops (ISSREW), 2019

Another important issues is the detection of known vulnerabilities, which are listed under the name Common Vulnerability Exposures (CVE). On one hand, CVE are discovered and listed every day. On the other, software are running and there are not always updated. As a result, an issue is to know which binaries are used as a component of software or simply which binaries are in a system. And then, the goal is to find whether or not those binaries contain publicly disclosed cybersecurity vulnerabilities and exposures (CVE). Indeed, a vulnerable binary may be the target of an attack. As a result, it is critical to know what is the nature of running binaries in a system in order to secure it. UL and CYD have join their forces in order to develop a proof of concept based on AI models.

Relation to Pilots

The research performed within Task 1.3 is linked to the CONCORDIA pilots in two ways (see Figure 7). First, there are papers that contribute to pilots:

1. The research on Isolating Critical Software carried out by RISE and SnT has potential for the e-health pilot (IoT-based healthcare). RISE, SnT and EESY Innovation are discussing this further.

2. The paper *Auto-encoding Robot State against Sensor Spoofing Attacks* has potentially interesting ideas for UAV anomaly detection.

Second, there are software and papers that contribute to pilots:

1. Software developed by BD, CYD and UL will be part of the Telecom pilot 2.1.

2. The proof of concept on CVE detection developed by CYD and UL has potentially applications for DDOS clearing house pilot 3.2.

Lastly, at least the malware part will receive data from the threat intelligence pilot 3.1, which is necessary to feed software.

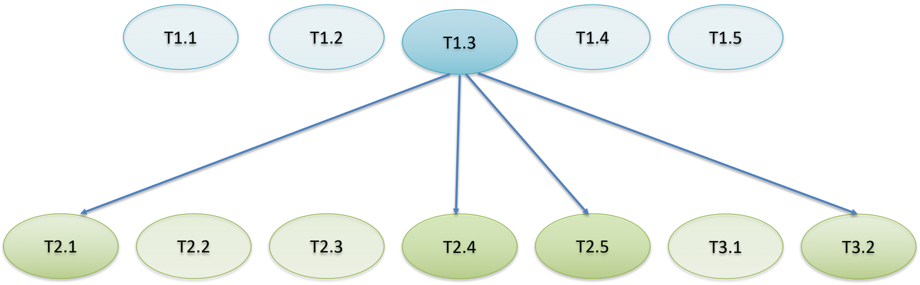

Figure 7: Links between task T1.3 and the CONCORDIA pilots

Highlights

Members of T1.3 have published 10 papers, including at top venues such as the Annual Computer Security Applications Conference (ACSAC), the ACM SIGSOFT International Symposium on Software Testing and Analysis (ISSTA) and the International Symposium on Model Checking of Software (SPIN). In addition, software like GravityZone and Gorille Pro have been developed, enabling the take-off of scientific ideas by industry.

5.1 Preventing File-less Attacks with Machine Learning Techniques

Contact: Bogdan Prelipcean (BD)

The cyber-threat detection problem is a complex one due to the large diversity of attacks, increasing number of prevalent samples and to the arms race between attackers and security researchers. A new class of attacks which appeared in the past years is modifying its spreading and action methods in order to become non-persistent. Being non-persistent, the usual detection and analysis methods which are file oriented, do not work anymore. Therefore, several solutions became available like memory introspection, process activity monitoring or application enforcement. However, these solutions are time consuming, therefore their usage impose some additional resources needs. In this paper we discuss an entry-level anomaly detection method of the command lines arguments which are passed to the most known system tools generally available in Windows and not only. Some of these tools are used for years in companies to automate tasks, but only in the recent period they became a powerful tool for the attackers. The main element that makes file-less so popular is the easiness in which the launches initial infection and assists with the attacks using legitimate Windows administrative tools already installed on the victim‘s system. Typical attacks exploit vulnerabilities in browsers and associated programs. The user may become a victim of this type of attack when browsing an exploit-kit affected website that exploits vulnerabilities in browsers, Java, PDF reader, Flash player, script-based malware using JavaScript/JScript and other client-side software. The method proposed by Bitdefender is based on a derived version of Perceptron algorithm. Besides the machine learning approach one of the main contributions is on the analysis and the extraction of features from the command line of a launching process. The purpose is to block a malicious chain of processes to act in a harmful way. The paper containing the results was accepted by the SYNASC 2019 Conference (September 2019) and published in the proceedings of the conference (December 2019).

5.2 Improving Detection of Malicious Office Documents using One-Side Classifiers

Contact: Bogdan Prelipcean (BD)

The current threat landscape is diverse and has lately been shifting from the binary executable application to a more light-coded and data-oriented approach. In recent years, the attacks based on malicious scripts (VBA) from Microsoft Office documents have been growing in popularity and so has the necessity of detecting them. While the VBA language was originally developed as a powerful scripting language to help users automate tasks and create macro-driven applications, it became a prevalent infection vector due to its extended capabilities, such as accessing the native Windows system calls. As a response to the growing threat of malicious macros, Microsoft has implemented several security features to prevent the execution of unwanted code. Nowadays, as a countermeasure, the attackers distribute the malicious documents through spear-phishing[4] emails and use social engineering techniques in order to trick the victim into enabling the execution of the malicious macros. Bitdefender developed a lightweight approach of a machine learning solution due to the limitations of it being used in an anti-malware solution where performance is a highly important requirement. Bitdefender's solution is based on a derived version of the Perceptron algorithm which builds a detection model focused on the properties of the macro code extracted from the VBA project of Microsoft Office files. The paper containing the results was accepted by the SYNASC 2019 Conference (September 2019) and published in the proceedings of the conference (December 2019).

5.3 LockerGoga Quickly Reversed

Contact: Jean-Yves Marion (UL), Werner Laurent (CYD)

Recent attacks of LockerGoga against Altran in France and Norsk Hydro in Norway illustrate the necessity to have advanced anti-malware defenses. The attack in France happened in January and the one in Norway in March. Those attacks should have been stopped.

Cyber-Detect is developing an anti-virus engine named Gorille. It is based on morphological analysis which is devised at LORIA, the computer science lab of Lorraine University. Today, Gorille is able to detect LockerGoga and its variants without any specific signature.

In a nutshell, Gorille identifies malicious threats embedded in Linux and Windows binary files. For this, Gorille knows a collection of malicious behaviors. Each binary file submitted to Gorille is then scanned and as soon as a set of malicious inter-link behaviors is detected, Gorille raises an alert. Gorille knows about 200,000 malicious behaviors. From these malicious behaviors, Gorille identifies 55 malicious behaviors in the submitted samples of LockerGoga. As a result, Gorille would have stop LockerGoga attacks.

The morphological analysis is based on the detection of malicious behaviors. Each malicious behavior consists of an abstraction of the control flow graph of the code instantiating this malicious behavior.

The result on LockerGoga was first published as a post on Cyber-Detect Blog [81]. The detailed of the results are accepted at Malcon conference (October 2019). The title of the paper [52] is LockerGoga quickly reversed.

[4] Spear-phishing is a phishing attempt targeted to certain individuals or companies. Phising is a fraudulent practice that aims to obtain sensitive information from users (such as passwords, credit card information, etc.) or trick them in performing certain actions by concealing this practice under the impression of communicating with a trustworthy entity.

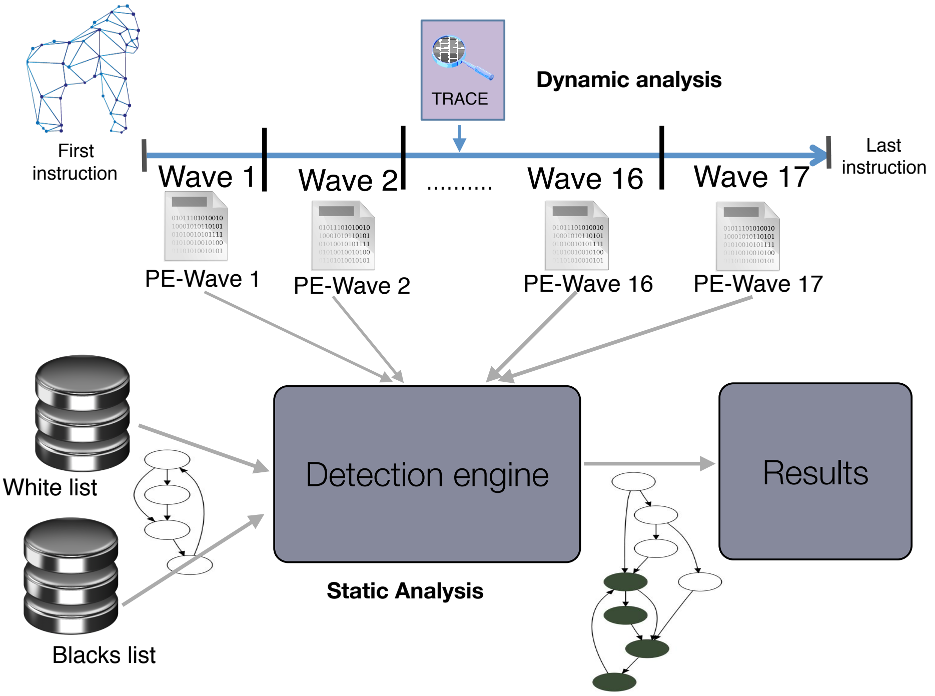

Figure 8: Architecture of Gorille Pro

CYD developed an orchestrator that sends requests to Gorille server together with a user-friendly interface in order to use Gorille, which works on Linux and Windows. This interface was used in the Malcon Conference paper and during the Cyberweek at Rennes, France. The main functionalities of the software called Gorille Pro are to detect matches with malicious behaviors and to see where code lines are similar and so to check by hand that Gorille outputs a correct answer.

Gorille Pro performs bot static analysis and dynamic analyses and it constitutes a *major development*. Figure 8 shows the overall architecture of the Gorille Pro software.

5.4 Software: GravityZone

Contact: Ovidiu Mihaila, Bogdan Prelipcean, Dragos Gavrilut (BD)

GravityZone is a business security solution built from ground-up for virtualization and cloud to deliver security services to physical endpoints, mobile devices, virtual machines in private, public cloud and Exchange mail servers. GravityZone is one product with a unified management console available in the cloud, hosted by Bitdefender, or as one virtual appliance to be installed on company's premises, and it provides a single point for deploying, enforcing and managing security policies for any number of endpoints and of any type, in any location. The product delivers multiple layers of security for endpoints and for Microsoft Exchange mail servers: anti-malware with behavioral monitoring, zero-day threat protection, application control and sandboxing, firewall, device control, content control, anti-phishing and antispam. Each role instance can be installed on a different appliance. Built-in role balancers ensure that the GravityZone deployment protects even the largest corporate networks without causing slowdowns or bottlenecks. Existing load balancing software or hardware can also be used instead of the built-in balancers, if present in the network. GravityZone incorporates techniques enabling device hardening and control, detection at both the pre- and on-execution stages, automatic clean-up and rollback action. Detection in pre- and on- execution is achieved by static and behavioral analysis with Machine Learning/AI models, along with powerful heuristics and process monitoring methods. Another feature is the possibility of sending suspicious files in an isolated sandboxing environment where the files behavior will be analyzed and registered. All these while delivering single-pane-of-glass visibility, alerts and notification, integration and reporting capabilities.

5.5 Software: Gorille Pro

Contact: Ovidiu Mihaila, Bogdan Prelipcean, Dragos Gavrilut (BD)

Gorille Pro is a tool that analyzes highly obfuscated binary codes (Windows PE X86) like malware. The technology used by Cyber-Detect leans on the concept of morphological analysis developed at Lorraine University. It is based on the automatic construction of multi-dimensional signatures, stored in a knowledge behavior graph databases, and so capturing binary code behavior and functionalities. Thanks to a homemade dynamic tracer, most of malware protections are bypassed. Gorille Pro has a strong detection engine thanks to the use of AI models and of algorithmics methods. See Figure 8 for the overall architecture.

5.6 Isolation of Critical Software

Contact: Anum Khurshid (RISE)

The increasing number of connected devices in the technological ecosystem has led to researchers and system manufacturers to deal with security threats in a more serious way. Conventional high end systems have a different threat landscape and security propositions to address them. With the magnitude of connectivity in Internet of Things (IoT), issues like trusted data, security of system components, secure communication are major concerns. Hardware-assisted Trusted Execution Environments (ARM TrustZone, Intel SGX) provide isolation of data and operations running in parallel with the operating system. They address a number of security requirements of low-end IoT nodes (like sensors and actuators):

- *Data protection*: Sensitive data can be stored in secure memory spaces and can only be accessed by secure software. Non-secure software can only gain access to secure APIs providing services to the non-secure domain, and only after security checks or authentication.

- *Firmware protection*: Firmware that is preloaded can be stored in secure memories to prevent it from being reverse engineered and compromised by malicious attacks.
- *Critical Operation protection*: Software for critical operations can be preloaded as secure firmware and the appropriate peripherals can be configured to permit access from the secure state only. In this way, the operations are protected from intrusion from the non-secure side.
- *Secure boot*: The secure boot mechanism enables confidence in the platform, as it will always boot from secure memory.

TrustZone-M is used to provide the above mentioned securities in IoT nodes based on Cortex-M processors. Software designed to secure IoT systems (such as secure boot, cryptography libraries, and authentication) can be isolated from untrusted software modules including user applications, third party libraries, device drivers, and protocol stacks. However, the untrusted software modules can access secure world resources via application programming interfaces (APIs) provided by the secure software.

Although, TrustZone for Cortex-M processors enables separation of security critical computation and data resources, it fails to offer secure communication guarantees between the two worlds. Non-secure program in the non-secure world can communicate with any secure software in the secure world using direct function calls, i.e., there is no way to limit or control access to security critical resources in the secure world. Furthermore, there is no way to determine if a request for security critical resources is coming from a certain application in the non-secure world. This means it would not be possible to stop malicious code running in the non-secure world, from making a request to any secure world resource. This allows malicious applications to falsify requests and repeatedly pass problematic data (a maliciously-crafted message) to discover the vulnerabilities of the secure software. TrustZone does not provide a way to protect messages transmitted during cross-world communication. In most TrustZone-based trusted execution environment (TEE) solutions, a non-secure program may use a shared memory zone in the non-secure memory area, as a channel for the transmission of messages between the two worlds. This channel is vulnerable to attacks if the non-secure world is infected with malware. As a result, messages transferred through this channel could be intercepted and manipulated by malware. This security issue limits the effectiveness of TrustZone-based TEE. Ongoing research in this task presents a framework for establishing a secure communication channel between the non-secure and secure domains of TrustZone-M.

RISE and SnT collaborate on research related to hardware-assisted trusted execution environment, see also Section 3.5. SnT plans to develop a verification and benchmarking framework which can be used to validate the secure communication mechanism (ongoing research at RISE). The e-Health pilot (T2.4) is focused on providing end-to-end security for e-Health communications via Internet, Web and Cloud and will be deployed and evaluated in various use cases in IoT based systems. Discussions are ongoing between RISE, SnT and EESY Innovation regarding TEE-based IoT solutions for IoT-based healthcare. In general, within CONCORDIA strong collaborations exist between RISE and ERICSSON.

5.7 Assessing the State and Improving the Art of Parallel Testing for Programs Written in C

Contact: Oliver Schwahn (Lancaster/TUD)

The execution latency of a test suite strongly depends on the degree of concurrency with which test cases are executed. However, if test cases are not designed for concurrent execution, they may interfere, causing result deviations compared to sequential execution. To prevent this, each test case can be provided with an isolated execution environment, but the resulting overheads diminish the merit of parallel testing. Our large-scale analysis of the Debian Buster package repository shows that existing test suites in C projects make limited use of parallelization. We present an approach to (a) analyze the potential of C test suites for safe concurrent execution, i.e., result invariance compared to sequential execution and (b) execute tests concurrently with different parallelization strategies using processes or threads if it is found to be safe. Applying our approach to nine C projects, we find that most of them cannot safely execute tests in parallel due to unsafe test code or unsafe usage of shared variables or files within the program code. Parallel test execution shows a significant acceleration over sequential execution for most projects. We find that multi-threading rarely outperforms multi-processing. Finally, we observe that the lack of a common test framework for C leaves make as the standard driver for running tests, which introduces unnecessary performance overheads for test execution.

The results from the research were published in Proceedings of the 28th ACM SIGSOFT International Symposium on Software Testing and Analysis (ISSTA) [127].

The research is coupled with task T1.3: Security by design as it focuses on identifying safe execution strategies for parallel testing.

5.8 How to Kill Symbolic Deobfuscation for Free

Contact: Jean-Yves Marion (UL)

Code obfuscation is a major tool for protecting software intellectual property from attacks such as reverse engineering or code tampering. Yet, recently proposed (automated) attacks based on Dynamic Symbolic Execution (DSE) shows very promising results, hence threatening software integrity. Current defenses are not fully satisfactory, being either not efficient against symbolic reasoning, or affecting runtime performance too much, or being too easy to spot. We present and study a new class of anti-DSE protections coined as path-oriented protections targeting the weakest spot of DSE, namely path exploration. We propose a lightweight, efficient, resistant and analytically proved class of obfuscation algorithms designed to hinder DSE-based attacks. Extensive evaluation demonstrates that these approaches critically counter symbolic deobfuscation while yielding only a very slight overhead.

Reverse engineering and code tampering are widely used to extract proprietary assets (e.g., algorithms or cryptographic keys) or bypass security checks from software. Code protection techniques precisely seek to prevent, or at least make difficult, such man-at-the- end attacks, where the attacker has total control of the environment running the software under attack. Obfuscation aims at hiding a program?s behavior by transforming its executable code in such a way that the behavior is conserved but the program becomes much harder to understand. Even though obfuscation techniques are quite resilient against basic automatic reverse engineering (including static attacks, e.g. disassembly, and dynamic attacks, e.g. monitoring), code analysis improves quickly. Attacks based on Dynamic Symbolic Execution (DSE, a.k.a. concolic execution) use logical formulas to represent input constraints along an execution path, and then automatically solve these constraints to discover new execution paths. DSE appears to be very efficient against existing obfuscations, combining the best of dynamic and semantic analysis.

We study path-oriented protections, a class of protections seeking to hinder DSE by substantially increasing the number of feasible paths within a program.

1. We detail a formal framework describing path-oriented protections. We characterize their desirable properties, namely tractability, strength, and the key criterion of single value path (SVP). The framework is predictive, in the sense that our classification is confirmed by experimental evaluation, allowing both to shed new light on the few existing path-oriented protections and to provide guidelines to design new ones. In particular, no existing protection achieves both tractability and optimal strength (SVP). As a remedy, we propose the first two obfuscation schemes achieving both tractability and optimal strength.

2. We highlight the importance of the anchorage policy, i.e., the way to choose where to insert protection in the code, in terms of protection efficiency and robustness. Especially, we identify a way to achieve optimal composition of path-oriented protections, and to completely prevent taint-based and slice-based attacks (two powerful code-level attacks against obfuscation), coined as resistance by design.

3. We conduct extensive experiments with two different attack scenarios, exhaustive path coverage and secret finding. Results confirm that path-oriented

> protections are much stronger against DSE attacks than standard protections (including nested virtualization) for only a slight overhead. Moreover, while existing techniques can still be weak in some scenarios (e.g., secret finding), our new optimal schemes cripple symbolic deobfuscation at essentially no cost in any setting. Finally, experiments against slice, pattern-matching and taint attacks confirm the quality of our robust-by-design mechanism.

The results from the research was published in the proceedings of the Annual Computer Security Applications Conference (ACSAC 2019) [96].

5.9 Inferring Performance Bug Patterns from Developer Commits

Contact: Yiqun Chen (Lancaster/TUD)

Performance bugs, i.e., program source code that is unnecessarily inefficient, have received significant attention by the research community in recent years. A number of empirical studies have investigated how these bugs differ from "ordinary" bugs that cause functional deviations and several approaches to aid their detection, localization, and removal have been proposed. Many of these approaches focus on certain sub-classes of performance bugs, e.g., those resulting from redundant computations or unnecessary synchronization, and the evaluation of their effectiveness is usually limited to a small number of known instances of these bugs. To provide researchers working on performance bug detection and localization techniques with a larger corpus of performance bugs to evaluate against, we conduct a study of more than 700 performance bug fixing commits across 13 popular open source projects written in C and C++ and investigate the relative frequency of bug types as well as their complexity. Our results show that many of these fixes follow a small set of bug patterns, that they are contributed by experienced developers, and that the number of lines needed to fix performance bugs is highly project dependent.

The results from the research was published in Proceedings of the 30th International Symposium on Software Reliability Engineering (ISSRE 2019) [146].

The research is connected to task T1.3: Security by design as it explores performance bugs empirically in real-world applications.

5.10 Gyro: A Modular Scale-Out Layer for Single-Server DBMS

Contact: Habib Saissi (Lancaster/TUD)

Scaling out database management systems (DBMSs) requires distributed coordination, which can easily become a bottleneck. Recent work on speeding up distributed transactions has addressed this problem by proposing scale-out techniques that are deeply integrated with the concurrency control mechanism of the DBMS. This paper explores the design of modular coordination layers, which encapsulate all scale-out logic and can be applied to scale out any unmodified single-server

DBMS. It proposes Gyro, a modular coordination layer that runs on top of a collection of single-server DBMS instances and interacts with them only through their client interface. Gyro distributes the load by ensuring that as many requests as possible are executed by only one DBMS instance. Our experiments show that modular distributed coordination is practically viable and can be much faster than traditional distributed transaction protocols using two-phase commit.

The results from the research were published in proceedings 38th International Symposium on Reliable Distributed Systems (SRDS 2019) [121].

The research is connected to task T1.3: Security by design and explores the design of modular coordination layers in database management systems.

5.11 Analyzing and Improving Customer-side Cloud Security Certifiability

Contact: Yiqun Chen (Lancaster/TUD)

Cloud services have become popular as an effective form to outsource computational resources. While providing cost efficiency on the one side, this outsourcing also causes a certain loss of control over the computational resources, which makes security risks difficult to predict and manage. To address such concerns, security service level agreements (secSLAs) have been proposed as contracts between Cloud service providers (CSPs) and Cloud service customers (CSCs) that cover security properties of Cloud services. SecSLAs cover a variety of different security properties, ranging from the availability of encrypted communication channels for accessing Cloud resources to the timely detection and removal of vulnerabilities in the CSP's infrastructure. As previous work has shown, and as is evident for the example of timely vulnerability removal, not all of these security properties can be assessed by the CSC, which limits their utility as a contract basis. In this paper we propose a new monitoring framework for Cloud services to support the monitoring and validation of security properties on the customer side that require infrastructure-internal knowledge. To obtain the security properties to be monitored by our framework, we have manually investigated 97 different quantifiable properties in 5 standards from both industry and academia. We identified only 21 measurable properties from those standards, out of which we implement measurements for 13 representative ones and evaluated our measurements on the OPENSTACK platform.

The results from the research were published in proceedings of the 9th IEEE International Workshop on Software Certification (WoSoCer) [150].

The research is connected to task T1.3: Security by design and explores certifying the Cloud services to justify the Cloud service provider's trust level.

5.12 Extracting Safe Thread Schedules from Incomplete Model Checking Results

Contact: Neeraj Suri (Lancaster/TUD)

Model checkers frequently fail to completely verify a concurrent program, even if partial-order reduction is applied. The verification engineer is left in doubt whether the program is safe and the effort towards verifying the program is wasted. We present a technique that uses the results of such incomplete verification attempts to construct a (fair) scheduler that allows the safe execution of the partially verified concurrent program. This scheduler restricts the execution to schedules that have been proven safe (and prevents executions that were found to be erroneous). We evaluate the performance of our technique and show how it can be improved using partial-order reduction. While constraining the scheduler results in a considerable performance penalty in general, we show that in some cases our approach somewhat surprisingly even leads to faster executions.

The results from the research were published in International Symposium on Model Checking Software (SPIN) [84].

Our works contribute to pilot in T2.5: Security of Unmanned Aerial Systems (UAS) by offering extensive support for this pilot in streamlining the complex operations in diverse settings. For instance, the separation of safety-critical assets and mission-critical assets is a common practice in the aviation industry, however, it is impractical to replica the entire ecosystem necessary for the proper operations of the UAS. Therefore, the methodologies developed in T1.3 is useful in determining the dependencies among the operations apriori.

5.13 Detection of Common Vulnerability Exposures (CVE)

Contact: Jean-Yves Marion (UL), Werner Laurent (CYD)

We are working on an original method to detect known vulnerabilities (CVE). Nowadays, an issue is to know which binaries are used as a component of software or simply which binaries are in a system. This raises the questions: (1) Which is the compiler name and its version given binary has been compiled with? (2) Are there binaries containing publicly disclosed cybersecurity vulnerabilities and exposures (CVE)? Indeed, a vulnerable binary may be the target of an attack. As a result, it is critical to know what is the nature of running binaries in a system in order to secure it.

About the first question, in order to identify the compiler that was used to obtain a binary file, we exploit Gorille engine based on morphological analysis (see section on malware analysis). For this, Gorille is now used to analyze binary files obtained from various compilers and to find similarities and dissimilarities. Then, Gorille is used as a classifier which outputs the correct compiler name with the correct options. This was experimented with 8 different versions of `OpenSSL` library complied with different compilers and options. Then, we retrieve the name of the compiler and the option used from a single compiled function of `OpenSSL` like AES.

Regarding the second question, we developed a new method based on a Random Forest classifier/detector. This is a join work Cyber-Detect and LORIA - Lorraine University.

The proposed solution works as follows: For a given CVE, the learning vector is defined as a sequence of 0 and 1 representing the presence or not of the characteristic hashes of the VEC (sorted in lexicographical order) and ordered by the corresponding UR, and has the same structure as the lines of the learning matrix.

The Random Forest classifier/detector is trained with a set of CVE compiled with different compilers with different options and different operating systems.

At the end of the training, the Random Forest, which for a vector provided as input, will return the value 0 (patched), 1 (vulnerable), or 2 (neutral), i.e., not affected by the CVE.

The analysis of a program file for a given CVE is carried out in three steps:

1. Gorille computes the hash of all possible vulnerable part of the program.
2. The hash is correctly formatted for the Random Forest classifier/detector.
3. The vector obtained in the previous step is the input of the Random Forest classifier/detector.

This approach has been implemented as a proof of concept and experimented with 11 CVEs on 139 binary files of a CentOS distribution. Results go in the right direction but there is still room for improvements. This proof of concept has been made by a cooperation between Lorraine University and the start-up Cyber-Detect. A user-friendly interface was developed by Cyber-Detect in order to easily see whether or not a binary file contains a CVE and what is the criticity level of vulnerability. This work has been presented in a workshop on cybersecurity in Kyoto, Japan (April 2019).

5.14 Auto-encoding Robot State against Sensor Spoofing Attacks

Contact: Sean Rivera, Antonio Ken Iannillo (SnT)

In this research, we focus on Light Detection and Ranging (LiDAR) systems. These systems are largely employed to achieve high positioning resolution, in substitution or support of other approaches that locate the robot in indoor environments. Any alteration of these sensor data can silently force the robot to initiate dangerous maneuvers for itself and the environment in which it operates in. Regardless of the purpose of the robot, sensor quality and robustness are highly requested to perform eventual mission- and safety-critical operations.

The main idea is to detect anomalies in the lidar data through an auto-encoder architecture, i.e., an artificial neural network which aims to learn efficient data representation in an unsupervised fashion. Initially the model is trained and the reconstruction error signals is analyzed. This should serve to create thresholds that will be used in the detection. In operational mode, the robot uses the encoder to compress the sensor input data. Then, the decoder is used to compute the reconstruction error and compare it with the thresholds for anomaly detection. This compressed data can also be efficiently sent to other robots or the fog/cloud for complex fingerprinting and overall monitoring. Then, the auto-encoder architecture splits and processes different time windows to more efficiently detect changes. This approach adds a temporal correlation to the normal spatial correlation of our solution.

Rivera et al. implemented this idea in a workshop paper at the 4th Workshop on Reliability and Security Data Analysis (RSDA). We implemented the anomaly detector and evaluated it against several types of spoofing attacks comparing four different compression rates for the auto-encoder. We first gathered the lidar sensor input data from two different simulations. Then, we duplicated them into several datasets and injected in each dataset a different spoofing attack at a random time. Our approach is effective with a 99% True Positive rate and a 10% False Negative rate for 83% space-saving, decreasing to a 50% False Negative rate while maintaining the 99% True Positive rate with the 11% space-saving. Detection is more effective with the 83% space-saving, however, the 41% space-saving could handle almost all of the same attacks while using half the data.

Detecting anomalies is robots is an important topic nowadays since their usage in civil, business, and military environment. In particular, Airbus is leading the implementation of a use case for UAVs. If implemented in these devices, the anomaly detection mechanism can provide a further level of security for this use case.

6 Data/Application-Centric Security (T1.4)

Task 1.4 (T1.4) of the CONCORDIA project is concerned with data and application-centric security. Figure 9 presents a general framework that was proposed by TUBS (as the task leader) to organize the research efforts in this task as well as to show the possible interaction with the different tasks in the other WPs. T2.1 (Telecom pilot), Task T2.2 (Finance pilot) and Task T2.4 (e-Health pilot) have agreed to adopt this framework. These pilots will contribute to the framework by providing (raw) data, which will be processed by different analysis mechanisms to detect any Indicator of Compromise (IoC) and to determine the credibility and seriousness of the potential threat. Later, we share information related to these IoCs with the CONCORDIA central threat intelligence platform that will be provided by T3.1. Within this framework, we are putting more emphasis on cloud security to solve the next issues:

- How to protect (big) data before and after storing it on the cloud.
- How to protect the cloud services themselves.

Another critical question is how to perform data behavioral analysis on the collected data to detect suspicious behaviors or attacks. Forensic data visualization is one of these mechanisms which we promised to use to detect malicious activities. However, we have taken a strategic decision that we would not put too much effort in this direction at this early stage of the project. We decided to wait until we collect enough raw-data and IoCs.

In the following sections, we report on the research activities performed within T1.4. Section 6.1 presents a solution for protecting cloud services using automated security enhancement strategies, with a particular focus on challenges related to the migration of resources composing these services. Section 6.2 discusses the continuous certification of composed cloud services as well as the cost prediction

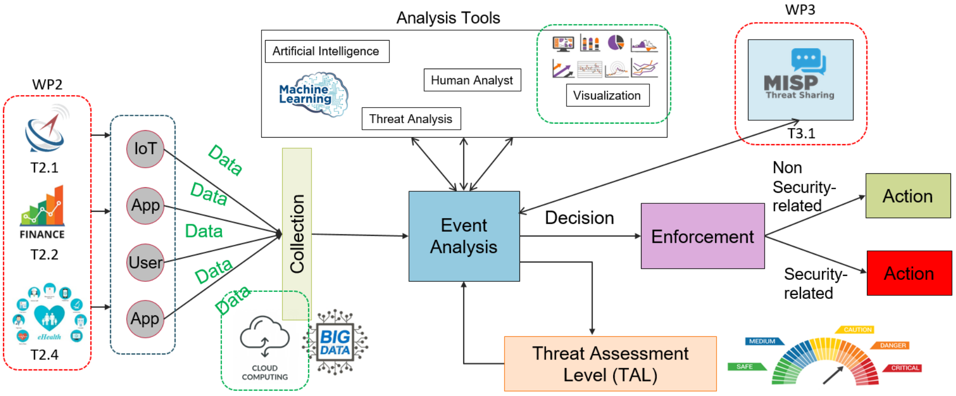

Figure 9: Framework for data-centric security

of this operation. Section 6.3 investigates the feasibility of the deployment of a 5G network core into OpenStack cloud in containerized environment. Section 6.4 discusses the use of blockchains to manage inter-organizational process execution and secure data-sharing. Section 6.5 addresses the problem of performing analysis on privacy information between organizations which cannot share their data by using blockchains. Section 6.6 presents a mechanism to enable the distribution and collecting of sensitive data securely using a policy-based trust management schema. Section 6.7 discusses a protocol which support transferring sensitive data between mobile application and trust execution environment in which the data can process securely. Section 6.8 presents a framework for creating and managing cyber insurance policy for cyber-systems. Section 6.9 presents a solution for ensuring the secure data exchange of IoT devices and applications by using cellular identity federation. Section 6.10 presents another solution for secure data exchange of IoT devices by using identity provider and management system. Section 6.11 presents a 5G smart home solution that could be used as a testbed for collecting the data to be used by the various analysis tools (e.g., ML) to detect the anomaly behavior of the IoT devices. Section 6.12 performs an overall risk assessment, providing the basis for discussions of security-relevant comparisons of Remote Postal Voting to Remote Electronic Voting or Internet Voting. Section 6.13 discusses the use of graph database mechanisms for cyber network analysis and visualization.

Figure 10 indicates how work done in T1.4 relates to the CONCORDIA pilots.

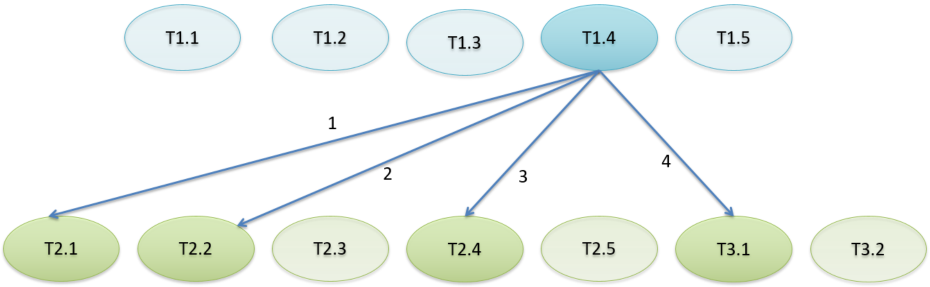

Figure 10: Links between task T1.4 and the CONCORDIA pilots

1. Papers [39, 38, 124, 36] have a strong connection with T2.1. These paper were published as a result of collaboration between OsloMET and TELENOR (T2.1 pilot leader). Work published by TUBS [109] and UMIL [9] are related to T2.1 too.

2. The secure data exchanging protocol which was proposed by SnT [29] has potential for the T2.2 (finance pilot). Also, there is a plan to have a collaboration between UP and T2.2 to integrate the UP proposed solution [75] with in the finance pilot.

3. Work published by UI [119] and UP [75] have potential for T2.4 (e-Health).

4. Lastly, as we have mentioned before, all the collected information about the IoCs will be shared using the T3.1 threat intelligence platform. TELENOR and TIM will establish and support the collaboration between T1.4 and T3.1.

6.1 Automating Security Enhancement for Cloud Services

Contact: Remi Badonnel, Olivier Festor (UL)

The design of value-added cloud services has been leveraged by the deployment of large data centers providing multiple and heterogeneous computing resources (software components, virtualized hardware equipments) that can be dynamically composed. These resources available in a metered manner are contributing to lower the operational costs of infrastructures and their services. Orchestration languages, such as as the TOSCA language (Topology and Orchestration Specification for Cloud Applications) [101], facilitate the specification and implementation of elaborated services from these resources. As these ones are available through broad network access, they constitute an attractive target that is concerned by a large variety of security attacks. These resources are often distributed over several cloud providers, and shared amongst different clients based on multi-tenancy models, causing isolation issues. They are also more exposed to service disruption attacks due to their on-demand self-service provisioning, and are subject to changes and updates over time. The dependencies amongst them may also permit their access and alteration by unauthorized parties due to improper access control. Therefore, it is of major importance to support, from a security viewpoint, the changes that may occur on elaborated cloud services and their resources.

The objective of this work is to elaborate automated security enhancement strategies for protecting them, with a particular focus on challenges related to the migration of resources composing these services. Let consider that an elementary resource is migrated from one cloud provider to another one, or is migrated to a different infrastructure of the same cloud provider. This contextual change may impact on the security of the elementary resource, but may also in turn impact on the whole orchestrated service built from this resource. In that context, automation methods and algorithms are required to dynamically adapt security mechanisms by taking into account the contextual constraints, and the specification of elaborated services. We are considering two different categories of complementary mechanisms to be combined in order to support such automated protection.

The first category corresponds to endogenous mechanisms, and consists in reducing the exposure to attacks, by modifying internally the cloud elementary resources. These mechanisms consist in hardening the configuration of virtualized resources at different levels, and they may be driven by security information data sources in order to prevent the exploitation of vulnerabilities. For instance, security management approaches based on the Open Vulnerability and Assessment Language (OVAL) [98], part of the Security Content Automation Protocol (SCAP) language family [141], consist in formalizing and avoiding vulnerable configurations that are described in the form of a logical combination of configuration conditions that if observed on one target resource, then such vulnerability is present on the considered system [97]. The hardening of resources may also go far beyond based on the usage of specific virtualization techniques. These techniques typically contribute to provide a better isolation [13], and to minimize the components that may be critical in the resources. At the extreme case, the use of unikernel resources may serve as a support to the generation of lightweight virtual machines embedding a single application, and its minimum runtime dependencies [34]. The purpose is to strictly contrain the resources to a minimal configuration, but this may lead to rethink the configuration management lifecycle, with a configuration change driving the re-building of a cloud resources. These endogenous techniques may not be sufficient due to the operational constraints or to the lack of patches, and may be restricted due to the environment offered by cloud providers.

The second category corresponds to exogenous mechanisms, and consists in complementing the cloud elementary resources based on external security functions. The programmability of network resources brings flexibility to the deployment and adaptation of such mechanisms. In particular, it enables to dynamically building and orchestrating chains that are composed of different security functions. These security functions may typically include firewalls, intrusion detection systems, and data leakage prevention mechanisms. They may be provided as middleboxes using network function virtualization, or be directly implemented on the software-defined networking (SDN) layer [60]. For instance, some efforts, such as Flowtag [44], have focused on extending middleboxes to make them SDN-capable, and using tags in network packets for determining how to process the corresponding traffic, so that it follows specific security treatment before reaching a given resource. The SIMPLE architecture [14] also defines a policy enforcement layer based on software-defined networking and flow correlation for middlebox traffic steering. High-level languages, such as the Frenetic family of languages [47], constitute also an important support to the functional specification of security chains and their implementation. They offer modular constructs that facilitate compositional reasoning, and are then translated into low-level configuration rules interpretable by programmable switches. The criticality of these chains and their functions requires to guarantee their consistencies through verification methods. For instance, the Kinetic language [69], based on finite state machines (FSM), has been proposed to automatically verify the correctness of such chains with respect to user-specified temporal properties. SMT solving and model checking techniques have also been experimented with that respect in [126]. These methods however do not take into account the configuration of resources, nor the orchestration that they may be part of, and are not correlated to endogenous mechanisms.

The considered automation strategies are directly related to the security mechanisms that are developed in the context of tasks T1.2 and T1.3, and are partially driven by the different alerts provided by the security monitoring platforms.

6.2 Security Assurance for Cloud (Composite) Services

Contact: Claudio A. Ardagna, Marco Anisetti (UMIL)

The maturity reached by cloud computing has fostered the implementation of a number of distributed infrastructure, platform, and application services available worldwide. Current trends in software distribution and provisioning envision services made available as commodities over distributed systems including the Internet or the cloud marketplace. At the same time, the trend towards coarse-granularity business services, which cannot be managed by a single entity, resulted in several approaches to service composition that maximize software re-use by dynamically composing single services on the basis of their functionalities. A major challenge faced by distributed service-based systems deployed on the cloud goes beyond the ability to guarantee the functionality of composite services, and must consider the importance of guaranteeing stable Quality of Service (QoS) in the form of non-functional properties requirements such as security, performance, and trust. Service compositions need to guarantee optimal and verifiable properties, managing different events that might change their structure such as component relocation, substitution, malfunctioning, versioning, adaptation. Continuous monitoring and verification of service non-functional properties is needed and usually achieved by means of assurance techniques. Recently, certification-based assurance techniques have been introduced to guarantee stable QoS in the cloud. They are based on continuous collection of evidence on the behavior of the system, which is used to verify whether the considered system holds a specific (set of) non-functional property and award a certificate proving it. To this aim, distributed agents are instrumented to connect to different endpoints in the cloud and retrieve evidence used to evaluate the non-functional status of the target cloud-based system. Current certification techniques mostly focus on the certification of single-service systems and often do not consider the cost of maintaining stable QoS. Even worse, a trend in service composition is to provide an ad hoc composite service for each request, with high costs on the cloud providers (CPs).

Our work in the first year of the project has been grounded on our past works on Certification of Cloud Services and focused on extending it towards the continuous certification of composed cloud services, on one hand, and the prediction of costs for the continuous service composition certification, on the other hand [7]. The idea is to face two colliding requirements. On one side, there is the need to guarantee non-functional properties of a service composition. This is a challenging task that requires continuous evaluation of compositions at cloud-provider side, to accomplish the dynamic and evolving nature of the cloud. On the other side, there is the need to take the costs observed by cloud providers for certified composition management under control. These costs, in fact, rapidly increase because the costs of continuous certification and verification become substantial. Current research on cloud computing has privileged solutions minimizing costs on the final users, neglecting the costs on the cloud providers that often represent a major source of fee increase. Our service composition is driven by certificates awarded to single services (e.g., using our past technique for single service certification) and by a fuzzy-based cost evaluation methodology, and assumes certified properties as must-have requirements for service selection and composition. This methodology aims to decrease the costs of cloud providers, also analyzing those costs introduced by the need of keeping the composition continuously monitored and certified. It contributes to the resolution of the long-standing problem of managing non-functional properties of distributed applications and composite services in a cost-effective way. It provides an approach that effectively relocates and refines service compositions in the cloud at run time guaranteeing stable QoS.

In addition to this work, in the first year, we have also investigated methodologies and approaches for building Certified DevOps pipelines of cloud microservices [9]. The idea is to produce a more advanced certification methodology for microservices in cloud that considers the entire development process in the framework of modern CI/CD pipelines. Our contribution in this area consists of the integration of a framework for assurance checks within a DevOps pipeline. The framework is capable to continuously and semi-automatically measure assurance metrics, which support certification activities.

This effort, carried out mainly in Task 1.4, is also relevant for activities in Task 1.1 aiming to evaluate the assurance of IT systems including smart sensors and devices. A proper assurance solution is also important for all pilots providing functionalities for evaluating the security status of a system and its level of compliance to predefined rules. In particular, with TELENOR and Pilot "Threat Intelligence for the Telco Sector", we started investigating the possibility of applying the proposed assurance techniques to 5G communications. A preliminary evaluation of a possible collaboration in this context has also been investigated with TUBS, to the aim of applying the proposed assurance framework in the context of embedded composite services.

6.3 Connecting Remote eNodeB with Containerized 5G C-RANs in OpenStack Cloud

Contact: Bruno Dzogovic (OsloMet), Thanh van Do (TELENOR)

5G mobile networks are softwarized and virtualized networks consisting of multiple virtual network functions (VNF) which are hosted in different data centers. In this work [39], we explore the feasibility of deployment of a 5G network core into OpenStack cloud in containerized environment. The Cloud-native Network Function (CNS) of the 5G core as a precursor for network slicing is achieved via instantiating the core elements in Docker containers in OpenStack, where the CNS is provided via secure integration between the OpenStack's Neutron networking module and the container in which the 5G elements are running. In order for this to be feasible, there has to be a way how to make the containers behave as if they are virtual machines running directly through the Neutron networking in OpenStack and exploit the OvS SDN. For this purpose, there is a special plugin called Kuryr, which has the function of making the Docker container behave as if it's part of the Neutron virtual router, with possibility of obtaining even a floating IP (public), while having its own MAC address.

Furthermore, the security part comes with the usage of tokens. If we want the container to authenticate to the OpenStack user, within the specific project and group where the VM belongs, then a specific token is issued for the container to attach to the Neutron virtual routing domain. This is explicated in one of the paragraphs, or more precisely as it's stated: "It provides access to the Neutron networking with the specific user that is assigned to (in this case Admin, within the Admin network, the same project name and relevant password). Notably, in this version of OpenStack, the actual version of Keystone is v2.0, that has slightly different API than v3.0 and uses less authentication parameters, disregarding the need for the "project_domain_name", "project_domain_id", "user_domain_name" and "user_domain_id" fields." To better illustrate the method, this is represented as in Figure 11.

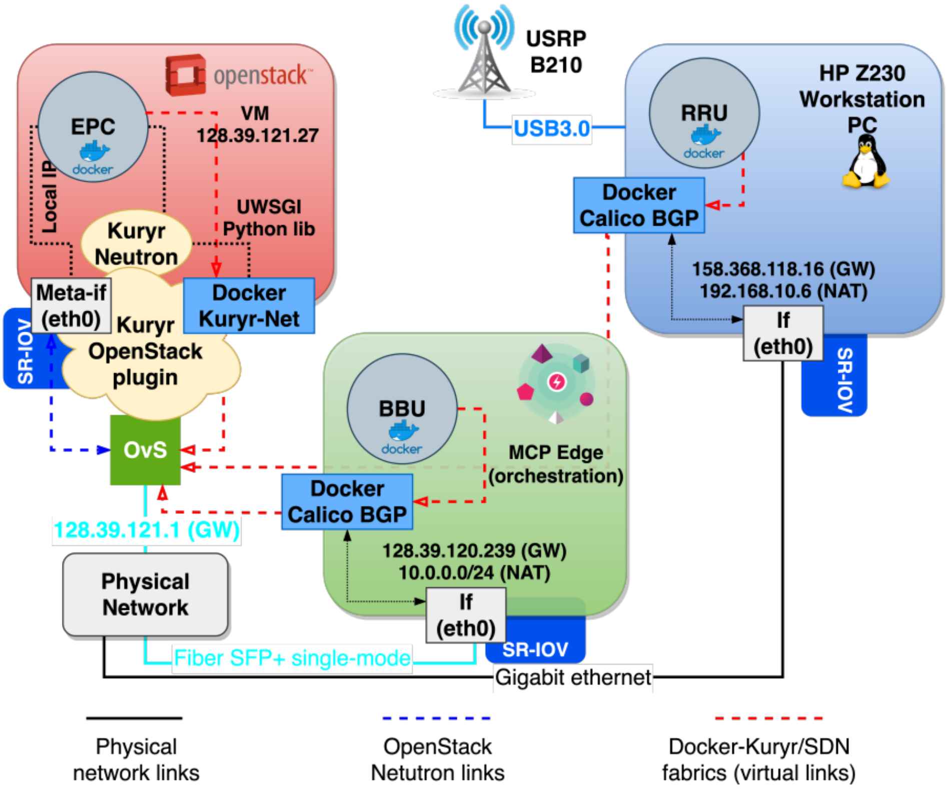

Figure 11: 5G4IoT Lab Infrastructure at the Oslo Metropolitan University

From here, we can proceed with the fact that the plugin utilizes the OpenStack's Keystone module for handling the users and authentication procedures. These are the whole cloud-related security details that this paper focuses on, where the main topic is the actual feasibility of achievement of connection with remotely-instantiated 5G core networks. This is a security at a lower level and is not the 5G-related security. When it comes to the 5G security, the users are authenticated in the core network via SIM authentication, which is not the matter of discussion in this specific research paper. That would be a bit more in-detail about this topic.

An earlier 5G testbed is built at the Secure 5G4IoT lab operated by TELENOR, OsloMet and Wolffia. This 5G testbed will pave the way for the TELENOR pilot (T2.1).

6.4 Blockchain-based Secure Execution of Collaborative Process

Contact: Barbara Carminati (UI)

Today we are living in an era where technology makes more easier the collaboration in many different applications domains (e.g., social networks, IoT, Big Data analytics). A relevant representative of such class of collaborative services is one of inter-organizational processes, where an organized group of joined activities is carried out by two or more organizations to achieve a common business goal. In this picture, collaboration brings many benefits in terms of resource usage optimization, improved QoS, knowledge sharing, and so on. However, it also poses serious security and privacy threats to the data each organization exposes during the process execution. This risk is mainly due to the weak trust relationships that may hold among the collaborating parties, which result in a potential lack of trust in how data/operations are managed.

A promising way to deal with this lack of trust is leveraging on blockchain to support secure inter-organizational business processes. In general, a blockchain is a distributed data structure, replicated and shared among members of a network, acting as a distributed ledger, used to keep track of every exchange of resources or assets between participants of a network. These changes are recorded into transactions, batched into time-stamped linked blocks, forming the so-called chain of blocks. Transactions are inserted into blocks only if they are considered valid by the network participants. Transaction validation is reached through a distributed consensus protocol that, in general, is considered secure if the majority of network participants are honest. An important aspect, when leveraging on a blockchain, is that the computation involved in transaction validation can be encoded into predefined programs. For instance, the well-known Bitcoin framework provides a set of programs tailored for cryptocurrency management. In contrast, more recent blockchain frameworks, like Ethereum, support the idea of running arbitrary user-defined programs, called smart contracts. The idea is to translate contractual clauses into smart contracts stored and executed on the blockchain. As such, blockchain can be seen as a distributed ledger storing results (i.e., transactions) of (smart) contracts whose correct evaluation have been validated by network participants. This view brings several benefits. The first is that it does not require a central trusted party to validate the correct execution of contracts. Moreover, the obtained transparency well overcomes the lack of trust among parties involved in the contracts. As a result, blockchain/distributed ledger have gained increasing interest from companies aiming at encoding with smart contracts their processes and collaborations.

In line with this trend, several works have recently investigated how to exploit blockchain technology to manage inter-organizational business process execution [78, 83, 143], where the most relevant issue is the lack of mutual trust among parties. The key idea of all these approaches is that organizations expose their services to be directly invoked by smart contracts, having the smart contract playing the logic of coordinating and monitoring the overall service composition. Distributed consensus ensures the correctness of smart contract execution, aka their transactions, ensuring thus the correctness of collaborations among different organizations. However, each data in the blockchain, e.g., stored on the chain or used in a smart contract, is public. Indeed, the distributed consensus protocol relies on the fact that validators have to re-execute the original smart contract to validate the resulting transaction. This poses severe consequences for the confidentiality and privacy of involved data, and thus to the feasibility of the integration of business process execution with blockchain (see [27] for a more detailed analysis).

To cope with these confidentiality requirements, in [28] we proposed a solution where sensitive data contained in the smart contract is encrypted in such a way that can be consumed (aka decrypted) only by authorized organizations (aka organization that indeed needs it for the completion of its task) by, at the same time, allowing simple computations over it. At this purpose, given a smart contract SC implementing the inter-organizational process, in [28] we generate a new Confidential Smart Contract - CSC, where all sensitive data contained in smart contract variables are encrypted exploiting PKI encryption, to enforce access only to authorized as well as homomorphic encryption scheme. We see several applications of the framework proposed in [28] to CONCORDIA tasks, some of them have been investigated in the first year of the project. A first relevant application is in the context of IoT devices (i.e., Task 1.1), where many of the most interesting services are offered through a collaboration of IoT devices. To deploy collaborative IoT services, the first step is to discover the appropriate set of devices that have to collaborate to reach the target goal. At this purpose, this discovery process has to be driven by a set of search requirements on device functionalities (e.g., temperature sensing) and device features (e.g., device location and/or maximum latency). A key point is therefore to have assurance that all requirements have been correctly considered during the discovery process. Indeed, an untrusted discovery process might fail in evaluating the specified requirements in several ways. It might return devices that do not satisfy the requirements (e.g., not placed in the required location) as well omit those that indeed satisfy them, with the final result of selecting IoT devices not implementing the required service. To avoid this situation, it is required that the discovery process provides a proof of the fact that it has retrieved all and only those devices satisfying the specified search requirements.

At this purpose, in [119], we exploited blockchain to ensure the correctness of the IoT discovery process execution. More precisely, we propose to have an IoT discovery process, driven by the specified search requirements, implemented via a smart contract. Blockchain network validation of the smart contract execution assures that all search requirements have been correctly evaluated. Furthermore, similar to the framework proposed in [28], we designed an approach that allows smart contract execution by preserving, at the same time, the confidentiality of involved sensitive data (e.g., device location).

Another relevant application, on which we are currently working on, is related to Task 1.4, where the blockchain framework proposed in [28] could be seen as the building block for a more general secure and controlled data-sharing among organizations, where data release is driven by an undergoing collaborative process.

6.5 Using Blockchains to Enable Big Data Analysis of Private Information

Contact: Kostas Lampropoulos (UP)

Big Data analysis can provide valuable insights and solutions across multiple sectors (Medical, Finance, Telcos, etc.). However, access to real, large datasets is often very difficult due to privacy reasons and the reluctance to share data. For example, telecommunication providers have large volumes of data with cybersecurity events and attacks they suffer. These datasets contain highly sensitive data (due to business, security, reputation etc.) and thus the providers are very reluctant to share with external parties. At the same time though, in order to produce high quality products that can help Telcos protect themselves against attacks, one must be able to process such large datasets (of security events and attacks) from as many different sources as possible (many Telco providers). The problem also applies in multiple other sectors as well (e.g., process medical records in health sector). The aim of this work is to provide a blockchain-based solution capable of facilitating big data analysis on private datasets located across different administrative domains.

Despite the fact that Big Data analysis and blockchains are subjects which are undergoing intense study, to the best of our knowledge there isn't any other work similar to our solution. Research efforts that examine both Big Data and blockchains, either describe the general capabilities of the two emerging technologies or examine cases which are not focused on the analysis of Big Data but in other processes like authentication and access control.

Our proposed solution [75] is based on the observation that companies which want to benefit from Big Data analysis on external private datasets, do not need to access these datasets. Instead, what they need to execute a specific code/algorithm (analysis) on them and collect the results. Based on this observation and with the use of blockchain technologies (Hyperledger Fabric) we designed a solution based on which a company can create a smart contract with the analysis (code/algorithm) that wants to perform on a private dataset and then request its execution on selected nodes (nodes that host these private datasets). Eventually, the company will receive the desired results without accessing the data. To make the whole process secure, we selected the Hyperledger Fabric due to its multiple security and privacy features. These features are:

- *Permissioned DLT technology*: Allows controlled access to the network.
- *Chaincode*: Smart contract support.
- *Channels*: Each channel supports a separate private ledger.
- *Private data collections*: Private data collections are datasets that are not stored on chain and can only be seen and accessed by authorized parties.
- *Endorsement and validation policy enforcement*: This feature gives the ability to define a set of rules/policies about who can run or endorse a specific chaincode.

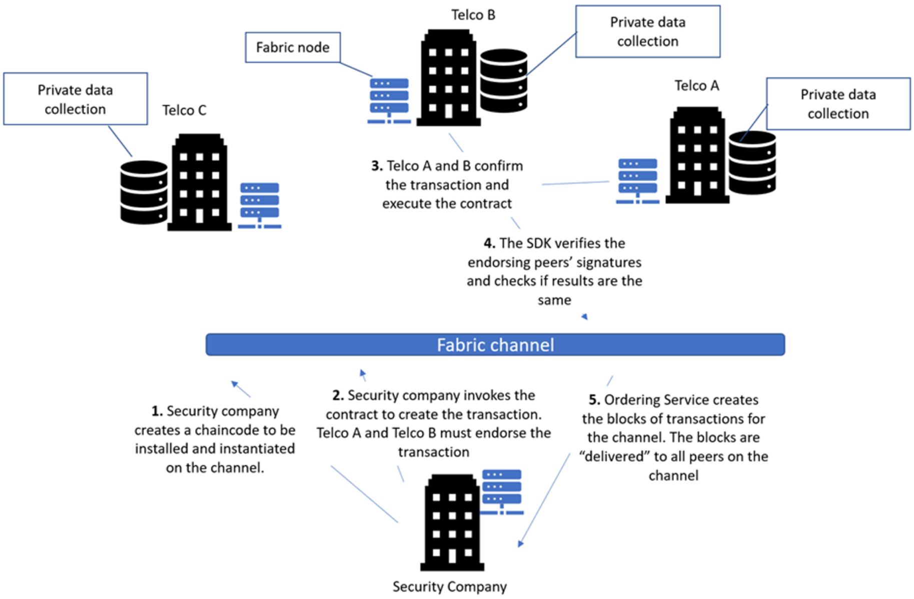
Telco B
Fabric node
Private data collection
Private data collection
Telco C
Telco A
Private data collection
3. Telco A and B confirm the transaction and execute the contract
4. The SDK verifies the endorsing peers' signatures and checks if results are the same
Fabric channel
1. Security company creates a chaincode to be installed and instantiated on the channel.
2. Security company invokes the contract to create the transaction. Telco A and Telco B must endorse the transaction
5. Ordering Service creates the blocks of transactions for the channel. The blocks are "delivered" to all peers on the channel
Security Company

Figure 12: Blockchains for Big Data Analysis on private datasets

Figure 12 presents how a security company can benefit from our solution and perform Big Data analysis on private datasets which are stored in two Telcos. Telcos

A and B hold large datasets of security related information (logs) and share these datasets via a bilateral agreement. Using our solution, a security company can create a smart contract asking for a specific analysis on their logs. The results are published as private data collections and only authorized nodes (security company) can access them. Also, a hash of the results is stored on the chain as proof. *Note 1*: Telcos A and B have the same datasets, and the SDK can cross reference the produced results in order to ensure that both Telcos executed the algorithm on the desired data. *Note 2*: In cases of datasets hosted by only one company, the added value of our proposed solution is the creation of a solid (immutable) proof of the results, which can be used as evidence in a legal dispute.

Within CONCORDIA, this work has attracted the interest of the Financial (T2.2) use case for both business and security scenarios. Furthermore, it is our goal to further explore its usefulness in other pilots like, e.g., the Mobility or e-Health pilots (T2.3, T2.4).

6.6 Policy-based Secure Data Sharing

Contact: Vassilis Prevelakis, Mohammad Hamad (TUBS)

As the use of technology increases, so will the amount of data it produces. At the same time, the need for protecting this data becomes a critical need, especially if we deal with confidential data (e.g., biometrics, financial, telecom, etc.). Many real attacks showed how hackers could reach this data and stole it (e.g., In July 2019, FBI arrested one person on suspicion 30 GB of Capital One credit application data from a rented cloud data server [2]). In many cases, the main reasons behind such successful attacks were that the data was stored in a non-secure repository without proper encryption or exchanged over insecure over-the-air channels. The situation even becomes worse when we know that the devices involved in exchanging such data, in most cases, are resource-constrained and hence are not in a position to employ traditional security protocols such as TLS and IPsec.

We present in [109] a new model for secure exchange and store information that is based on off-line symmetric key encryption of the data and the use of policy-based credentials for the release of the encryption keys. Fig 13 presents the proposed model. By using this model, the sender (data source) does not need to establish a direct network connection with the receiver (consumer, data analyzing process). Instead, it can employ a store-and-forward transmission method by storing the data in a repository (e.g., cloud) and authorize a remote consumer to use it. End-to-end data security is achieved by encrypting the data before the transfer. Any bulk encryption algorithm can be used. However, we still need to address the problem of sending the encryption key to the consumer, especially if they are more than one, and their identity is not known in advance. Therefore we have adopted Aegis Secure Key Repository (ASKR) server, which used as a key repository as well as an enforcement point that controls access to the keys. Each time one sender

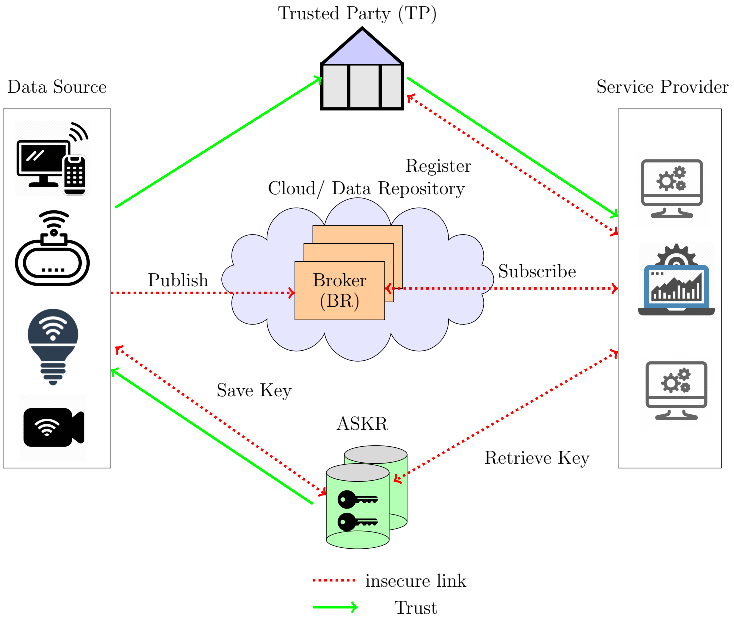

Figure 13: ASKR System Model.

needs to transmit the data securely, it saves the secret key, which is used to encrypt this data within the ASKR. Later, ASKR receives a request from one consumer to retrieve that key and use it to decrypt the data before processing it. The consumer needs to be authorized (trusted) to retrieve the encryption key. This authorization or the trust relation is implemented through a chain of credentials that delegates the trust from the data source to the data consumer by using a third trusted party, as shown in Fig 13. By using this mechanism, we eliminate the need for trusting the data repository. Note that even If ASKR is compromised, the intruder will only have access to the keys, not to the contents of the files. The keys, or key access credentials do not contain any information that may allow the file itself or its contents to be identified.

In addition to this work, we have also implemented a preliminary mechanism for securing higher-level applications through a robust, underlying, custom Linux kernel which could be used by cloud providers to build the virtual machines on that they offer to their clients [134]. This is ongoing research and will be investigated further in the future.

This effort is also relevant to Task T2.1 with TELENOR, where we have a plan to use their 5G test-bed to collect data and store it securely using our proposed mechanism.

6.7 Secure Exchange Between Mobile Application and Trusted Environment (SGX)

Contact: Federico K. Carvalho Ota (SnT)

Recurrently, events such as data breaches put the user's privacy on the spotlight. At the business level, evolution of data protection regulations (e.g., GDPR in Europe) are pressuring service providers but there is an urge. The lack of enhanced secure user authentication, specially in banking environments, increases privacy concerns. Some previous leakages of worldwide players, such as Facebook, show that the trade-off between security and convenience seems to always end in big losses, i.e., identify theft and confidentiality failure. If leaking personal conversations or pictures are considered big issues, the problem become even more critical when using sensitive personal data, such as biometric data. In the context of this paper, we focus exclusively on biometric data as a use case. Indeed, mobile banking is now a reality and so is the fingerprint reader embedded on most of the mobile phones. There is no doubt that biometric, implicit and continuous authentication, improve security of mobile applications. However, it is paramount to take into consideration the threats that a potential exposition of this data might cause. The main manufacturers of smartphones have implemented the extractor, matcher, database and decision system embedded in the operational system to enable collection of biometric data with the hardware sensor shipped within the device. In most of the cases, the embedded authentication method works fine. However, in case that a more complex process is required, the biometric data must be transferred from the device to an external server, which raises several concerns related to potential data breaches.

In this research, we propose a protocol to transfer sensitive data (e.g., biometric data) from a mobile device to a trusted execution environment, in which the data can be processed securely. Furthermore, when it comes to biometric systems, we must consider vulnerabilities that could potentially be explored by adversarial machine learning as the matcher is usually based on training models. What all these attacks have in common is that they are initiated from the data source because we consider, like most of today web applications, that the certificate authority of the service provider is not compromised. Indeed, certificate based mechanisms are usually provided to ensure that the requests are coming from an attested third party. However, when they come directly from the mobile apps it is not possible to assure the origin of the requests. This mainly happens because the certificate, usually pinned in the app code, can be retrieved by a reverse engineering process. In possession of the certificate, an attacker can build his own app to send malicious requests to the bank API. With the enforcement of the Payments Services Directive 2 (PSD2), banks are opening their API to third party. This threat is serious because it impacts security measure on HTTP requests. PSD2 shall increase the HTTP requests coming from different third parties applications making more difficult to decide in a short time frame if the request is legitimate. Then, if the data is processed in a non-reliable cloud environment, it is also necessary to protect the data from the provider in order to prevent unauthorized disclosure and tampering. Therefore, in order to provide a trusted execution environment, the Intel Software Guard Extensions (SGX) was chosen as the root of trust for the proposed protocol. In the server side, all operations are processed inside a SGX enclave to guarantee confidentiality and integrity.

Our contribution is a protocol to perform a secure exchange of data between a mobile application and a trusted execution environment, that has been presented at the 4th Workshop on Reliability and Security Data Analysis (RSDA) [29].

This research cover one of the research pillars of the project, namely the data-centric security. It improves the data protection when the exchange of sensitive information, in this case biometric data, happens among a mobile device and a remote server. Furthermore, this research can be used in the context of Finance sector pilot, where the KYC paradigm can be protected from attack on the communication side.

6.8 A Framework for Liability Based Trust

Contact: Sotiris Ioannidis (FORTH)

Security certification is a way to evaluate and certify than a cyber infrastructure is in line with security standards. The certification process results in a certificate, stating that an ICT system is compliant with a certification scheme, which defines the security properties and standards that should be complied with. The evaluation of the infrastructure can be either self assessed, i.e., the process is carried out by the owner of the infrastructure, or be conducted by a third-party organization, i.e., a certification lab. A notable certification scheme which spans certification process from the early stages of development up to the evaluation of the final infrastructure is the Common Criteria for Information Technology Security Evaluation [130]. Common Criteria evaluation results in a rating against an Evaluation Assurance Level (EAL). EALs are characterized by:

- The types of evidence that have been taken into account for evaluation (e.g., inspection, testing, formal verification).
- The agent who carried out the evaluation.
- The documentation/evidence that has to be produced and assessed for a successful evaluation.

This work focused on the development of a framework for creating and managing cyber insurance policy for cyber systems. Creating such policies will enhance the trust- worthiness of cyber systems and provide a sound basis for liability in cases of security and privacy breaches in them. The framework is supported by a platform of tools enabling an integrated risk cyber system security risk analysis, certification and cyber insurance, based on the analysis of objective evidence during the operation of such systems. The development of the insurance platform is build upon and integrating state of the art tools, methods and techniques. The development of the platform is driven by certification, risk analysis and cyber insurance scenarios for cyber system pilots providing cloud and e-health services. Through these, the platform addresses the conditions required for offering effective cyber insurance for interoperable service chains cutting across application domains and jurisdictions. The platform aims to tackle the challenges of offering cyber insurance for interoperable service chains cutting across application domains and jurisdictions.

During 2019 we developed the tools enabling an integrated risk cyber system security risk analysis.

This works is connected with with the data-centric security enhancing the trustworthiness of cyber systems and provide a sound basis for liability in cases of security and privacy breaches in them and with CONCORDIA T2.1 “Telecom Sector: Threat Intelligence for the Telco Sector” and T2.3 “Transport E-Mobility Sector: Security of the e-Charging Infrastructure” pilots.

6.9 Enhancing Security of Cellular IoT with Identity Federation

Contact: Bernardo Santos (OsloMET), Thanh van Do (TELENOR)

This work [124] presents a Cellular Identity Federation solution, which has been designed to both strengthens and simplifies the authentication of Internet of Things (IoT) devices and applications by providing single sign-on between the network layer and IoT applications. They are hence relieved of the burden of authentication and identity management, which could be both technically and economically challenging. The work aims at clarifying how IoT authentication can be skipped without compromising security. The proposed solution is described thoroughly, and the authentication process is depicted step by step. Last but not least is the comprehensive description of the proof-of-concept which shows the feasibility of the Cellular Identity Federation.

In the proposed Cellular Identity Federation solution (see Figure 14), the current 4G mobile network consisting of standardized network elements such as eNodeBs (base stations), MME (Mobility Management Entity), HSS (Home Subscriber Server) on the control plane and S-GW(Serving Gateway), P-GW (Packet Data Network Gateway) on the user plane) is now interfaced with the IoT Server such that information about successfully authenticated IMSIs can be shared with the IoT Server. Three new entities are introduced in the solution as follows:

- *The Identity Provider (IDP/IDMS)*: makes use of the authentication information from the HSS to carry out authentication of the IoT devices on behalf of the IoT Server. In our solution, OpenID Connect is selected and used because it is by far the most popular and simple standard.

- *The Secure Database (DB)*: To keep the security protection of the HSS at the same level as before, instead of introducing an IP interface on the HSS and allowing direct interactions with it from the IDP, we introduce a partially mirrored database, which gets transferred from the HSS only relevant parameters about successfully authenticated devices such as IMSI, IMEI (International Mobile Equipment Identity), MSISDN (Mobile Subscriber ISDN Number), PDP (Packet Data Protocol) type, PDP address (IP address), APN (Access Point Name), etc. These parameters will again be sent to the IDP for storage and use in the authentication of the IoT devices.
- *The OTP Generator*: To ensure that an IoT Device is really the one belonging to the IoT Owner, it has to prove that it actually is the device carrying a SIM that has been successfully authenticated by the HSS. For that, it is not sufficient that it presents its IMSI or IMEI to the IoT Server upon login because both IMSI and IMEI can be easily sniffed for replay. A more secure token is required. Upon receipt of authentication parameters of an IoT Device from the DB, the IDP will request the OTP Generator to produce a one-time password and send it to the IoT Device using the PDP address such that it can use it to login onto the IoT Server.

To clarify how an IoT device is authenticated let us now consider an IoT_Dev1 hosting a SIM1. The authentication process is as follows:

- At power on, the SIM on an IoT Device participates in the authentication of User Equipment.
- Upon successful authentication, the HSS stores the state of the IoT Device as registered and notifies the Replica Database that initiates the duplication and transfer of the parameters of the IoT Device to the IDP.
- The IDP stores the data and send a request for the generation of a one-time password to be sent to the PDP address of the IoT Device.
- The OTP Generator generates an OTP and sends it to both the IDP and IoT_Dev1.
- The IoT Client on the IoT Device fetches the OTP and presents it to the IoT Server upon login.
- The IoT Server redirects the IoT Client to the IDP.
- The IDP compares the presented OTP and if it matches with the stored one the IoT Client is considered authenticated and directed back to the IoT Server, which grants access to IoT Client. The authentication process is hence completed without active participation of the IoT Server which does not have to administrate passwords of its IoT clients while strong security is still ensured.

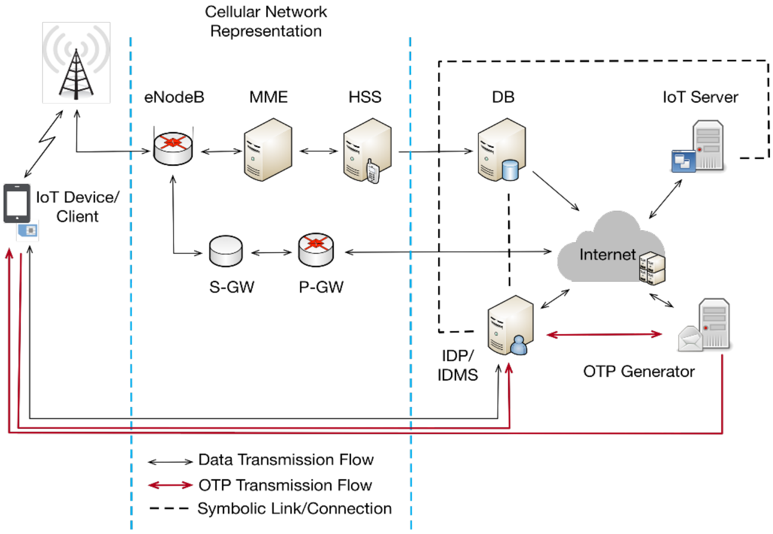

Figure 14: Overall architecture of the Cellular Identity Federation solution.

The proposed and implemented Identity Federation will enable strong authentication of IoT device using the strong SIM authentication. Further, the federation of the IoT IDs with the mobile IDs will enable the usage of data from both the IoT application layers and the network layers.

Thanks to the proposed Identity Federation, data from both the IoT application layer and the network layer can be used together in the anomaly detection allowing the detection of flooding attacks even before they are launched. This has been an important contribution to the Telenor pilot.

6.10 A Secure Authentication Mechanism for IoT Devices in 5G Networks

Contact: Bernardo Santos (OsloMET), Thanh van Do (TELENOR)

Upon the new paradigm of Cellular Internet of Things, through the usage of technologies such as Narrowband IoT (NB-IoT), a massive amount of IoT devices will be able to use the mobile network infrastructure to perform their communications. However, it would be beneficial for these devices to use the same security mechanisms that are present in the cellular network architecture, so that their connections to the application layer could see an increase on security. As a way to approach this, an identity management and provisioning mechanism, as well as an identity federation between an IoT platform and the cellular network is proposed as a way to make an IoT device deemed worthy of using the cellular network and perform its actions.

Our main objective is to have an authentication mechanism that can be applied to IoT devices so that their communications can be secure and be protected from possible exploits. This means that such a device has USIM capabilities and can directly use the known cellular network mechanisms.

It is important to provide an identity to each IoT device connected to the network so that they can be authenticated when trying to access the application layer (which in this case is where an IoT platform will reside). In order to do so, we are adding to our network architecture an identity management/identity provider server that will take care and manage the identities of all IoT (mobile) devices that are or will be registered in the network [36].

The Identity Provider (IDP) and Management System (IDMS) will be responsible for issuing the identities for every IoT (mobile) device that will be registered in the network. In order to have this link between the IDP server and the cellular network, it is necessary to achieve an agreement and a consensus as to how to identify the IoT devices so that they can communicate with the IoT platform in a more secure and trustable form. In order to identify and authenticate an IoT device, the IDP has to check if there's such a SIM registered in the cellular network (that is associated to an IoT device), by consulting the database managed by the HSS, which has all the records of registered devices. This access between the IDP and the HSS cannot have a direct repercussion on the normal behavior of the cellular network, so it cannot be a simple yet direct link between both components. To overcome this, an interface between the cellular network and the IDP is deployed - in the form of an Application Programming Interface (API) or by creating a secure replica of the HSS's database exclusive to the IDP's usage – so that it is possible to achieve a federation from the network side.

If there's a match, an identity will be issued for such device and will have access to the network and its resources. Fail to do so and if there isn't an alternative for a device to prove itself, an identity won't be issued. This identity is none other than a set of key-pair values that can be defined by an administrator, and by using the OpenID Connect protocol, it gives the versatility needed upon identity creation and enough flexibility to be used in a mobile/IoT context.

This work provides a strong authentication mechanism for IoT devices connected to a cellular network. Instead of using passwords, which is a weak form for authentication the IoT devices can now be strongly authenticated using the SIM authentication. This work is essential to realize the Telenor pilot "Preventing IoT flooding attacks on cellular network".

6.11 Bringing 5G Into User's Smart Home

Contact: Bruno Dzogovic (OsloMet), Thanh van Do (TELENOR)

A successful 5G Smart Home must be able to support 5G devices along with WLAN devices. This is not possible with the 5G specifications, and to remedy the situation, this work proposes a 5G Smart Home Solution that combines the 5G network slice concept with a femtocell home router, supporting both 5G and WLAN access technologies. While making use of standard technologies, the solution has proven to be efficient and reliable. For that purpose, in this work [38], we propose a femtocell network slice architecture for 5G, utilizing WLAN in order to connect the user's smart home with the distributed cloud services.

As there will be both WLAN and mobile network devices in the future home our Smart Home solution will have to able to support both types of devices. The 5G4IoT Smart Home can be typically realized by two network slices as follows (see Figure 15):

- *mMTC (massive Machine Type Communication) Network Slice*: OpenAirInterface deployment of a Core Network and eNB base station, powered by USRP B210 software-radio at Band 3 – 1.8GHz and Band 7 – 2.6GHz (mobile network).
- *WLAN Network Slice*: LWIP deployment of a containerized WLAN function for dual-connectivity architecture to multiple base stations simultaneously, powered by the same USRP B210 software-radio at 2.4GHz.

The femtocell home router offers both WLAN and 4G/5G connectivity in the home where the eNB macro-cell can struggle to reach the indoor devices. In this case, the back-haul link to the service provider is a fibre or cable line, which is routed through proprietary gateway. However, the backbone link can also be 5G fixed wireless if necessity arises, or if the mobile operator does not have a terrestrial deployment solution in the particular area. Additionally, the operator can offer local cloud services for specific applications and thus the users can access these with minimal latency and optimal performance. Such applications are usually smart-home related, as for the home appliances so as for the electric cars, gadgets or various contrivances. The femtocell home router should be able to operate as a multi-purpose device, namely having the following characteristics:

- Interfaces for fibre/Ethernet networking in order to connect to any external network gateway provided by the service provider.
- Wireless capability for both Wi-Fi and 4G/5G by using multiband antennas, as well as NB-IoT (NarrowBand IoT).
- Utilization of software-defined radio entities, i.e., being able to work independently on a generic Linux PC or as a standalone device, while connecting to the operator's core network.

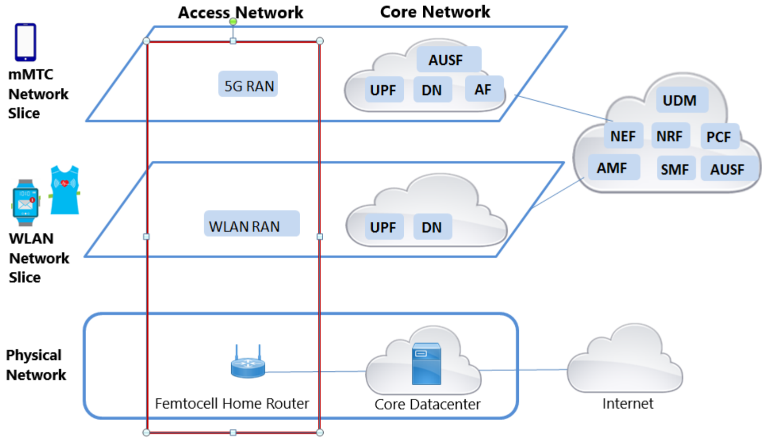

Figure 15: Architecture of the 5G4IoT Smart Home Solution

- Capability to communicate with AAA server (EAP-AKA for the WLAN and the DIAMETER authentication service in the core network of the 4G/5G).
- Capability to implement routing protocols such as BGP, OSPF, AODV.

The work realizes a 5G Smart Home Solution capable of supporting simultaneously 5G and WLAN, which ease the adoption of 5G as wireless infrastructure. This work enables the collection of data that can be used by the Machine Learning platform in the anomaly detection to combat IoT devices flooding attacks, which is a relevant for the Telenor pilot.

6.12 Postal Voting Process and Security Analysis

Contact: Christian Killer, Burkhard Stiller (UZH)

Around the globe, government services are becoming increasingly digitized [6]. Naturally, these efforts include electoral processes. In Switzerland, the federal government defined strategies enabling digitization for public authorities and processes, including Electronic Voting (EV) [132, 37]. Private companies collaborate with Swiss authorities to actively define standards across e-Government processes [137]. The Swiss EV typically refers to *Remote* EV (REV) carried out over the internet, which is also often referred to as Internet Voting (I-Voting) [72].

Even though the majority of Swiss citizens would prefer REV, a political position proposes a moratorium on EV in Switzerland [41]. According to their initiative [41], a REV system has to be ”*at least as secure as the current remote postal voting (RPV) system*”. Thus, the key question is: what exactly does a minimal level of security involve? Which security metrics and mechanisms are mandatory? In the general public perception, EV often provokes a fear of change, presuming the current RPV system to be mostly analog and tamper-proof. However, it can be argued that the current Swiss RPV system is already partially EV, since many steps already involve distributed electronic systems. Thus, defining and comparing the security properties of a REV also requires an analysis of the current RPV system in Switzerland.

Due to the federal and decentralized structure of Switzerland, each canton and municipality autonomously manages their respective jurisdictional electoral procedures. Cantons and municipalities execute a degree of independence in decisions on how to handle certain parts of the voting process. Therefore, the current RPV system in Switzerland is neither universally documented or specified, nor homogeneous across entities. This research, therefore, formalizes the Postal Voting Process Flow (PVPF) in Switzerland. The approach taken formalizes the PVPF in a step-based model, for which major assumptions made, such as trust, people involved, and technology applied, are made explicit, if known. Finally, this research performs an overall risk assessment, providing the basis for discussions of security-relevant comparisons to REV or I-Voting.

This research was published and presented at E-VOTE ID 2019 in Bregenz [67], as well as the Swiss Cyber Storm security conference in Berne [133].

6.13 Forensic Data Visualization

Contact: Martina Šestak, Muhamed Turkanović (UM)

Using graph databases as a fraud detection media, whereby exploring structural anomalies in subgraphs with the help of community detection algorithms, such as Louvain in order to detect fraudulent nodes. Also, using the generated information to prepare integrity constraint mechanisms, which can be set up in order to prevent further frauds. Graph database technology has seen an increasing usage in large enterprises and complex domains. Our research focuses on exploring graph database mechanisms, which can be used for cyber network analysis and visualization. Their underlying property graph data model enables complex analysis to be performed, such as root cause analysis of an attack and finding anomalous subgraphs within the network. Furthermore, network anomalies can be prevented by employing a constraint model able to notify users about anomalous behavior of network components.

Since one of the growing areas of graph databases application is fraud detection, the goal of our research was to explore how graph database technology and graph-based algorithms can be used to prevent and detect network anomalies in different domains (e.g., insurance frauds, cyber threats, etc.). We performed a literature review to gain insights into state-of-the-art techniques, methods and metrics used to analyze networks and detect anomalies as potential frauds. Our research showed that, in the context of graphs, frauds can be represented as subgraphs (motifs) exhibiting unusual behavior in terms of structural anomalies, such as higher density or node degree. So far, most research has been performed on the topic of anomalous subgraph detection, for which authors have used a variety of methods (classification, clustering, graph scan statistic, neighborhood analysis, etc.). Some of the modern Graph Database Management Systems, such as Neo4j, already include a set of pre-built graph algorithms, which can partly be used for the purpose. Specifically, we have tested the performance of community detection algorithms available in Neo4j (e.g., Louvain algorithm) on several real datasets, which required us to transform source data from CSV format to a property graph database model. Once the model was built, we were able to load data from source CSV files as nodes and relationships in a Neo4j database, and perform further analysis by visually querying the database through Neo4j browser interface. Furthermore, our research showed that community detection algorithms, such as the Louvain algorithm, are able to divide nodes into communities in a way that some communities have a higher percentage of fraudulent nodes (more than 80%), which proves that they can be used to speed up the anomaly detection process. Apart from graph analysis, we also tackled the challenge of preventing anomalies by using the integrity constraints mechanism available in graph databases. Specifically, we developed a constraints model flexible enough to specify higher order cardinality constraints (e.g., limit the number of given relationships considered as “normal” between any number of nodes of given types), which is to be published as part of our future work [139]. The results of the model performance testing show that it does not bring significant overhead to query performance time. Overall, our research results show that graph databases provide mechanisms, which can be used in different areas of cybersecurity, such as attack visualization, detecting anomalous network components or preventing attacks by timely notifications about network changes.

7 User-Centric Security (T1.5)

Task 1.5 (T1.5) of the CONCORDIA project is focused on user–centric security. Specifically, the main research pillars of T1.5 are as follows:

- **Privacy:** This task aims on developing techniques for Personal Identifiable Information (PII) leakage detection to the advertising ecosystem and the general Web.
- **Identity Management:** The objective of this task is the developing blockchain based methods for creating digital identities. This will allow users of Online Social Networks (OSN) to verify their real-world identities without a centralized authority for storing and managing their personal information.
- **Social Networks and Fake News:** This task focuses on investigating techniques for identifying fake news in Online Social Networks as well as developing blockchain based methods for suppression of fake News.

In the following sections we present the results of the research activities related to Task 1.5 three main objectives.

Privacy

Section 7.1 discusses the results of a large-scale analysis regarding the Personal identifiable Information (PII) leakage via registration pages of thousands of popular websites. Section 7.2 presents the findings of five research activities related to methodologies and tools that allow the end-users to monitor advertising ecosystem. Section 7.3 is related to user security in information systems. Specifically, it presents the CyberSure, a cyber–insurance framework with applications to the health–care sector. Section 7.4 investigates how machine learning algorithms can be used to support individuals in specifying and managing their privacy preferences. Moreover, a privacy matching mechanism able to relax the conditions in users' privacy preferences in order to match service providers privacy policies, is also presented. Section 7.5 summarizes the research on two topics; (i) privacy in online social networks, and (ii) the complexity of decentralized application development. Furthermore, a Cybersecurity Education Model is proposed.

Identity Management

Section 7.6 discuses a new approach based on the role of human behavior in guaranteeing the trustworthiness of IoT environments and defines an assurance methodology based on policies to evaluate the compliance of human behavior and the trustworthiness of corresponding smart devices. Section 7.7 presents the TradeMap, an integrated architecture, designing and enabling an online end–to–end (e2e) trading marketplace, while supporting anonymous management features. Section 7.8 describes the importance of data visualization on critical decisions regarding cyber–threats. Hence, a security management dashboard for interactive use by cybersecurity analysts is proposed. Section 7.9 presents a large-scale blockchain testbed, called BlockZoom with adaptable environment built on top of the Grid5000 platform. The platform is a large-scale environment for experiment-driven research with large amount of resources scattered over 8 geographic sites in France and Luxembourg. Section 7.10 describes the DIMANDS p2p, an autonomous discovery system designed to organize the identities that a given user may own. It is a solution to the identity management (IdM) problem, namely the management of diverse identities that users own in online networks and services.

Social Networks and Fake News

Section 7.11 discusses the results of three studies related to the behavior of state–sponsored trolls in social networks as well as the influence that the trolls disinformation campaigns had during the 2016 US presidential election.

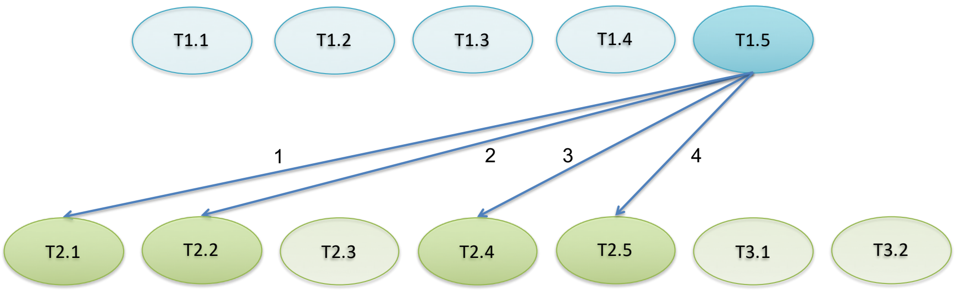

Figure 16: Links between task T1.5 and the CONCORDIA pilots

The research performed within Task 1.5 is linked to the CONCORDIA pilots as follows (see Figure 16):

1. The papers [104, 100, 130, 62] (see Section 7.2) which are related to ad-ecosystem have active interaction with Telefonica's pilot (T2.1).
2. The papers [56, 129, 94] (see Sections 7.3 and 7.9) are in close alignment with "eFinance" pilot T2.2.
3. The papers [55, 8] (see Sections 7.3 and 7.6) have potential for the "e–health" pilot T2.4.
4. The paper [8] has potential for the "Security of Unmanned Aerial Systems (UAS)" pilot T2.5.

7.1 You Shall Not Register! Detecting Privacy Leaks across Registration Forms

Contact: Sotiris Ioannidis (FORTH)

Most of the modern web services offer to users the ability to be registered to them via dedicated registration pages. Most of the times, they use this method so the users can profit by accessing more content or privileged inside items. In these pages, users are typically requested to provide their names, email addresses, phone numbers and other personal information in order to create an account. As the purpose of the tracking ecosystem is to collect as many information and data from the user, this kind of PII (personally identifiable information) might leak on the 3rd-Parties, when the users fill in the registration forms.

As the purpose of online advertising is to increase the market share of the companies by promoting their products and services, the advertising industry continuously designs new mechanisms to deliver more effective and highly targeted ads. In order to serve highly targeted ads, advertisers employ various, often questionable and privacy intrusive, techniques for collecting and inferring users' personal information. They typically employ techniques, both stateful and stateless, for tracking users visits across different websites, which allow them to collect user's data, and reconstruct parts of their browsing history.

In this work, we conduct a large-scale measurement analysis of the PII leakage via registration pages of the 200,000 most popular websites of the web. We design and implement a scalable and easily replicable methodology, for detecting and filling registration forms in an automated way. Our analysis shows that a number of websites ($\approx 5\%$) leak PIIs to 3rd-Party trackers without any user's consent, in a non-transparent fashion. Furthermore, we explore the techniques employed by 3rd-Parties in order to harvest user's data, and we highlight the implications on user's privacy.

In the first year we conducted the analysis mentioned above and submitted our work in IOSec2019 workshop of ESORICS 2019.

This study is related to T1.5 User-Centric Security and specifically on the online advertisement ecosystem. The analysis inspection of the content in this study is associated with privacy leakage, Personally Identifiable Information and Web Tracking and CONCORDIA's 2.2 pilot: "Finance Sector: Assessing Cyber Risks, Threat Intelligence for the Finance Sector."

7.2 User Privacy Within the Advertising Ecosystem

Contact: Nicolas Kourtellis (TID)

The work performed by Telefonica (TID) together with internal CONCORDIA collaborators as well as external collaborators is generally aiming to provide methodologies and tools to end-users, privacy researchers and policy makers and auditors, so that they are able to monitor and audit the advertising ecosystem, shed light on the protocols used by its various entities regarding user tracking and data sharing practices, and increase transparency on the web.

Research results have been published in various top conferences and are summarized in the following paragraphs.

User data is the primary input of digital advertising, fueling the free Internet as we know it. As a result, web companies invest a lot in elaborate tracking mechanisms to acquire user data that can sell to data markets and advertisers. However, with same-origin policy, and cookies as a primary identification mechanism on the web, each tracker knows the same user with a different ID. To mitigate this, Cookie Synchronization (CSync) came to the rescue [104], facilitating an information sharing channel between third parties that may or not have direct access to the website the user visits. In the background, with CSync, they merge user data they own, but also reconstruct a user's browsing history, bypassing the same origin policy. In this paper, we perform a first to our knowledge in-depth study of CSync in the wild, using a year-long weblog from 850 real mobile users. Through our study, we aim to understand the characteristics of the CSync protocol and the impact it has on web users' privacy. For this, we design and implement CONRAD, a holistic mechanism to detect CSync events at real time, and the privacy loss on the user side, even when the synced IDs are obfuscated. Using CONRAD, we find that 97% of the regular web users are exposed to CSync: most of them within the first week of their browsing, and the median userID gets leaked, on average, to 3.5 different domains. Finally, we see that CSync increases the number of domains that track the user by a factor of 6.75.

In recent years, Header Bidding (HB) has gained popularity among web publishers, challenging the status quo in the ad ecosystem. Contrary to the traditional waterfall standard, HB aims to give back to publishers control of their ad inventory, increase transparency, fairness and competition among advertisers, resulting in higher ad-slot prices. Although promising, little is known about how this ad protocol works: What are HB's possible implementations, who are the major players, and what is its network and UX overhead?

To address these questions, we design and implement HBDetector: a novel methodology to detect HB auctions on a website at real-time [100]. By crawling 35,000 top Alexa websites, we collect and analyze a dataset of 800k auctions. We find that: (i) 14.28% of top websites utilize HB. (ii) Publishers prefer to collaborate with a few Demand Partners who also dominate the waterfall market. (iii) HB latency can be significantly higher (up to 3× in median case) than waterfall.

Although digital advertising fuels much of today's free Web, it typically does so at the cost of online users' privacy, due to the continuous tracking and leakage of users' personal data. In search for new ways to optimize the effectiveness of ads, advertisers have introduced new advanced paradigms such as cross-device tracking (CDT), to monitor users' browsing on multiple devices and screens, and deliver (re)targeted ads in the most appropriate screen. Unfortunately, this practice leads to greater privacy concerns for the end-user.

Going beyond the state-of-the-art, we propose a novel methodology for detecting CDT and measuring the factors affecting its performance, in a repeatable and systematic way [130]. This new methodology is based on emulating realistic browsing activity of end-users, from different devices, and thus triggering and detecting cross-device targeted ads. We design and build Talon, a CDT measurement framework that implements our methodology and allows experimentation with multiple parallel devices, experimental setups and settings. By employing Talon, we perform several critical experiments, and we are able to not only detect and measure CDT with average AUC score of 0.78-0.96, but also to provide significant insights about the behavior of CDT entities and the impact on users' privacy. In the hands of privacy researchers, policy makers and end-users, Talon can be an invaluable tool for raising awareness and increasing transparency on tracking practices used by the ad-ecosystem.

Being able to check whether an online advertisement has been targeted is essential for resolving privacy controversies and implementing in practice data protection regulations like GDPR, CCPA, and COPPA. In [62] we describe the design, implementation, and deployment of an advertisement auditing system called eyeWnder that uses crowdsourcing to reveal in real time whether a display advertisement has been targeted or not. Crowdsourcing simplifies the detection of targeted advertising, but requires reporting to a central repository the impressions seen by different users, thereby jeopardizing their privacy.

We break this deadlock with a privacy preserving data sharing protocol that allows eyeWnder to compute global statistics required to detect targeting, while keeping the advertisements seen by individual users and their browsing history private. We conduct a simulation study to explore the effect of different parameters and a live validation to demonstrate the accuracy of our approach. Unlike previous solutions, eyeWnder can even detect indirect targeting, i.e., marketing campaigns that promote a product or service whose description bears no semantic overlap with its targeted audience.

Recent work has demonstrated that by monitoring the Real Time Bidding (RTB) protocol, one can estimate the monetary worth of different users for the programmatic advertising ecosystem, even when the so-called winning bids are encrypted. In [99] we describe how to implement the above techniques in a practical and privacy preserving manner. Specifically, we study the privacy consequences of reporting back to a centralized server, features that are necessary for estimating the value of encrypted winning bids.

We show that by appropriately modulating the granularity of the necessary information and by scrambling the communication channel to the server, one can increase the privacy performance of the system in terms of K-anonymity. We have implemented the above ideas on a browser extension and disseminated it to some 200 users. Analyzing the results from 6 months of deployment, we show that the average value of users for the programmatic advertising ecosystem has grown more than 75% in the last 3 years.

The main achievements from this line of work in 2019 were six academic publications (4 published and 2 under review) to top conferences and two user tools (web browser plugins) with back-end systems for increase user transparency on the web. Furthermore, the work included active interaction with the telco pilots, and especially Telefonica's pilot (T2.1) and active research and development and collaboration with CONCORDIA partners: TID, FORTH, CUT.

This research and development performed under CONCORDIA is directly related to the Telco pilots in Task 2.1. In particular, the methods proposed and the tools developed by TID and the collaborating partners in these research papers will be useful in the Telefonica pilot, which focuses on detecting and blocking leakage of PII and other sensitive or personal data to unauthorized 3rd parties, proposing new, privacy-preserving machine learning methods for modeling end-users without exposing private information and without losing their anonymity on the web, and deploying such methods on user devices or even at the network level and on edge computing nodes, for network-wide application of the methods.

7.3 User Security in Information Systems

Contact: George Hatzivasilis (FORTH)

Insurance of digital assets is becoming an important aspect nowadays, in order to reduce the investment risks in modern businesses. GDPR and other legal initiatives makes this necessity even more demanding as an organization is now accountable for the usage of its client data. In [55], we present a cyber insurance framework, called CyberSure. The main contribution is the runtime integration of certification, risk management, and cyber insurance of cyber systems. Thus, the framework determines the current level of compliance with the acquired policies and provide early notifications for potential violations of them. CyberSure develops CUMULUS certification models for this purpose and, based on automated (or semi-automated) certification carried out using them, it develops ways of dynamically adjusting risk estimates, insurance policies and premiums. In particular, it considers the case of dynamic certification, based on continuous monitoring, dynamic testing and hybrid combinations of them, to adapt cyber insurance policies as the conditions of cyber system operation evolve and new data become available for adjusting to the associated risk. The applicability of the whole approach is demonstrated in the healthcare sector, for insuring an e-health software suite that is provided by an IT company to public and private hospitals in Greece. The overall approach can reduce the potential security incidents and the related economic loss, as the beneficiary deploys adequate protection mechanisms, whose proper operation is continually assessed, benefiting both the insured and the insurer.

This activity is connected with CONCORDIA's 2.4 pilot "e-Health Sector: Privacy and Data Protection".

Nowadays, more-and-more aspects of our daily activities are digitalized. Data and assets in the cyber-space, both for individuals and organizations, must be safeguarded. Thus, the insurance sector must face the challenge of digital transformation in the 5G era with the right set of tools. In [56], we present CyberSure – an insurance framework for information systems. CyberSure investigates the interplay between certification, risk management, and insurance of cyber processes. It promotes continuous monitoring as the new building block for cyber insurance in order to overcome the current obstacles of identifying in real-time contractual violations by the insured party and receiving early warning notifications prior the violation. Lightweight monitoring modules capture the status of the operating components and send data to the CyberSure backend system which performs the core decision making. Therefore, an insured system is certified dynamically, with the risk and insurance perspectives being evaluated at runtime as the system operation evolves. As new data become available, the risk management and the insurance policies are adjusted and fine-tuned. When an incident occurs, the insurance company possesses adequate information to assess the situation fast, estimate accurately the level of a potential loss, and decrease the required period for compensating the insured customer. The framework is applied in the ICT and healthcare domains, assessing the system of medium-size organizations. GDPR implications are also considered with the overall setting being effective and scalable.

This activity is connected with CONCORDIA's 2.2 pilot "Finance Sector: Assessing Cyber Risks, Threat Intelligence for the Finance Sector".

7.4 Supporting Individual in Specifying and Managing their Privacy Preferences

Contact: Barbara Carminati (UI)

Nowadays, individuals are becoming increasingly dependent on numerous online services to make their lives more comfortable and convenient. To offer such services, service providers collect, store, and process a massive amount of personal information about individual users. However, although individuals voluntarily provide such personal information to service providers, they often have no idea how their information is subsequently used. This may also cause serious privacy threats as users lose control over their data.

In general, service providers collect these data in accordance with their privacy policies. Privacy policy defines how a service provider collects and manages customers' personal information, the purpose for which that information might be used, and the length of time it can be retained. To have more control over how their data are used, individuals can specify their privacy settings (aka privacy preferences), stating how their data should be used and managed by the service provider. Unfortunately, the average user might find it difficult to properly set up privacy settings due to lack of knowledge and subsequent lack of decision-making abilities regarding the privacy of their data.

To overcome this problem, during the first year of CONCORDIA, we have investigated how machine learning algorithms could be used to support individuals in specifying and managing their privacy preferences. This results with the proposals of two different approaches. The first one copes with the consideration that, in general, privacy checking is handled by hardly matching users' privacy preferences against service providers' privacy policies. This hard matching implies denying all services whose privacy policies do not fully match an individual's privacy preferences. If individuals do not choose suitable privacy settings, they may not be able to access many services. This could represent a problem if we consider that many users experience difficulty when selecting privacy preferences. Let us consider, for instance, a user's privacy preference stating that his health-related data (e.g., blood-pressure and heart-beat data) can be shared with a healthcare service provider if it retains the data only for 100 days. A healthcare service provider whose privacy policy states that it will store data for 110 days does not satisfy the user's privacy preference, even though it is close to satisfying them. In a real-life scenario, if a user has a health condition, he would likely not care about data retention for a further 10 days. He would likely make an exception to his privacy preference to access the service. However, a hard privacy matching mechanism would deny the user access to the service without considering the possible benefits for the user.

To cope with this scenario, in [4] we have investigated a more flexible privacy checking system for IoT-based environments. More precisely, we propose a soft privacy matching mechanism able to relax, in a controlled way, some of the conditions in users' privacy preferences to match service providers' privacy policies. The key idea is to adopt machine learning to infer, for each user, which component of a privacy preference (e.g., purpose, data, retention, and recipient) she/he is willing to relax and how much (i.e., relaxation range). At this purpose, we exploit supervised machine learning algorithms over a labeled training dataset, consisting of provider service requests, their related privacy policies, and user decisions on joining/denying these services. The learning algorithm takes as features the decision and the distance values between each component of the user's privacy preferences and provider' policy. Then it builds a classifier able to state when privacy preference can be relaxed, which of its components has to be modified and how much it can be relaxed.

We had further investigated the problem by acknowledging that individual's privacy decision may vary based on the situation he/she was living when the decision was taken. For instance, a user may feel comfortable accessing entertainment services while at home but not during office hours when he or she is at work. Thus, we extended our approach by taking into account individuals' contextual information, where contextual information refers to any piece of data of an individual that can be used to define his or her current situation.

In general, to achieve more fine-grained control, users might specify privacy preferences stating conditions on how personal data has to be used also based on the current situation (e.g., no access to entertainment services when the location is office). This brings the nice benefit of increasing user control over his/her data. However, since a user may interact with several contexts, it also increases the number of preferences that (s)he has to specify and manage, resulting in a very complicated and time-consuming task. To address this issue, in [5], we proposed a service that helps users to manage their privacy settings when they move from one context to another. The main idea is to infer the best privacy preferences for the new context leveraging on privacy preferences previously specified by the user for different contexts. At this aim, we select those contexts for which user has already specified a preference and that are similar to the new one. In doing this, we do not merely rely on a similarity measure between contexts, but we also keep into account user's perspective. Indeed, given a similarity measure, two contexts could have the same distance to the new one but differ on a few fields (e.g., time, location) that are very relevant for that user. In [5], we exploited machine learning to learn, for a target user, which context field is more informative for the privacy-decision and thus should have more relevance in the similarity measure. Once the most similar and most relevant contexts have been selected, the system has to retrieve the corresponding privacy preference. Here, the basic assumption is that this preference might represent a good match for the new context, but some slight modifications might be needed as well. To understand whether and how the identified privacy preference has to be adapted for the new context, we want to take into account once again user's perspective. More precisely, to learn which fields of the identified privacy preference need to be modified (e.g., purpose, data, retention, and recipient), we exploit machine learning to infer, from each user, which component and how is willing to adapt it.

We plan to continue this research by also considering the trustworthiness of the entities involved in the privacy-decision process (e.g., data, context information, users). At this purpose, we are working with the University of Milano to design privacy management services empowered with their assurance methodology based on human behavior compliance.

These research activities are related to "privacy" pillar of task 1.5.

7.5 Social Network Privacy and Complexity of Distributed Applications

Contact: Blaž Podgorelec, Lili Nemec Zlatolas, Muhamed Turkanović (UM)

The research was performed on two topics, privacy in online social networks and the complexity of decentralized application development. Furthermore, an analysis of a Cybersecurity Education Model was performed. The former [92] was performed with the goal to understand the issues (e.g., privacy, trust, self-disclosure) concerning social network's users, particularly Facebook users. Within this research, a model was built, which includes privacy value, privacy risk, trust, privacy control, privacy concerns, and self-disclosure. The model was evaluated through structural equation modeling, whereby an online survey was conducted where a total of 602 respondents participated. The research results (i.e., proposed model) provide new knowledge about privacy issues, trust, and self-disclosure on online social networks and can thus be helpful for other researchers on that area or developers of online social networks. The results of the research also show significant relationships between the constructs, e.g., privacy risk, privacy value, trust in Facebook, self-disclosure, privacy control. Presented results of the research directly provide knowledge related to User-Centric security, especially to the field of privacy with the correlation to the area of social networks. In the article [113], an analysis of decentralized application development complexity was performed. The results can have direct implications on the development process of decentralized applications, whose primary purpose is to enable decentralization of the web, while also fostering the security features of distributed and decentralized storage like blockchains. With this, the developers and researchers in the field of decentralized applications, can better understand and perform additional studies on how the complexity of the development process impacts the usability of the decentralized applications for the end-users, and how the usability can have implications on the user security (i.e., privacy, data leakage, identity fraud, etc.), since higher complexity can lead to a higher number of flaws in the design and code. Because of the reasons mentioned above (i.e., impacts of development complexity to the usability of the decentralized applications), knowledge obtained from research indirectly contributes to the privacy aspects related to User-Centric security. Paper [135] presents an overview of a cybersecurity education model from the Information Systems and Information Technology perspective, together with a good example. The presented education model is shaped according to the guidelines by the Joint Task Force on Cybersecurity Education and the expectations of the Slovene industry regarding the knowledge and skills their future employees should possess, which relates to the User(Employee)-Centric security.

7.6 Assessing Trust Assurance based on Human Behavior Compliance

Contact: Claudio A. Ardagna

The success of the Internet of Things (IoT) is increasingly pushing towards the development of Cyber–Physical Systems (CPSs) that go beyond traditional IT boundaries, combining physical and digital environments in a single one. Today systems have an opaque perimeter, where a mixture of platforms, software, services, things, and people collaborate in an opportunistic way. Computations are then moving from the center (e.g., cloud) to the periphery, supporting analytics and knowledge extraction partially executed at the edge of the network, near the physical environment and sensors where data are collected. In this scenario, people carry/manage a plethora of smart devices (often integrated in their smartphones or in their homes) sensing the surrounding environment and communicating with edge nodes without even noticing. Like for the Web 2.0 revolution when user-generated content entered the loop, people are not only passive consumers of pervasive services; they rather become (often unintentional) service providers that distribute user-collected data. The exponential growth of smart devices (200 billions of connected objects with a mean of 26 objects for every human predicted by 2020) and connected people (2.87 billion users carrying a smartphone by 2020), coupled with pervasive mobility and the proliferation of IoT applications, make the trustworthiness of collected data and the users' privacy the most important requirements to the success of these applications. Data trustworthiness is mandatory to build a chain of trust on a decision process taken according to IoT data and edge computations. This is very similar to the Web 2.0 scenario, where the plague of fake news and fake data substantially decreased the quality of decision processes, such as, for instance, in case of Twitter bots able to support or defame a specific product with high rates of success.

Traditional solutions to data validation mostly focused on assurance or reputation techniques. On one side, the system generating data undergoes an assurance process based on testing or monitoring; on the other side, the reputation of the entity owning the system is evaluated. Both approaches are however not viable in complex cyber–physical scenarios based on smart devices for different reasons: i) devices are usually resource constrained and cannot therefore be the target of extensive testing/monitoring; ii) IoT systems have fuzzy perimeters that are difficult to evaluate using traditional assurance techniques; iii) devices are often owned by unknown users, whose reputation cannot be easily evaluated; iv) people are often unintentional data providers, differently from Human-Provided Services where users consciously distribute services. Other approaches have provided solutions based on behavioral analysis, aimed to understand and differentiate human behaviors. In this scenario, there is an increasing need of privacy-aware solutions that evaluate the compliance of people to behavioral policies and, in turn, the trustworthiness of data collected through their devices.

Our work in the first year of the project focused on the need of a new human/user–centric approach, where the behavior of people carrying/managing smart devices is an indicator of the trustworthiness of the generated data and corresponding devices [8]. We defined an assurance methodology based on data analytics that selects trustworthy devices according to the behavior of involved people. First, a technology–independent assurance policy template specifies requirements on correct human behaviors and data collection processes. Second, the assurance policy template is translated in a technology-dependent assurance policy instance specifying corresponding activities for human-behavioral modeling and data collection. Finally, trustworthy smart devices (and corresponding data) are selected according to a specific evaluation function and used to build a chain of trust underpinning a specific decision/inference. The proposed approach i) clarifies the role of human behavior in guaranteeing the trustworthiness of IoT environments; ii) defines an assurance methodology based on policies to evaluate the compliance of human behavior and, in turn, the trustworthiness of collected data and corresponding smart devices.

The assumption that the behavior of people owning smart devices can contribute to the evaluation of the trustworthiness of collected data and, in turn, of the whole decision process radically changes the context in which assurance solutions must execute. The intuition behind our approach is that the more people behavior is compliant, the higher the trustworthiness of data collected through their smart devices. On the other side, privacy protection is unavoidable and must be carefully considered. We are then working within Task 1.5 with University of Insubria to integrate our approach with solutions for users' privacy protection. The idea is to provide new privacy-aware security approaches (e.g., access control, identity management, authorization mechanisms) with verifiable behavior. We will continue working on this topic considering assurance techniques, blockchain, and the like, mainly in the context of Pilot "Security of Unmanned Aerial Systems (UAS)" to the aim of providing a privacy-aware location-based access control for UAV, and Pilot "e–Health Sector: Privacy and Data Protection" for managing emergency scenarios in smart homes.

7.7 TradeMap: A FINMA-compliant and Blockchain-based End-2-end Trading MarketPlace

Contact: Sina Rafati, Burkhard Stiller (UZH)

Data leaks and privacy scandals have been a growing concern of the last decade. While most traditional, i.e., centralized, online platforms (such as online marketplaces) require users to register with their personal data, they potentially expose the user's identity and data to be used for unintended purposes. Recently a great attention is raised to use Blockchains (BCs) to overcome such issues. However, at a first glance, eCommerce approaches relying on BCs either for KYC or for the e2e trading, suffer from several deficits. These include: **(1)** Lacking a sufficient and reliable user identification, **(2)** neglecting regulations, e.g., the Swiss Financial Market Supervisory Authority (FINMA), **(3)** lacking of ZKP-based mechanisms in support of privacy, **(4)** using fully decentralized BCs, which by default do not bring privacy as data stored by them can be accessed publicly, **(5)** neglecting costs of storing large data on BCs, **(6)** overloading an SC, **(7)** missing the use of standardized interfaces, and **(8)** missing a secure back-end and SC implementation. To address these deficits, we introduce TradeMap, an integrated architecture, designing and enabling an online end-to-end (e2e) trading marketplace, while supporting anonymous management features.

TradeMap addresses the Swiss Financial Market Supervisory Authority (FINMA) regulations by designing a FINMA-complaint Know Your Customer (KYC) platform. Additionally, TradeMap is based on blockchains and employs Ethereum smart contracts (SC). Thus, trust and anonymity between the marketplace and the KYC system relies on zero knowledge proof-based SCs used for user identification processes. With this management approach proposed, the user authentication is only verified within the KYC platform, providing a legally valid and fully anonymous online trading platform.

The SC developed brings requirements such as accessibility, anonymity and usability. The accessibility requirement for users could be met by having a browser extension installed allowing the user to create a new wallet. Since sending transactions to the Ethereum network always cost an amount of Ether, it is important that the browser extension has access to the user's wallet to transfer the funds. Blockchain transactions are transparent and visible for all users in the network. Since the goal is to allow user's complete anonymity, it must be considered to make SC requests anonymous. A design proposal is to use a unique hashed key (KYC key) with which users could register themselves at connected platforms anonymously while providing consent that the user's data could be shared with the platform. To prevent that the KYC key is copied once a SC request was processed, it is suggested that the mentioned key be hashed.

Designing multi-component systems based on blockchains (BC) can only make IT systems better, if all features supported by traditionally centralized applications, such as KYC or eCommerce, are offered unchanged and improvements as of new features and functionality is being added. From a governance perspective of managing user interactions with centralized systems, especially with respect to the case of relying on a decentralized trust factor of BCs, TradeMap propose a novel design, that meets facilitating traditional eCommerce as a marketplace of e2e trading. In this sense TradeMap integrates the BC with care, undertaking several steps to meet FINMA regulations of user registrations.

TradeMap performs such a registration without storing all user data in the BC. With respect to the trading platform, the permission for verifying the user's identity is only done by a method, which itself does not need any identity disclosure. Thus, before and after the identity verification, the trading platform's administrator will not receive any identity-related information. In that sense the anonymous management feature had been achieved for the e2e trading marketplace. TradeMap enables data management for users, e.g., businesses, to trace back the history of products before purchase. With TradeMap each product is now traceable, safe trading is enabled, and a healthy and ethical marketplace is designed to offer full independence to users from trusted third parties, such as banks.

A short paper presented at the poster session of IFIP/IEEE CNSM 2019 [111] and the technical report elaborating the processes and operations of the developed system including smart contracts is presented in [110] .

This work directly relates to privacy enhancement, zero knowledge proofs and blockchains, which are all integrated into a marketplace platform connected to a KYC system which established a FINMA compliant user verification. This method enabled anonymous and at the same time legal activities in a trading platform. All the key principles followed in this work are in close alignment to the Task 1.5 of WP1. This work is conducted under the task 1.5 goals such as "identity management" and "privacy", where adhering to the existing regulations is one of the main concerns. Moreover, this work might be of interest in the finance pilot with the peer to peer trading features being developed based on blockchains and cryptocurrencies.

7.8 Dashboard for Collaborative Threat Management Visualization

Contact: Christian Killer, Bruno Rodrigues, Burkhard Stiller (UZH)

Cybersecurity concerns have globally risen to one of the top priorities in both research and development. Distributed Denial-of-service (DDoS) attacks are one major concern in cybersecurity, because their perpetration requires little effort on the attacker's side, while massive damage is inflicted to the victim. One approach to defend against DDoS attacks is sharing hardware and defense capabilities with other systems, an approach called cooperative DDoS mitigation. In a context where Autonomous Systems (AS) rely on cybersecurity specialists to make critical decisions regarding threats, it is necessary to structure and categorize data such that visualization "makes sense" to the analyst. As a cooperative defense involves multi-disciplinary concepts and the decision-making process usually requires a low response time from the user, selecting an appropriate type of graphical representation and flow of interaction is not a straightforward task [65, 66]. Thus, this work covered the design of a security management dashboard for a cooperative defense, especially designed for interactive use by cybersecurity analysts.

7.9 Large-Scale Blockchain Testbed and A Case Study of Automation of KYC Result Sharing

Contact: Wazen Shbair, Radu State (SnT)

Blockchain platforms are anticipated to serve millions of users. Therefore, the performance evaluation of new solutions have to consider large-scale assessment of the technologies behind the scene. Currently, testing and evaluating blockchain apps are conducted using simulation tools, private and public testnet networks. The simulation tools prove the validity of an approach. Private network facilitates adapting blockchain configuration, so it is possible to reproduce experiments under different settings. Whereas testing using public networks mimics conditions that are similar to the production environment. However, public and private networks have drawbacks; public networks obstructs a deep understanding of the employed technology, while private networks make the testing environment incomparable with real world conditions. To address this issue, we develop a large-scale blockchain testbed, called BlockZoom with adaptable environment built on top of the Grid5000 platform. The platform is a large-scale environment for experiment-driven research with large amount of resources scattered over 8 geographic sites in France and Luxembourg. Through different configuration scenarios developers can evaluate the apps behavior at a scale comparable to the production environment and in the same time has the flexibility of private blockchain networks for testing under different conditions. As a direct application of using the testbed, the KYC problem has been explored. A blockchain-based KYC prototype has been designed and developed. The prototype has been evaluated via the testbed under different configuration scenarios and stress testing.

Fake client account is a serious issue in Online Social Networks (OSN) and that explains why OSNs are plagued by online fraud. To avoid reinvent the wheel, we learned the lessons from the financial industry where the verified client identity is crucial to avoid financial transaction fraud. The term ”Know Your Customer” (KYC) is widely used in banks to identify their clients before doing financial business with them. Therefore, banks need to collect data on their customers, for example on their identity and on their home address. This data is then used to perform security and background checks on the customers. KYC is required by EU law, in order to ensure that banks deal only with legitimate customers. However, the KYC process is cumbersome and costly, and required to be performed repeatedly by every bank for each of their customers. This includes redundant work that can be alleviated by trusted automation. In this context, blockchain technology can solve many of the problems related to KYC like the on-boarding issue. Blockchain-based approach can disseminate the client's information across various parties once verified and provides the trustability and tractability of the provided information. Thus, a KYC once performed can be accessed by other financial institutions with unique authorization from the client. This will make the KYC process much easier, simpler, less time consuming and cost-effective. In this sense we will work on transferring this technique to the OSN sector.

Working prototypes of the BlockZoom testbed [129] and Blockchain-based KYC system [94] were presented at the demo session of IEEE/ICBC 2019.

This work is directly related to the “Identity Management” pillar of Task 1.5. The developed KYC system is in close alignment with the requirements of the financial pilot - Caixa Bank in terms of digital on-boarding and automating the verification of new clients identities. In addition BlockZoom testbed provides rational evaluation environment for the blockchain-based solutions, which may be developed under the umbrella of CONCORDIA project.

7.10 Cross-federation Identity Solutions using Blockchain Technologies

Contact: Kostas Lampropoulos (UP)

One of the most challenging issues of today's networks and services (in both the physical and cyber world) is the management of diverse identities that users own across different contexts. This is known as the identity management (IdM) problem and involves the management of users' diverse identities and identity related data that are scattered across different network locations and are only meaningful within the service context in which they are used. Across different service contexts they appear as unassociated identity data that cannot be attributed to the same user or exchanged among users. The aim of this work is to propose a global identity management solution that

1. can provide IdM solutions with the use of existing identities and does not introduce any kind of new global IDs and
2. is not limited inside the context of a federation border but instead connects the various federations and identity schemes facilitating the easy creation of new private business environments which can operate under their own technological and business requirements.

Various solutions have been developed in the area of IdM, but the proposed systems are usually the outcome of independent efforts, addressing the IdM problem from different perspectives, narrow contexts, and inevitably different requirements. Their applicability highly depends on proprietary features like the use of specific global identity formats, while their scope is confined in closed trusted groups (federations) with statically defined procedures and mechanisms. The result is a diversification of IdM solutions with serious interoperability issues among emerging IdM islands. Any attempts to integrate these islands aggravate rather than solve the problem. For instance, the idea to introduce a commonly accepted global identifier leads to the proliferation of global IDs from parties that try to enforce their own solutions. Efforts to integrate different federations into bigger ones recursively introduce new identity formats and procedures.

Our work is based on a system called DIMANDS designed a few years ago based on a large-scale peer to peer architecture. DIMANDS is an autonomous discovery system designed to organize all kinds of identities that a single user may own (e.g., government organizations, applications, service or network providers). It uses existing identifiers and does not introduce new identity formats which affect existing services or network procedures. On-going work is evolving the DIMANDS architecture to allow easier and wider adoption. DIMANDS p2p overlay is now redesigned and imported inside a blockchain (ethereum). However, with the adoption of blockchain technologies many of the initial architectural decisions are no longer valid due to the open access of the network. Thus, we are re-designing many parts of the system in order to meet the strict privacy and security requirements of a global IdM system.

For the first year, this work has not been published yet. The system is still in the development phase and we are completing a first updated version of the architecture, which can support all strict security and privacy requirements of an IdM system with the use of a public blockchain.

This work is associated with T1.5 User-Centric Security and addresses various security and privacy aspects of the problem. Furthermore, it facilitates the creation of new innovative business cases since it operates with existing identity schemes and does not introduce any kind of new ID that must be adopted by service providers. For CONCORDIA, this work is tightly connected to the Financial (T2.2) use case (KYC) where different unassociated identities must be connected to provide better validation of an end user. Another pilot that might also benefit from our solution is the Telco (T2.1), which requests solutions for identifying a specific user from the large variety of end devices he/she may be using in the near future and with a wider adoption of 5G deployment.

7.11 Analyzing the Behavior and Influence of State–Sponsored Trolls in Social Networks

Contact: Nikos Salamanos, Michael Sirivianos (CUT)

Over the past couple of years, anecdotal evidence has emerged linking coordinated campaigns by state-sponsored actors with efforts to manipulate public opinion on the Web, often around major political events, through dedicated accounts, or “trolls”. Although they are often involved in spreading disinformation on social media, there is little understanding of how these trolls operate, what type of content they disseminate, and most importantly their influence on the information ecosystem. A leading example of state–sponsored troll disinformation campaigns is found in the Russian efforts to interfere in and manipulate the outcome of the 2016 US presidential election which were unprecedented in terms of the size and scope of the operations. Millions of posts across multiple social media platforms gave rise to hundreds of millions of impressions targeting specific segments of the population in an effort to mobilize, suppress, or shift votes. Trolls were particularly focused on the promotion of identity narratives, though that does not distinguish them from many other actors during the election. Hence, there is a considerable debate as to whether state–sponsored disinformation campaigns that operated on social media were able to affect the outcome of the 2016 US Presidential election. While there is a large body of work which tried to address this question from distinct disciplinary angles, a conclusive result is still missing.

In the first year of the project, we analyzed the behavior and influence of state–sponsored trolls in social networks. This study has led to two academic publications [147, 148] and a work in progress [122].

In [122], we measured the impact of troll activities on the virality of the ambiguous political information that had been shared on Twitter during the 2016 US Presidential election. For that purpose, a very large directed graph has been constructed by the interactions between the users (tweet replies and mentions). The graph consists of 9.3 million nodes and 169 million edges and it has been constructed based on two Twitter datasets: (i) A collection of 152.5 million tweets that was downloaded using the Twitter API during the US presidential election period (from September 21 to November 7, 2016). Hence, we have access to original troll tweets that have yet to be deleted by Twitter. (ii) A collection of original troll tweets which have been released by Twitter itself as part of the investigation on foreign interference in the 2016 US election – the misinformation campaigns of 8,275 state–sponsored accounts linked to Russia, Iran and Venezuela states. By using graph analysis techniques and classification, we are able to identify the group of users who were most probably the driving force of the viral cascades that have been shared by the troll accounts. Our primary contributions are as follows:

1. We construct one of the largest graphs in the literature, which represents the interactions between state–sponsored troll accounts and authentic users in Twitter during the period of 2016 US Presidential election. This is an approximation of the original followers–followees social graph.

2. We present strong evidence that the trolls' activity was not the main cause that led to viral cascades of web and media material in Twitter. The experimental results clearly show that the authentic users who had close proximity to the trolls in the graph were the most active and influential part of the population and their activity was the driving force of the viral materials. A possible scenario is that instead of injecting new content, these trolls "resonate" with communities online with which they sought to form relationships. They do so, particularly by targeting opinion leaders in these communities. This is consistent with previous literature on information warfare tactics.

In [147], we addressed the following questions: (i) How do state–sponsored trolls operate? (ii) What kind of content do they disseminate? (iii) Is it possible to quantify the influence they have on the overall information ecosystem on the Web? The study is based on the set of 2.7K accounts released by the US Congress as ground truth for Russian state–sponsored trolls. From a dataset containing all tweets released by the 1% Twitter Streaming API, we search and retrieve 27K tweets posted by 1K Russian trolls between January 2016 and September 2017. We characterize their activity by comparing it to a random sample of Twitter users. Then, we quantify the influence of these trolls on the greater Web, looking at occurrences of URLs posted by them on Twitter, 4chan, and Reddit. Finally, we use *Hawkes Processes* to model the influence of each Web community (i.e., Russian trolls on Twitter, overall Twitter, Reddit, and 4chan) on each other. Our study leads to several key observations:

1. Trolls actually bear very small influence in making news go viral on Twitter and other social platforms alike. A noteworthy exception are links to news originating from RT (Russia Today), a state-funded news outlet: indeed, Russian trolls are quite effective in "pushing" these URLs on Twitter and other social networks.

2. The main topics discussed by Russian trolls target very specific world events (e.g., Charlottesville protests) and organizations (such as ISIS), and political threads related to Donald Trump and Hillary Clinton.

3. Trolls adopt different identities over time, i.e., they "reset" their profile by deleting their previous tweets and changing their screen name/information.

4. Trolls exhibit significantly different behaviors compared to other (random) Twitter accounts. For instance, the locations they report concentrate in few countries like the USA, Germany, and Russia, perhaps in an attempt to appear "local" and more effectively manipulate opinions of users from those countries. Also, while random Twitter users mainly tweet from mobile versions of the platform, the majority of the Russian trolls do so via the Web Client.

Finally, in [148] we studied many aspects of state–sponsored disinformation campaigns that remain unclear, that is, how do state–sponsored trolls operate? what kind of content do they disseminate? how does their behavior change over time? is it possible to quantify the influence they have on the overall information ecosystem on the Web? The study is based on two ground truth data about state–sponsored actors. First, we use 10M tweets posted by Russian and Iranian trolls between 2012 and 2018. Second, we use a list of 944 Russian trolls, identified by Reddit, and find all their posts between 2015 and 2018. We analyze the two datasets across several axes in order to understand their behavior and how it changes over time, their targets, and the content they shared. For the latter, we leverage word embeddings to understand in what context specific words/hashtags are used and shed light to the ideology of the trolls. Also, we use *Hawkes Processes* to model the influence that the Russian and Iranian trolls had over multiple Web communities; namely, Twitter, Reddit, 4chan's Politically Incorrect board (/pol/), and Gab. Our main findings are as follows:

1. Our influence estimation experiments reveal that Russian trolls were extremely influential and efficient in spreading URLs on Twitter. Also, when we compare their influence and efficiency to Iranian trolls, we find that Russian trolls were more efficient and influential in spreading URLs on Twitter, Reddit, Gab, but not on /pol/.

2. By leveraging word embeddings, we find ideological differences between Russian and Iranian trolls. For instance, we find that Russian trolls were pro–Trump, while Iranian trolls were anti-Trump.

3. We find evidence that the Iranian campaigns were motivated by real-world events. Specifically, campaigns against France and Saudi Arabia coincided with real-world events that affect the relations between these countries and Iran.

4. We observe that the behavior of trolls varies over time. We find substantial changes in the use of language and Twitter clients over time, for both Russian and Iranian trolls. These insights allow us to understand the targets of the orchestrated campaigns for each type of trolls over time.

5. We find that the topics of interest and discussion vary across Web communities. For example, we find evidence that Russian trolls on Reddit were extensively discussing about cryptocurrencies, while this does not apply, to a great extent, for the Russian trolls on Twitter.

This work is directly related to the “Social Networks and Fake News” research pillar of Task 1.5.

8 Organization of the Scientific Community and Events

The Description of Action identifies four (SMART) objectives for WP1. Two of these relate to the organization of the scientific community and events:

- Organization of scientific events in the area of cybersecurity, including a dedicated annual European cybersecurity conference.
- Leading role in the organization of the scientific community, outreach to different target audiences, including public media and the general public.

The subsections below will discuss progress and achievements for these activities. An interesting observation is that CONCORDIA activities seem to focus on venues where security is *applied* within a specific context. A good example of such venue is the ACM Internet Measurement Conference (IMC), which publishes excellent measurement research regarding (the security of) the Internet. This focus on *applied* or *usable security* is a strong indicator that CONCORDIA research is inspired by, and wants to have impact on, real-world security problems.

8.1 European Cybersecurity Conference

At the time of writing of the CONCORDIA project proposal the initial idea was to develop a European cybersecurity conference, similar to the AIMS (International Conference on Autonomous Infrastructure, Management, and Security) or TMA (Traffic Measurement and Analysis) conferences. After discussions with the three other pilot projects (ECHO, SPARTA and Cybersec4Europe) it was decided that, instead of creating independent and therefore competing conferences per project, a joint event supported by all four pilot projects would be a far better approach. Therefore, together with the three other pilot projects, CONCORDIA is co-organizing two PhD schools: the 4th International NeCS PhD Winter School and the IFIP Summer School on Privacy and Identity Management.

NeCS PhD Winter School

NeCS (European Network for Cybersecurity, www.necs-project.eu) was a H2020 Marie Curie ITN project that run till September 2019. NeCS was formed in response to the increased need for highly qualified experts, to address the issues of training and development of talented junior researchers, as indicated in the European Cyber-security strategy and highlighted in the EC's Digital Agenda. In the last three years NeCS organized in February PhD Winter Schools in Trento. Such schools lasted for the entire week and had several speakers and more than 30 students.

During the summer, the NeCS organizers contacted all four EU cybersecurity pilots (CONCORDIA, Cybersec4Europe, ECHO and SPARTA) with the request whether we could together continue the organization of these Winter Schools and continue the training of young researchers. CONCORDIA strongly supported this initiative, and an organizing committee was formed with two representatives per pilot project. The representatives from CONCORDIA are Prof. Burkhard Stiller (UZH) and Dr. Anna Sperotto (UT).

At the moment of writing this deliverable, the plan is that the 4th International NeCS PhD Winter School will take place in Trento from 20-24 of January. The first day will be dedicated to CyberRange platforms. The second till the fourth day will be filled by lectures. Each day there will be 4 lectures and each pilot will organize three of them. On Friday morning there will be an invited speaker and in the afternoon there will be PhD presentations.

IFIP Summer School on Privacy and Identity Management

The IFIP Summer School on Privacy and Identity Management 2020 is planned to take place between August 17-21, 2020 in Brno at the premises of the Masaryk University. Masaryk University is one of the partners within CONCORDIA; the local host and one of the general co-chairs will be Prof. Vashek Matyas. To avoid duplication, the organizers of the Summer School strongly synchronize with the NeCS Winter School organizers (for CONCORDIA: Prof. Burkhard Stiller and Dr. Anna Sperotto).

The school's lead theme will be: *Training and Awareness*, covering in particular:

- privacy and security evaluation,
- privacy and security certifications,
- usable privacy and security,
- privacy and security training and education,
- privacy and security awareness,
- privacy and security cyber ranges.

8.2 Organization of Scientific Events

CONCORDIA members are or have been very active in organizing scientific conferences. These conferences include (in alphabetic order):

- ADBIS 2019 (General chair and local chair)
  European Conference on Advances in Databases and Information Systems
- AISec 2019 (Workshop Chairs)
  12th ACM Workshop on Artificial Intelligence and Security
- C&TC 2019 (Program co-chair)
  International Symposium on Secure Virtual Infrastructures Cloud and Trusted Computing 2019
- CNSM 2019 (Program co-chair)
  IFIP/IEEE/In Assoc. with ACM SIGCOMM International Conference on Network and Service Management 2019

- CODASPY 2020 (Program co-chair)
  ACM Conference on Data and Application Security and Privacy 2020
- DISSECT 2020 (General (co)chair)
  6th IEEE/IFIP Workshop on Security for Emerging Distributed Network Technologies
- ENASE 2019 (Program co-chair)
  14th International Conference on Evaluation of Novel Approaches to Software Engineering
- ICDE 2019 (Research PC Vice Chair)
  35th IEEE International Conference on Data Engineering
- ICDE 2020 (General co-chair and Research PC Vice Chair)
  36th IEEE International Conference on Data Engineering
- IEEE BigData 2019 (Honorary General Chair)
  2019 IEEE International Congress on Big Data
- IEEE CLOUD 2019 (Program co-chair)
  11th IEEE International Conference on Cloud Computing
- IEEE ICWS 2019 ()
  11th IEEE International Conference on Web Services
- IEEE Services 2019 (Vice-Program Chair in Chief)
  2019 IEEE World Congress on Services
- IEEE TPS 2019 (Program co-chair)
  IEEE International Conference on Trust, Privacy and Security in Intelligent Systems, and Applications
- IFIP SEC 2020 (General chair and Local chair)
  International Conference on ICT Systems Security and Privacy Protection
- IM 2019 (Experience Program co-chair)
  IFIP/IEEE International Symposium on Integrated Network Management
- IMC 2019 (General chairs)
  ACM Internet Measurement Conference 2019
- Manifesto for Today 2019 (Organizers)
  Manifesto for Today - workshop co-located at ACM WomENcourage 2019
- MSTEC 2019 (General co-Chair)
  1st Model-Driven Simulation and Training Environments for Cybersecurity
- NetSoft 2019 ()Program co-Chair
  IEEE Conference on Network Softwarization
- OTM 2019 (General co-Chair)
  OnTheMove Federated Conferences & Workshops
- PAM 2020 (TPC (co)chair)
  Passive and Active Measurement conference
- SERVICES 2019 (Program co-chair)
  IEEE SERVICES Doctoral Symposium 2019
- SITIS 2019 (Honorary Chair)
  15th International Conference on Signal Image Technology & Internet based Systems

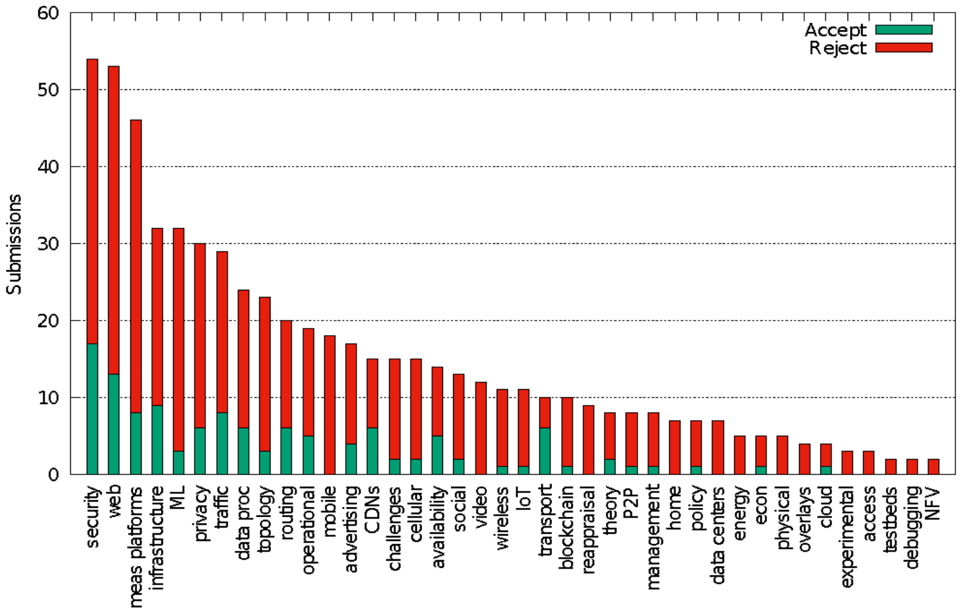

Figure 17: Most popular topics for the IMC 2019 Conference

The URLs and details of these conferences can be found in Appendix B. An interesting observation is that the 28 positions are equally divided between men and women (14 each).

Membership of Technical Program Committees (TPC)

As shown in Appendix C, CONCORDIA researchers have been active as TPC members for more than 40 conferences. For some of these conferences, e.g., the IEEE/IFIP Network Operations and Management Symposium taking place in Budapest in April 2020, up to six CONCORDIA members serve as TPC members.

ACM Internet Measurement Conference 2019

One of the key events organized by CONCORDIA members last year was the ACM Internet Measurements Conference (IMC), which was held between 21-23 October 2019 in Amsterdam, the Netherlands. This conference was attended by 234 attendees from 28 countries. In total 197 papers were submitted, nearly 50% came from the US and roughly one third from Europe. Papers were reviewed in two rounds. In the first round each paper was reviewed by at least three reviews. After this round 94 papers remained, and each of these remaining papers was reviewed by at least two additional reviewers. After the second round, 65 papers remained for discussions during the face to face TPC meeting, which was attended by most of the 42 TPC members. After this meeting 39 papers were selected, which is just below 20% acceptance rate. Sheppards were assigned to each of these papers, to further improve quality. For the first time this year IMC followed a ”double-blind” review process. Another important issue for IMC is that measurements should be performed in an ethical defendable way; whenever appropriate papers have to discuss the ethical implications of the research.

Although IMC has a broad scope on measurements, as can be seen in Figure 17 the most prominent category of papers focuses on (Internet) security. Therefore IMC is amongst the key conferences CONCORDIA targets at.

8.3 Organization of the Scientific Community

The organization of the scientific community involves the following activities:

- Chairing professional organizations (ACM, IEEE and IFIP).
- Editing scientific journals.
- Steering committee membership.

Chairing Professional Organizations

The scientific community is organized by professional organizations like ACM, IEEE and IFIP. CONCORDIA researchers hold the following positions in these organizations:

- Chair of IFIP TC6[5] (Burkhard Stiller, UZH)
  The Technical Committee 6 (TC6 - Communications Systems) is one of the largest TCs within IFIP in terms of activities and revenues. TC6 has nine Working Groups (WGs) as well as a number of Special Interest Groups (SIGs), the majority of which are concerned either with specific aspects of communications systems themselves or with the application of communications systems. In addition, one WG focuses on communications in developing countries. TC6 meets twice a year, in spring and fall, usually co-locating its meetings with a related conference.
- Chair of IFIP TC6 Working Group 6.6[6] (Rémi Badonnel, UL)
  The Working Group 6 of IFIP TC6 focusses on management of network and distributed systems. Management is defined in five functional areas: Fault management, Configuration management, Accounting management, Performance management and Security management (FCAPS). Of these, security management has become the greatest challenge.

[5] https://ifip.informatik.uni-hamburg.de/ifip/tc/6

[6] https://ifip.informatik.uni-hamburg.de/ifip/tc/6/wg/6.6/officers

Editing Scientific Journals

In total nine CONCORDIA researchers act as Series Editor, Associate Editor, Area Editor, Guest Editor, Editorial Board member or Editorial Advisory Board member for the following journals:

- ACM SIGCOMM Computer Communication Review
- ACM Transactions on Data Science
- IEEE Access
- IEEE Communications Magazine - Network and Service Management series
- IEEE Computing
- IEEE Transactions on Network and Service Management
- IEEE Transactions on Service Computing
- Springer - Data science and engineering
- Springer - Journal of Network and Systems Management
- Wiley - International Journal of Network Management
- Hindawi - Mobile Information Systems
- Frontiers Media SA - Cybersecurity and Privacy of Frontiers in Big Data
- World Scientific - International Journal of Cooperative Information Systems

Details are shown in Appendix D. Again, it is interesting to note that of the 22 positions, 10 are taken by women. This is again an indication that women involved in CONCORDIA have a strong leadership role in research.

In addition to (permanent) positions within editorial boards, in 2019 CONCORDIA researchers have also acted as guest editors for:

- IEEE Transactions on Network and Service Management - Special Issue on Cybersecurity Techniques for Managing Networked Systems[7] (Rémi Badonnel, UL)
- Elsevier Pervasive and Mobile Computing - Security and Privacy in Edge Computing-Assisted Internet of Things (IoT)[8] (Claudio Ardagna, UMIL)
- Future Generation Computer Systems - Trusted Cloud-Edges (CE) Computations (Claudio Ardagna, UMIL)

[7] https://www.comsoc.org/publications/journals/ieee-tnsm/cfp/cybersecurity-techniques-managing-networked-systems

[8] https://www.sciencedirect.com/journal/pervasive-and-mobile-computing/special-issue/1020ZZFQS7T

Steering Committee Membership

CONCORDIA researchers are member of the Steering Committee for the following conferences:

- ACM SACMAT - Symposium on Access Control Models and Technologies[9] (Elena Ferrari, UI)
- IFIP/IEEE IM - International Symposium on Integrated Network Management[10] (Aiko Pras, UT and Olivier Festor, UL)
- IFIP TMA - Traffic Measurement and Analysis Conference[11] (Anna Sperotto, UT)

[9] http://www.sacmat.org/2020/officers.php

[10] https://im2019.ieee-im.org/steering-committee

[11] https://tma.ifip.org/

9 Contributions to Standards and Open Research Data

One of the objectives of WP1 is to contribute to standardization, open research data and code, shared via systems like GitHub. This section summarizes the achievements in this area *by WP1 researchers*.

It should be noted that CONCORDIA has special tasks for standardization as well as open data:

- Task 5.3: Certification and Standardization activities
- Task 6.4: Data management

The complete overview of all CONCORDIA activities related to standardization and open data is therefore included in deliverables of WP5 and WP6.

9.1 Standardization Activities Performed by WP1 Researchers

Standardization is an effort that takes many years from initial research till full standards. The activities identified in this subsection are therefore activities that started already years before, but obtained important outcome in 2019 and have some relationship with current WP1 research. The results include three Request for Comments (RFCs) and one Internet-Draft.

YANG Library (RFC 8525)

There is a need for a standard mechanism to expose which YANG modules [19], datastores [20], and datastore schemas [20] are in use by a network management server.

RFC 8525 [17] defines the YANG module ”ietf-yang-library” that provides this information. This version of the YANG library is compatible with the Network Management Datastore Architecture (NMDA) [20]. The previous version of the YANG library, defined in [RFC7895], is not compatible with the NMDA since it assumes that all datastores have exactly the same schema. This is not necessarily true in the NMDA since dynamic configuration datastores may have their own datastore schema. Furthermore, the operational state datastore may support non-configurable YANG modules in addition to the YANG modules supported by conventional configuration datastores.

The old YANG library definitions have been retained (for backwards-compatibility reasons), but the definitions have been marked as deprecated. For backwards compatibility, an NMDA server SHOULD populate the deprecated ”/modules-state” tree in a backwards-compatible manner. The new ”/yang-library” tree will be ignored by legacy clients but will provide all the data needed for NMDA-aware clients (which will ignore the ”/modules-state” tree). The recommended approach to populate ”/modules-state” is to report the YANG modules with ”config true” data nodes that are configurable via conventional configuration datastores and the

YANG modules with ”config false” data nodes that are returned via a Network Configuration Protocol (NETCONF) `<get>` operation or equivalent.

The YANG library information can be different on every server, and it can change at runtime or across a server reboot. If a server implements multiple network management protocols to access the server’s datastores, then each such protocol may have its own conceptual instantiation of the YANG library.

If a large number of YANG modules are utilized by a server, then the YANG library contents can be relatively large. Since the YANG library contents change very infrequently, it is important that clients be able to cache the YANG library contents and easily identify whether their cache is out of date.

NETCONF Extensions to Support the Network Management Datastore Architecture (RFC 8526)

RFC 8526 [21] extends the NETCONF protocol defined in [42] in order to support the Network Management Datastore Architecture (NMDA) defined in [20].

This RFC updates [42] in order to enable NETCONF clients to interact with all the datastores supported by a server implementing the NMDA. The update both adds new `<get-data>` and `<edit-data>` operations and augments existing `<lock>`, `<unlock>`, and `<validate>` operations.

This RFC also updates [19] in order to enable NETCONF clients to both discover which datastores are supported by the NETCONF server and determine which modules are supported in each datastore. The update requires NETCONF servers implementing the NMDA to support the YANG library [17].

RESTCONF Extensions to Support the Network Management Datastore Architecture (RFC 8527)

RFC 8527 [22] extends the RESTCONF protocol defined in [18] in order to support the Network Management Datastore Architecture (NMDA) defined in [20].

This RFC updates [18] in order to enable RESTCONF clients to discover which datastores are supported by the RESTCONF server, determine which modules are supported in each datastore, and interact with all the datastores supported by the NMDA. Specifically, the update introduces new datastore resources, adds a new query parameter, and requires the usage of the YANG library [17] by RESTCONF servers implementing the NMDA.

The solution presented in this RFC is backwards compatible with [18]. This is achieved by only adding new resources and leaving the semantics of the existing resources unchanged.

Considerations for Large Authoritative DNS Servers Operators

This Internet-Draft[12] summarizes recent research work exploring DNS configurations and offers specific tangible considerations to DNS authoritative server operators (DNS operators hereafter). The considerations (C1-C5) presented in this Internet-Draft are backed by previous research work, which used wide-scale Internet measurements upon which to draw their conclusions. This Internet-Draft describes the key engineering options, and points readers to the pertinent papers for details and other research works related to each consideration here presented.

These considerations are designed for operators of “large” authoritative servers. In this context, “large” authoritative servers refers to those with a significant global user population, like top-level domain (TLD) operators, run by a single or multiple operators. These considerations may not be appropriate for smaller domains, such as those used by an organization with users in one city or region, where goals such as uniform low latency are less strict.

It is likely that these considerations might be useful in a wider context, such as for any stateless/short-duration, anycasted service. Because the conclusions of the studies don’t verify this fact, the wording in this Internet-Draft discusses DNS authoritative services only. This Internet-Draft is not an IETF consensus document: it is published for informational purposes.

9.2 Open Research Data Provided by WP1 Researchers

WP1 researchers are responsible for the creation and maintenance of several open network data repositories. This subsection identifies these repositories. The first is a major dataset containing 3.6 trillion data points collected since the start in 2015.

OpenINTEL

The goal of the OpenINTEL[13] measurement platform is to capture daily snapshots of the state of large parts of the global Domain Name System (DNS). Because the DNS plays a key role in almost all Internet services, by analyzing DNS information we are able to detect various types of security issues, such as the creation of SPAM and phishing domains, or the adaption of DDoS protection services. By capturing DNS data and recording it over longer periods of time, OpenINTEL allows us to track the evolution of the DNS system. By performing active measurements, rather than passively collecting DNS data, we build consistent and reliable time series of the state of the DNS. 227 Million domains are measured on a daily basis, resulting in 2.4 billion data points per day. Since the start of OpenINTEL in 2015 3.6 trillion data points have been collected.

[12] https://datatracker.ietf.org/doc/draft-moura-dnsop-authoritative-recommendations/

[13] https://openintel.nl/

The motivation and details regarding OpenINTEL can be found in the paper "A High-Performance, Scalable Infrastructure for Large-Scale Active DNS Measurements" [136].

Trace-Share

Research validation and verification are fundamental principles of good scientific work. In terms of research in the area of network traffic measurement and analysis, however, these principles pose a great challenge. The research heavily depends not only on the correct processes of data usage but also on the availability of network traffic datasets that meet the common requirements and are publicly available. Without these datasets, we will never be able to reliably repeat, validate, and analyze research results.

The main idea of Trace-Share[14] is based on annotated units of network traffic that can be synthetically generated, or derived from real-world traffic. These units typically contain only a minimum of personal data, so they can be shared and, thanks to the restrictions on the inclusion of interest-related traffic only, be easily annotated. They can be also easily normalized and combined with each other or with a real-world traffic to create semi-labeled datasets.

Details and main ideas of the project are available in the publicly available publication "Towards Provable Network Traffic Measurement and Analysis via Semi-Labeled Trace Datasets" [31].

Ethereum Datasets

A Zenodo hosted repository[15] contains a dataset of transactions of 10 Ethereum addresses controlled by a private key. Each has at least 2000 output transactions, which include a transfer of cryptocurrency. All transactions are performed within no longer than three months period.

[14] https://github.com/Trace-Share/

[15] https://doi.org/10.5281/zenodo.3557460

10 Conclusions and Outlook

In many cases EU projects have, already before their start, a clear vision of what should be the outcome of the research. CONCORDIA is different, in the sense that its goal is not to design some novel architectures and prototypes, but to build a European secure, resilient and trusted *ecosystem* for innovation in the area of cybersecurity. For this purpose it brings together 55 partners, of which 28 from academia, and stimulates collaboration between these partners to increase impact.

WP1 has therefore used the first year of CONCORDIA to get organized and find collaborations. The email aliases and the information sharing portal (Confluence and GitLab) have been created. Management of the WP has been set up, which involved some changes in the leadership of two tasks since a) a task leader moved to another employer (T1.2) and b) to reflect the wish to make organizations with most PMs responsible for coordinating their task (T1.3). To allow participants to learn each other and to establish collaboration, meetings have been organized in Munich, Bremen, Brussels, Luxembourg and Vienna. Some of these meetings focused on seeking internal WP1 collaboration, whereas others focused on seeking collaboration between WP1 and the seven CONCORDIA pilots. In addition to the above meetings, dedicated meetings with specific pilots took place in July in Munich with WP3, in November in Ottobrunn with T2.5 and one more will take place early 2020 in Hamburg with T3.1.

At the end of year one we may conclude that WP1 is doing well. The objectives for the first year have been reached or, with respect to the number of publications, even exceeded. For example, the SMART objective for WP1 is to publish at least 100 scientific papers over a period of 4 years; after the first year already more than 50 papers have been published at workshops, conferences and in journals and transactions. Some of these papers were published at top venues, such as the *ACM Internet Measurement Conference* (ACM-IMC), *ACM Transactions on Cyber-Physical Systems* (ACM-TCPS), *IEEE Transactions on Sustainable Computing* (IEEE-T-SUSC) and *IEEE Transactions on Network and Service Management* (IEEE-TNSM). CONCORDIA researchers are active in organizing conferences and journals as well as the scientific community. Top-level conferences, such as ACM-IMC, have been sponsored by CONCORDIA. An interesting observation is that women within CONCORDIA seem to be specifically successful in editorial boards and steering committees. CONCORDIA researchers also contributed to the development of Internet Standards (3 RFCs and 1 Internet Draft).

Collaboration with the pilots have been established and will be strengthened in the coming years. Whereas the first year was used to explore collaboration opportunities for the whole set of research activities, the next years will focus on selected collaborations with a high potential for impact. For example, WP1 will strengthen its collaboration in the area of DDoS protection with T3.2, researchers within WP1 will start analyzing the threat intelligence data that is collected in T3.1, and collaboration will be strengthened with the sector specific pilots of T2.1 (Telco), T2.2 (Finance), T2.4 (eHealth) and T2.5 (Defense). To increase focus, partners with PMs spread over multiple WP1 tasks will concentrate these PMs in a smaller number of tasks, preferably one.

Now that the quantity of papers is above expectation, the emphasis will shift towards higher quality. Such shift seems realistic, since the project has started only recently and publishing top research at top venues generally takes a long time. At this moment CONCORDIA's top publications focus on venues where security is *applied* within a specific context (like ACM-IMC, ACM-TCPS, IEEE-T-SUSC and IEEE-TNSM), which indicates that CONCORDIA research is inspired by, and wants to have impact on, real security problems. In the next year(s) we expect to have papers *also* at the *traditional* security venues, like the USENIX Security Symposium, IEEE Symposium on Security and Privacy (S&P), the ACM Conference on Computer and Communications Security (CCS) and Network and Distributed System Security Symposium (NDSS).

References

[1] Arduino Cryptography Library. https://github.com/rweather/arduinolibs. Last visit: August 29, 2019.

[2] Capital one data theft impacts 106m people. https://krebsonsecurity.com/2019/07/capital-one-data-theft-impacts-106m-people/. Accessed: 2019-11-08.

[3] S. Abt and H. Baier. Are We Missing Labels? A Study of the Availability of Ground-Truth in Network Security Research. In *2014 Third International Workshop on Building Analysis Datasets and Gathering Experience Returns for Security (BADGERS)*, pages 40–55. IEEE, sep 2014.

[4] M. Z. Alom, B. Carminati, and E. Ferrari. Adapting users' privacy preferences in smart environments. In *2019 IEEE International Congress on Internet of Things (ICIOT)*, pages 165–172, July 2019.

[5] M. Z. Alom, B. Carminati, and E. Ferrari. Helping users managing context-based privacy preferences. In *2019 IEEE International Conference on Services Computing (SCC)*, pages 100–107, July 2019.

[6] K. Andersen, R. Medaglia, R. Vatrapu, H. Zinner Henriksen, and R. Gauld. *The Forgotten Promise of E-Government Maturity: Assessing Responsiveness in the Digital Public Sector. Government Information Quarterly, Vol. 28, Issue 4*, pages 439–445, October 2011.

[7] M. Anisetti, C. Ardagna, E. Damiani, F. Gaudenzi, and G. Jeon. Cost-effective deployment of certified cloud composite services. *Journal of Parallel and Distributed Computing*, 2019. to appear.

[8] M. Anisetti, C. Ardagna, E. Damiani, and A. Sala. A trust assurance technique for internet of things based on human behavior compliance. *Concurrency Computation: Practice and Experience*, 2019.

[9] M. Anisetti, C. A. Ardagna, E. Damiani, and F. Gaudenzi. A continuous certification methodology for devops. In *Proc. of The 11th ACM International Conference on Management of Digital EcoSystems (ACM MEDES 2019)*, Limassol, Cyprus, November 2019. Invited paper.

[10] M. Anisetti, C. A. Ardagna, E. Damiani, F. Gaudenzi, and G. Jeon. Cost-effective deployment of certified cloud composite services. *Journal of Parallel and Distributed Computing*, 2019.

[11] M. Apostolaki, A. Zohar, and L. Vanbever. Hijacking Bitcoin: Routing Attacks on Cryptocurrencies. In *2017 IEEE Symposium on Security and Privacy (SP)*, pages 375–392, San Jose, CA, USA, May 2017. IEEE.

[12] C. Ardagna, R. Asal, E. Damiani, N. El Ioini, and C. Pahl. Trustworthy IoT: An evidence collection approach based on smart contracts. In *Proc. of the 15th IEEE International Conference on Services Computing (SCC 2019)*, Milan, Italy, July 2019.

[13] S. Arnautov, B. Trach, F. Gregor, T. Knauth, A. Martin, C. Priebe, J. Lind, D. Muthukumaran, D. O'keeffe, M. Stillwell, and others. SCONE: Secure Linux Containers with Intel SGX. In *OSDI*, volume 16, pages 689–703, 2016.

[14] Z. Ayyub and R. Miao. Simple-fying middlebox policy enforcement using sdn. In *ACM SIGCOMM Computer Communication Review*, volume 43, Oct. 2013.

[15] V. Bajpai, O. Bonaventure, K. claffy, and D. Karrenberg. Encouraging Reproducibility in Scientific Research of the Internet. *Dagstuhl Reports*, 8(10):41–62, Jan 2019.

[16] D. J. Bernstein, N. Duif, T. Lange, P. Schwabe, and B. Yang. High-speed High-security Signatures. *Journal of Cryptographic Engineering*, 2(2):77–89, September 2012.

[17] A. Bierman, M. Björklund, J. Schönwälder, K. Watsen, and R. Wilton. YANG Library. RFC 8525, YumaWorks, Tail-f Systems, Jacobs University, Juniper Networks, Cisco Systems, Jan. 2019.

[18] A. Bierman, M. Björklund, and K. Watsen. RESTCONF Protocol. RFC 8040, YumaWorks, Tail-f Systems, Juniper Networks, Jan. 2017.

[19] M. Björklund. The YANG 1.1 Data Modeling Language. RFC 7950, Tail-f Systems, Aug. 2016.

[20] M. Björklund, J. Schönwälder, P. Shafer, K. Watsen, and R. Wilton. Network Management Datastore Architecture (NMDA). RFC 8342, Tail-f Systems, Jacobs University, Juniper Networks, Cisco Systems, Mar. 2018.

[21] M. Björklund, J. Schönwälder, P. Shafer, K. Watsen, and R. Wilton. NETCONF Extensions to Support the Network Management Datastore Architecture. RFC 8526, Tail-f Systems, Jacobs University, Juniper Networks, Cisco Systems, Jan. 2019.

[22] M. Björklund, J. Schönwälder, P. Shafer, K. Watsen, and R. Wilton. RESTCONF Extensions to Support the Network Management Datastore Architecture. RFC 8527, Tail-f Systems, Jacobs University, Juniper Networks, Cisco Systems, Jan. 2019.

[23] B. Blanchet. Automatic verification of security protocols in the symbolic model: The verifier proverif. In *Foundations of Security Analysis and Design VII*, pages 54–87. Springer, 2013.

[24] C. Bormann, M. Ersue, and A. Keranen. RFC 7228: Terminology for Constrained-Node Networks. *IETF Request For Comments*, 2014.

[25] A. Bucevschi, G. Balan, and D. Prelipcean. Preventing file-less attacks with machine learning techniques. In *Proc. of the 21st International Symposium on Symbolic and Numeric Algorithms for Scientific Computing, (SYNASC'19)*, 2019.

[26] B. Carminati, P. Colombo, E. Ferrari, and G. Sagirlar. Enhancing user control on personal data usage in internet of things ecosystems. In *2016 IEEE International Conference on Services Computing (SCC)*, pages 291–298, June 2016.

[27] B. Carminati, E. Ferrari, and C. Rondanini. Blockchain as a platform for secure inter-organizational business processes. In *2018 IEEE 4th International Conference on Collaboration and Internet Computing (CIC)*, pages 122–129, Oct 2018.

[28] B. Carminati, C. Rondanini, and E. Ferrari. Confidential business process execution on blockchain. In *2018 IEEE International Conference on Web Services (ICWS)*, pages 58–65, July 2018.

[29] F. K. Carvalho Ota, J. A. Meira, C. R. Cassagnes, and R. State. Mobile app to sgx enclave secure channel. In *2019 IEEE International Symposium on Software Reliability Engineering Workshops (ISSREW)*. IEEE, 2019.

[30] R. L. Castro, C. Schmitt, and G. Dreo. Aimed: Evolving malware with genetic programming to evade detection. In *2019 18th IEEE International Conference On Trust, Security And Privacy In Computing And Communications/13th IEEE International Conference On Big Data Science And Engineering (TrustCom/BigDataSE)*, pages 240–247. IEEE, 2019.

[31] M. Cermak, T. Jirsik, P. Velan, J. Komarkova, S. Spacek, M. Drasar, and T. Plesnik. Towards provable network traffic measurement and analysis via semi-labeled trace datasets. In *2018 Network Traffic Measurement and Analysis Conference (TMA)*, pages 1–8, June 2018.

[32] T. Chung, E. Aben, T. Bruijnzeels, B. Chandrasekaran, D. Choffnes, D. Levin, B. Maggs, A. Mislove, R. Rijswijk-Deij, J. Rula, and N. Sullivan. RPKI is Coming of Age: A Longitudinal Study of RPKI Deployment and Invalid Route Origins. In *Proc. of the 2019 ACM Internet Measurement Conference (IMC'19)*, pages 406–419, 2019.

[33] T. Chung, E. Aben, T. Bruijnzeels, B. Chandrasekaran, D. Choffnes, D. Levin, B. M. Maggs, A. Mislove, R. v. Rijswijk-Deij, J. Rula, and N. Sullivan. Rpki is coming of age: A longitudinal study of rpki deployment and invalid route origins. In *Proceedings of the Internet Measurement Conference*, IMC '19, pages 406–419, New York, NY, USA, 2019. ACM.

[34] M. Compastié, R. He, M. Kassi-Lahlou, R. Badonnel, and O. Festor. Unikernel-based Approach for Software-Defined Security in Cloud Infrastructures. In *Proc. of 2018 IEEE/IFIP Network Operations and Management Symposium (NOMS)*, Taipei, Taiwan, Apr. 2018.

[35] M. Conti, L. V. Mancini, R. Spolaor, and N. V. Verde. Analyzing android encrypted network traffic to identify user actions. *IEEE Transactions on Information Forensics and Security*, 11(1):114–125, 2015.

[36] B. de Matos Patrocínio dos Santos, B. Dzogovic, B. Feng, V. T. Do, N. Jacot, and T. V. Do. Towards Achieving a Secure Authentication Mechanism for IoT Devices in 5G Networks. In *6th IEEE International Conference on Cyber Security and Cloud Computing, CSCloud 2019 / 5th IEEE International Conference on Edge Computing and Scalable Cloud, EdgeCom 2019, Paris, France, June 21-23, 2019*, pages 130–135, 2019.

[37] Die Schweizerische Bundeskanzlei. *Vote électronique*. [Online] http://pvpf.ch/ve, last visit December 7, 2019.

[38] B. Dzogovic, P. d. S. Bernardo de Matos, T. van Do, D. van Thuan, B. Feng, and N. Jacot. Bringing 5G into User's Smart Home. In *17th IEEE Intl Conf on Pervasive Intelligence and Computing*, Fukuoka, 2019.

[39] B. Dzogovic, B. Santos, V. T. Do, B. Feng, N. Jacot, and T. Van Do. Connecting remote enodeb with containerized 5g c-rans in openstack cloud. In *2019 6th IEEE International Conference on Cyber Security and Cloud Computing (CSCloud)/ 2019 5th IEEE International Conference on Edge Computing and Scalable Cloud (EdgeCom)*, pages 14–19, June 2019.

[40] B. Dzogovic, D. Thanh, B. Santos, V. Do, B. Feng, and N. Jacot. Thunderbolt-3 backbone for augmented 5g network slicing in cloud-radio access networks. In *2019 IEEE 2nd 5G World Forum (5GWF)*, pages 415–420. IEEE, Sept. 2019.

[41] E-Voting-Moratorium. *Initiativtext*. [Online] http://pvpf.ch/evmor, last visit December 7, 2019.

[42] R. Enns, M. Bjorklund, J. Schönwälder, and A. Bierman. Network Configuration Protocol (NETCONF). RFC 6241, Juniper Networks, Tail-f Systems, Jacobs University, Brocade, June 2011.

[43] I. Eyal and E. G. Sirer. Majority Is Not Enough: Bitcoin Mining Is Vulnerable. In N. Christin and R. Safavi-Naini, editors, *Financial Cryptography and Data Security*, volume 8437, pages 436–454. Springer Berlin Heidelberg, Berlin, Heidelberg, 2014.

[44] S. K. Fayazbakhsh, L. Chiang, V. Sekar, M. Yu, and J. C. Mogul. Enforcing Network-Wide Policies in the Presence of Dynamic Middlebox Actions using FlowTags. In *Proc. 11th USENIX Symposium on Networked Systems Design and Implementation (NSDI'14)*, pages 543–546, Seattle, WA, USA, 2014.

[45] T. M. Fernández-Caramés and P. Fraga-Lamas. A Review on the Use of Blockchain for the Internet of Things. *IEEE Access*, 6:32979–33001, 2018.

[46] Forum of Incident Response and Security Teams (FIRST). Computer Security Incident Response Team (CSIRT) Services Framework Version 2.0. https://www.first.org/education/csirt_services_framework_v2.0, June 2019.

[47] N. Foster, M. J. Freedman, R. Harrison, C. Monsanto, and D. Walker. Frenetic, a Network Programming Language. In *Proceedings of the 16th ACM SIGPLAN International Conference on Functional Programming (ICFP'11)*, 2011.

[48] L. P. Fraile, A. P. Fournaris, and O. Koufopavlou. Revisiting rowhammer attacks in embedded systems. In *2019 14th International Conference on Design Technology of Integrated Systems In Nanoscale Era (DTIS)*, pages 1–6, April 2019.

[49] M. Franco, B. Rodrigues, and B. Stiller. MENTOR: The Design and Evaluation of a Protection Services Recommender System. In *Proceedings of the 15th IFIP/IEEE CNSM 2019*, Piscataway, New Jersey, US, October 2019. IEEE.

[50] M. Franco, B. Rodrigues, and B. Stiller. Mentor: The design and evaluation of a protection services recommender system. In *15th International Conference on Network and Service Management (CNSM 2019)*, pages 1–7, Halifax, Canada, October 2019. IFIP.

[51] C. Grunspan and R. Pérez-Marco. On profitability of selfish mining. *arXiv:1805.08281 [cs, math]*, May 2018.

[52] C. J. Guillaume Bonfante and, J.-Y. Marion, and F. Sabatier. Lockergoga quickly reversed. In *Malcon Conference*, 2019.

[53] M. Hamad, M. Tsantekidis, and V. Prevelakis. Red-zone: Towards an intrusion response framework for intra-vehicle system. In *5th International Conference on Vehicle Technology and Intelligent Transport Systems (VEHITS)*, Crete, Greece, May 2019.

[54] G. Hatzivasilis, I. Askoxylakis, G. Alexandris, D. Anicic, A. Bröring, V. Kulkarni, K. Fysarakis, and G. Spanoudakis. The interoperability of things: Interoperable solutions as an enabler for IoT and Web 3.0. In *2018 IEEE 23rd International Workshop on Computer Aided Modeling and Design of Communication Links and Networks (CAMAD)*, pages 1–7. IEEE, 2018.

[55] G. Hatzivasilis, P. Chatziadam, A. Miaoudakis, E. Lakka, S. Ioannidis, A. Alessio, M. Smyrlis, G. Spanoudakis, A. Yautsiukhin, M. Antoniou, and N. Stathiakis. Towards the insurance of healthcare systems. In *1st Model–driven Simulation and Training Environments for Cybersecurity (MSTEC), ESORICS, Springer, Luxembourg, 27 September, 2019*, pages 1–14, 09 2019.

[56] G. Hatzivasilis, P. Chatziadam, N. Petroulakis, S. Ioannidis, M. Mangini, C. Kloukinas, A. Yautsiukhin, M. Antoniou, D. G. Katehakis, and M. Panayiotou. Cyber insurance of information systems: Security and privacy cyber insurance contracts for ict and helathcare organizations. In *2019 IEEE 24th International Workshop on Computer Aided Modeling and Design of Communication Links and Networks (CAMAD)*, pages 1–6, Sep. 2019.

[57] G. Hatzivasilis, E. L. O. Soultatos, S. Ioannidis, D. Anicic, L. C. A. Broring, and G. S. M. Falchetto, K. Fysarakis. Secure semantic interoperability for IoT applications with linked data. In *Proceedings of the IEEE Global Communications Conference (GLOBECOM 2019),9-13December,2019,pp.1-7.*, Waikoloa, HI, USA, 2019.

[58] G. Hatzivasilis, O. Soultatos, S. Ioannidis, G. Spanoudakis, G. Demetriou, and V. Katos. Mobiletrust: Secure knowledge integration in vanets. *ACM Transactions on Cyber-Physical Systems*, 4(2), 2019.

[59] E. Heilman, A. Kendler, A. Zohar, and S. Goldberg. Eclipse attacks on bitcoin's peer-to-peer network. In *Proceedings of the 24th USENIX Conference on Security Symposium*, SEC'15, pages 129–144, Berkeley, CA, USA, 2015. USENIX Association.

[60] G. Hurel, R. Badonnel, A. Lahmadi, and O. Festor. Towards Cloud Based Compositions of Security Functions for Mobile Devices. In *IFIP/IEEE International Symposium on Integrated Network Management (IM'15)*, 2015.

[61] A. K. Iannillo and R. State. A proposal for security assessment of trustzone-m based software. In *2019 IEEE International Symposium on Software Reliability Engineering Workshops (ISSREW)*. IEEE, 2019.

[62] C. Iordanou, N. Kourtellis, J. Carrascosa, C. Soriente, R. C. Rumin, and N. Laoutaris. Beyond content analysis: Detecting targeted ads via distributed counting. In *Proceedings of the International Conference on emerging Networking EXperiments and Technologies (CONEXT)*. ACM, 2019.

[63] W. A. Jansen. Directions in security metrics research. Technical Report NISTIR 7564, National Institute of Standards and Technology (NIST), 2009. https://nvlpubs.nist.gov/nistpubs/Legacy/IR/nistir7564.pdf.

[64] A. Khalimonenko, O. Kupreev, and E. Badovskaya. DDoS Attacks in Q1 2018. [Online] https://securelist.com/ddos-report-in-q1-2018/85373/, April 2018. last visit September 14, 2019.

[65] C. Killer, B. Rodrigues, and B. Stiller. Security management and visualization in a blockchain-based collaborative defense. In *1st IEEE International Conference on Blockchain and Cryptocurrency (ICBC 2019)*, pages 108–111, Seoul, South Korea, May 2019. IEEE.

[66] C. Killer, B. Rodrigues, and B. Stiller. Threat management dashboard for a blockchain collaborative defense. In *IEEE GLOBECOM Workshop No. 27 on "Blockchain in Telecommunications: Emerging Technologies for the Next Decade and Beyond" (GLOBECOM 2019)*, pages 1–6, Waikoloa, U.S.A., December 2019. IEEE.

[67] C. Killer and B. Stiller. The swiss postal voting process and its system and security analysis. In R. Krimmer, M. Volkamer, V. Cortier, B. Beckert, R. Küsters, U. Serdült, and D. Duenas-Cid, editors, *Electronic Voting*, pages 134–149, Cham, 2019. Springer International Publishing.

[68] C. Killer and B. Stiller. The swiss postal voting process and its system and security analysis. In *4th International Joint Conference on Electronic Voting (E-Vote-ID 2019)*, pages 1–16, Bregenz, Austria, October 2019. Springer.

[69] H. Kim, J. Reich, A. Gupta, M. Shahbaz, N. Feamster, and R. Clark. Kinetic: Verifiable Dynamic Network Control. In *Proceedings of the 12th USENIX Conference on Networked Systems Design and Implementation (NSDI'15)*, 2015.

[70] D. Kopp, M. Wichtlhuber, I. Poese, J. Santanna, O. Hohlfeld, and C. Dietzel. Ddos hide & seek: On the effectiveness of a booter services takedown. In *Proceedings of the Internet Measurement Conference*, IMC '19, pages 65–72, New York, NY, USA, 2019. ACM.

[71] D. Kopp, M. Wichtluber, I. Poese, J. Santanna, O. Highlfeld, and C. Dietzel. DDoS Hide & Seek: On the Effectiveness of a Booter Services Takedown. In *Proc. of the 2019 ACM Internet Measurement Conference (IMC'19)*, pages 65–72, 2019.

[72] R. Krimmer, S. Triessnig, and M. Volkamer. *The Development of Remote E-Voting Around the World: A Review of Roads and Directions*. *First International Joint Conference on Electronic Voting and Identity (E-VOTE ID 2007). Bochum, Germany, October 2008*, pages 1–15, 2007.

[73] R. Labaca-Castro, B. Biggio, and G. Dreo Rodosek. Poster: Attacking malware classifiers by crafting gradient-attacks that preserve functionality. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, CCS '19, pages 2565–2567, New York, NY, USA, 2019. ACM.

[74] E. Lakka, N. E. Petroulakis, G. Hatzivasilis, O. Soultatos, M. Michalodimitrakis, U. Rak, K. Waledzik, D. Anicic, and V. Kulkarni. End-to-end semantic interoperability mechanisms for IoT. In *2019 IEEE 24th International Workshop on Computer Aided Modeling and Design of Communication Links and Networks (CAMAD)*, pages 1–6. IEEE, 2019.

[75] K. Lampropoulos, G. Georgakakos, and S. Ioannidis. Using blockchains to enable big data analysis of private information. In *2019 IEEE 24th International Workshop on Computer Aided Modeling and Design of Communication Links and Networks (CAMAD)*, pages 1–6. IEEE, 2019.

[76] V. G. Li, M. Dunn, P. Pearce, D. McCoy, G. M. Voelker, and S. Savage. Reading the tea leaves: A comparative analysis of threat intelligence. In *28th USENIX Security Symposium (USENIX Security 19)*, pages 851–867, Santa Clara, CA, Aug. 2019. USENIX Association.

[77] S. Limnaios, N. Sklavos, and O. Koufopavlou. Lightweight efficient simeck32/64 crypto-core designs and implementations, for iot security. In *27th IFIP/IEEE International Conference On Very Large Scale Integration (VLSI-SoC'19)*, 10 2019.

[78] O. López-Pintado, L. García-Bañuelos, M. Dumas, I. Weber, and A. Ponomarev. Caterpillar: A business process execution engine on the ethereum blockchain. *Software: Practice and Experience*, 49(7):1162–1193, 2019.

[79] S. Mannhart, B. Rodrigues, E. Scheid, S. S. Kanhere, and B. Stiller. Toward mitigation-as-a-service in cooperative network defenses. In *2018 IEEE 4th Intl Conf on Big Data Intelligence and Computing and Cyber Science and Technology Congress (CyberSciTech 2018)*, pages pp. 362–367, Aug 2018. Athens, Greece.

[80] Y. Marcus, E. Heilman, and S. Goldberg. Low-resource eclipse attacks on ethereum's peer-to-peer network. *IACR Cryptology ePrint Archive*, 2018:236, 2018.

[81] J.-Y. Marion and F. Sabatier. Lockergoga is smashed by gorille. [Online] http://www.cyber-detect.com/fr-blog.html, last visited December 16, 2019.

[82] Y. Meidan, M. Bohadana, Y. Mathov, Y. Mirsky, A. Shabtai, D. Breitenbacher, and Y. Elovici. N-baiot—network-based detection of iot botnet attacks using deep autoencoders. *IEEE Pervasive Computing*, 17(3):12–22, 2018.

[83] J. Mendling, I. Weber, W. V. D. Aalst, J. V. Brocke, C. Cabanillas, F. Daniel, S. Debois, C. D. Ciccio, M. Dumas, S. Dustdar, A. Gal, L. García-Bañuelos, G. Governatori, R. Hull, M. L. Rosa, H. Leopold, F. Leymann, J. Recker, M. Reichert, H. A. Reijers, S. Rinderle-Ma, A. Solti, M. Rosemann, S. Schulte, M. P. Singh, T. Slaats, M. Staples, B. Weber, M. Weidlich, M. Weske, X. Xu, and L. Zhu. Blockchains for business process management - challenges and opportunities. *ACM Trans. Manage. Inf. Syst.*, 9(1):4:1–4:16, Feb. 2018.

[84] P. Metzler, N. Suri, and G. Weissenbacher. Extracting safe thread schedules from incomplete model checking results. In *International Symposium on Model Checking of Software (SPIN)*, 2019.

[85] S. Morgan. 2019 Official Annual Cybercrime Report. *Herjavec Group*, 2019. https://www.herjavecgroup.com/wp-content/uploads/2018/12/CV-HG-2019-Official-Annual-Cybercrime-Report.pdf, last visit June 15, 2019.

[86] A. Mortensen, F. Andreasen, R. Tirumaleswar, C. Gray, R. Compton, and N. Teague. Distributed-Denial-of-Service Open Threat Signaling (DOTS) Architecture. Internet-Draft draft-ietf-dots-architecture-06, Internet Engineering Task Force, Mar. 2018. Work in Progress.

[87] M. Müller, M. Thomas, D. Wessels, W. Hardaker, T. Chung, W. Toorop, and R. Rijswijk-Deij. Roll, Roll, Roll your Root: A Comprehensive Analysis of the First Ever DNSSEC Root KSK Rollover. In *Proc. of the 2019 ACM Internet Measurement Conference (IMC'19)*, pages 1–14, 2019.

[88] M. Müller, M. Thomas, D. Wessels, W. Hardaker, T. Chung, W. Toorop, and R. v. Rijswijk-Deij. Roll, roll, roll your root: A comprehensive analysis of the first ever dnssec root ksk rollover. In *Proceedings of the Internet Measurement Conference*, IMC '19, pages 1–14, New York, NY, USA, 2019. ACM.

[89] C. Natoli and V. Gramoli. The Balance Attack or Why Forkable Blockchains are Ill-Suited for Consortium. In *2017 47th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*, pages 579–590, Denver, CO, USA, June 2017. IEEE.

[90] K. Nayak, S. Kumar, A. Miller, and E. Shi. Stubborn Mining: Generalizing Selfish Mining and Combining with an Eclipse Attack. In *2016 IEEE European Symposium on Security and Privacy (EuroS&P)*, pages 305–320, Saarbrucken, Mar. 2016. IEEE.

[91] L. Nemec Zlatolas, A. Kamišalić, and M. Turkanović. Correlation between students' background and the knowledge on conceptual database modelling. In T. Welzer, J. Eder, V. Podgorelec, R. Wrembel, M. Ivanović, J. Gamper, M. Morzy, T. Tzouramanis, J. Darmont, and A. Kamišalić Latifić, editors, *New Trends in Databases and Information Systems*, pages 45–51, Cham, 2019. Springer International Publishing.

[92] L. Nemec Zlatolas, T. Welzer, M. Hölbl, M. Heričko, and A. Kamišalić. A model of perception of privacy, trust, and self-disclosure on online social networks. *Entropy*, 21(8):772, 2019.

[93] S. R. Niya, S. Allemann, A. Gabay, and B. Stiller. Trademap: A finma-compliant anonymous management of an end-2-end trading market place. In *15th International Conference on Network and Service Management CNSM 2019*, pages 1–4, Halifax, Canada, October 2019. IEEE.

[94] R. Norvill, M. Steichen, W. Shbair, and R. State. Blockchain for the simplification and automation of kyc result sharing. In *IEEE International Conference on Blockchain and Cryptocurrency (ICBC 2019)*. IEEE Xplore, 2019.

[95] G. Oikonomou, J. Mirkovic, P. Reiher, and M. Robinson. A Framework for a Collaborative DDoS Defense. In *2006 22nd Annual Computer Security Applications Conference (ACSAC'06)*, pages 33–42, December 2006.

[96] M. Ollivier, S. Bardin, R. Bonichon, and J. Marion. How to kill symbolic deobfuscation for free (or: unleashing the potential of path-oriented protections). In *Proceedings of the 35th Annual Computer Security Applications Conference, ACSAC 2019, San Juan, PR, USA, December 09-13, 2019*, pages 177–189, 2019.

[97] X. Ou, S. Govindavajhala, and A. W. Appel. Mulval: A logic-based network security analyzer. In *Proceedings of the 14th Conference on USENIX Security Symposium - Volume 14*, SSYM'05, pages 8–8, Berkeley, CA, USA, 2005. USENIX Association.

[98] OVAL. Open Vulnerability and Assessment Language. http://oval.mitre.org/. Last visited on September 2019.

[99] M. Pachilakis, P. Papadopoulos, N. Laoutaris, E. P. Markatos, and N. Kourtellis. Measuring ad value without bankrupting user privacy. *arXiv preprint arXiv:1907.10331*, 2019.

[100] M. Pachilakis, P. Papadopoulos, E. P. Markatos, and N. Kourtellis. No more chasing waterfalls: A measurement study of the header bidding ad-ecosystem. In *Proceedings of the Internet Measurement Conference (IMC)*. ACM, 2019.

[101] D. Palma and T. Spatzier. Topology and orchestration specification for cloud applications (TOSCA). *Organization for the Advancement of Structured Information Standards (OASIS), Tech. Rep*, 2013.

[102] E. Papadogiannaki, C. Halevidis, P. Akritidis, and L. Koromilas. Otter: A scalable high-resolution encrypted traffic identification engine. In *International Symposium on Research in Attacks, Intrusions, and Defenses*, pages 315–334. Springer, 2018.

[103] E. Papadogiannaki and S. Ioannidis. Gpu-accelerated encrypted network traffic inspection. Poster, Sept. 2019. ACM Celebration of Women in Computing: womENcourage.

[104] P. Papadopoulos, N. Kourtellis, and E. Markatos. Cookie synchronization: Everything you always wanted to know but were afraid to ask. In *Proceedings of the World Wide Web Conference (WWW)*. ACM, 2019.

[105] T. Peng, C. Leckie, and K. Ramamohanarao. Survey of Network-based Defense Mechanisms Countering the DoS and DDoS Problems. *ACM Computing Surveys (CSUR)*, 39(1):pp. 03–15, 2007.

[106] N. E. Petroulakis, E. Lakka, E. Sakic, V. Kulkarni, K. Fysarakis, I. Somarakis, J. Serra, L. Sanabria-Russo, D. Pau, M. Falchetto, et al. Semiotics architectural framework: End-to-end security, connectivity and interoperability for industrial IoT. In *2019 Global IoT Summit (GIoTS)*, pages 1–6. IEEE, 2019.

[107] A. Pinto and A. Sieira. Data-Driven Threat Intelligence: Useful Methods and Measurements for Handling Indicators. https://www.first.org/resources/papers/conf2015/first_2015_pinto-alex_sierira-alex_data-driven-threatintelligence_20150619.pdf, 2015.

[108] B. Podgorelec, V. Keršič, and M. Turkanović. Analysis of fault tolerance in permissioned blockchain networks. In *2019 XXVII International Conference on Information, Communication and Automation Technologies (ICAT) special session blockchain technology and its application (BCTA*. IEEE, 2019.

[109] V. Prevelakis, M. Hamad, J. Najar, and I. Spais. Secure data exchange for computationally constrained devices. In *International workshop on Information &amp; Operational Technology (IT &amp; OT) security systems (IOSec 2019)*, Luxembourg, 2019.

[110] S. Rafati-Niya, S. Allemann, A. Gabay, and B. Stiller. Blockchain-based Anonymous P2P Trading System. Technical report, Universitaet Zuerich, Zurich, Switzerland, June 2019.

[111] S. Rafati-Niya, S. Allemann, A. Gabay, and B. Stiller. TradeMap: A FINMA-compliant Anonymous Management of an End-2-end Trading Market Place. In *15th International Conference on Network and Service Management CNSM 2019*, Halifax, Canada, October 2019. IEEE.

[112] S. Rafati-Niya, E. Schiller, I. Cepilov, F. Maddaloni, K. Aydinli, T. Surbeck, T. Bocek, and B. Stiller. Adaptation of proof-of-stake-based blockchains for IoT data streams. In *2019 IEEE International Conference on Blockchain and Cryptocurrency (ICBC)*, pages 15–16, May 2019.

[113] P. Rek, B. Podgorelec, and M. Turkanović. Complexity analysis of decentralized application development using integration tools. In *Proceedings of the Eighth Workshop on Software Quality Analysis, Monitoring, Improvement, and Applications (SQAMIA 2019)*, 2019.

[114] S. Rivera, S. Lagraa, A. K. Iannillo, and R. State. Auto-encoding robot state against sensor spoofing attacks. In *2019 IEEE International Symposium on Software Reliability Engineering Workshops (ISSREW)*. IEEE, 2019.

[115] B. Rodrigues, T. Bocek, A. Lareida, D. Hausheer, S. R. Niya, and B. Stiller. A blockchain-based architecture for collaborative ddos mitigation with smart contracts. In *11th IFIP International Conference on Autonomous Infrastructure, Management and Security (AIMS 2017)*, pages 16–29. Springer, Heidelberg, July 2017.

[116] B. Rodrigues, T. Bocek, and B. Stiller. Enabling a Cooperative, Multi-domain DDoS Defense by a Blockchain Signaling System (BloSS). In *Demonstration Track*, pages 1–3, Singapore, Singapore, Oct 2017. IEEE.

[117] B. Rodrigues, M. Franco, G. Parangi, and B. Stiller. SEConomy: a Framework for the Economic Assessment of Cybersecurity. In *16th International Conference on the Economics of Grids, Clouds, Systems, and Services (GECON 2019)*, pages 1–13, Leeds, UK, September 2019.

[118] B. Rodrigues and B. Stiller. Cooperative Signaling of DDoS Attacks in a Blockchain-based Network. In *Proceedings of the ACM SIGCOMM 2019 Conference Posters and Demos*, SIGCOMM Posters and Demos ’19, pages 39–41, New York, NY, USA, 2019. ACM.

[119] C. Rondanini, B. Carminati, and E. Ferrari. Confidential discovery of iot devices through blockchain. In *2019 IEEE International Congress on Internet of Things (ICIOT)*, pages 1–8, July 2019.

[120] G. Sagirlar, B. Carminati, and E. Ferrari. Decentralizing privacy enforcement for internet of things smart objects. *Computer Networks*, 143:112–125, 2018.

[121] H. Saissi, M. Serafini, and N. Suri. Gyro: A modular scale-out layer for single-server dbmss. In *Symposium on Reliable Distributed Systems (SRDS)*, 2019.

[122] N. Salamanos, M. J. Jensen, X. He, Y. Chen, and M. Sirivianos. On the influence of twitter trolls during the 2016 us presidential election. *Under submission. Available at* *https://arxiv.org/abs/1910.00531v2*, 2019.

[123] J. Santanna. DDoSDB: Collecting and Sharing information of DDoS attacks, 2019. https://ddosdb.org/, last visit June 15, 2019.

[124] B. Santos, B. Dzogovic, B. Feng, V. T. Do, N. Jacot, and T. Van Do. Enhancing Security of Cellular IoT with Identity Federation. In *Advances in Intelligent Networking and Collaborative Systems*, pages 257–268. Springer, 2019.

[125] E. Schiller, S. Rafati-Niya, T. Surbeck, and B. Stiller. Scalable Transport Mechanisms for Blockchain IoT Applications. In *44th IEEE Conference on Local Computer Networks (LCN 2019)*, Piscataway, New Jersey, US, October 2019. IEEE.

[126] N. Schnepf, S. Merz, R. Badonnel, and A. Lahmadi. Automated Verification of Security Chains in Software-Defined Networks with Synaptic. In *Proceedings of the 3rd IEEE Conference on Network Softwarization (NetSoft'17)*, 2017.

[127] O. Schwahn, N. Coppik, S. Winter, and N. Suri. Assessing the state and improving the art of parallel testing for c. In *Proceedings of the ACM SIGSOFT International Symposium on Software Testing and Analysis (ISSTA)*, pages 123–133. ACM, 2019.

[128] K. Shah, A. Salunke, S. Dongare, and K. Antala. Recommender Systems: An Overview of Different Approaches to Recommendations. In *International Conference on Innovations in Information, Embedded and Communication Systems (ICIIECS 2017)*, pages 1–4, Coimbatore, India, March 2017.

[129] W. Shbair, M. Steichen, J. Francois, and R. State. Blockzoom: Large-scale blockchain testbed. In *IEEE International Conference on Blockchain and Cryptocurrency (ICBC 2019)*. IEEE Xplore, 2019.

[130] K. Solomos, P. Ilia, S. Ioannidis, and N. Kourtellis. Talon: An automated framework for cross-device tracking detection. In *22nd International Symposium on Research in Attacks, Intrusions and Defenses (RAID 2019)*. USENIX, 2019.

[131] O. Soultatos, M. Papoutsakis, K. Fysarakis, G. Hatzivasilis, M. Michalodimitrakis, G. Spanoudakis, and S. Ioannidis. Pattern-driven security, privacy, dependability and interoperability management of IoT environments. In *2019 IEEE 24th International Workshop on Computer Aided Modeling and Design of Communication Links and Networks (CAMAD) (IEEE CAMAD 2019)*, Limassol, Cyprus, September 2019.

[132] Staatssekretariat für Wirtschaft SECO, Schweiz. *Nationale E-Government Studie 2019*. [Online] http://pvpf.ch/egov19, last visit December 7, 2019, March 2019.

[133] Swiss Cyber Storm Conference 2019. *Programt*. [Online] http://pvpf.ch/scs19, last visit December 7, 2019.

[134] M. Tsantekidis and V. Prevelakis. Efficient monitoring of library call invocation. In *The 2nd IEEE International Symposium on Future Cyber Security Technologies (FCST)*, 2019.

[135] M. Turkanović, T. Welzer, and M. Hölbl. An example of a cybersecurity education model. In *2019 29th EAEEIE Annual Conference (EAEEIE)*. IEEE, 2019.

[136] R. van Rijswijk-Deij, M. Jonker, A. Sperotto, and A. Pras. A high-performance, scalable infrastructure for large-scale active dns measurements. *IEEE Journal on Selected Areas in Communications*, 34(6):1877–1888, June 2016.

[137] Verein eCH. *eCH-Standards*. [Online] http://pvpf.ch/ech. last visit July 9, 2019.

[138] S. Viţel, G. Balan, and D. Prelipcean. Improving detection of malicious office documents using one-side classifiers. In *Proc. of the 21st International Symposium on Symbolic and Numeric Algorithms for Scientific Computing, (SYNASC'19)*, 2019.

[139] M. Šestak, M. Heričko, T. Welzer Družovec, and M. Turkanović. Applying k-vertex cardinality constraints on a neo4j graph database. In submission, 2019.

[140] C. Wagner, A. Dulaunoy, G. Wagener, and A. Iklody. Misp: The design and implementation of a collaborative threat intelligence sharing platform. In *Proceedings of the 2016 ACM on Workshop on Information Sharing and Collaborative Security*, WISCS '16, pages 49–56, New York, NY, USA, 2016. ACM.

[141] D. Waltermire, S. Quinn, K. Scarfone, and A. Halbardier. The Technical Specification for the Security Content Automation Protocol (SCAP): SCAP version 1.2. *NIST Special Publication*, 800:126, 2011.

[142] H. Wang, L. Barriga, A. Vahidi, and S. Raza. Machine learning for security at the iot edge - a feasibility study. In *2019 International Workshop on Machine Learning Security and Privacy: Experiences and Applications, IEEE MASS*, 11 2019.

[143] I. Weber, X. Xu, R. Riveret, G. Governatori, A. Ponomarev, and J. Mendling. Untrusted business process monitoring and execution using blockchain. In M. La Rosa, P. Loos, and O. Pastor, editors, *Business Process Management*, pages 329–347, Cham, 2016. Springer International Publishing.

[144] M. J. West-Brown, D. Stikvoort, and K.-P. Kossakowski. Handbook for computer security incident response teams (csirts). Technical report, Carnegie Mellon University, Dec. 1998.

[145] A. H. Yaacob, I. K. T. Tan, S. F. Chien, and H. K. Tan. Arima based network anomaly detection. In *2010 Second International Conference on Communication Software and Networks*, pages 205–209, Feb 2010.

[146] C. Yiqun, S. Winter, and N. Suri. Inferring performance bug patterns from developer commits. In *Proceedings of the International Symposium on Software Reliability Engineering (ISSRE)*, 2019.

[147] S. Zannettou, T. Caulfield, E. De Cristofaro, M. Sirivianos, G. Stringhini, and J. Blackburn. Disinformation warfare: Understanding state-sponsored trolls on twitter and their influence on the web. In *Companion Proceedings of The 2019 World Wide Web Conference*, WWW '19, pages 218–226. ACM, 2019.

[148] S. Zannettou, T. Caulfield, W. Setzer, M. Sirivianos, G. Stringhini, and J. Blackburn. Who let the trolls out?: Towards understanding state-sponsored trolls. In *Proceedings of the 10th ACM Conference on Web Science*, pages 353–362. ACM, 2019.

[149] S. T. Zargar, J. Joshi, and D. Tipper. A Survey of Defense Mechanisms Against Distributed Denial of Service (DDoS) Flooding Attacks. *IEEE Communications Surveys Tutorials*, 15(4):pp. 2046–2069, Fourth 2013.

[150] S. Zhao, C. Yiqun, S. Winter, and N. Suri. Analyzing and improving customer-side cloud security certifiability. In *IEEE International Workshop on Software Certification (WoSoCeR)*, 2019.

[151] B. Zhou, D. He, and Z. Sun. Network Traffic Modeling and Prediction with ARIMA/GARCH, 2005.

A Publications

The table on the next pages show all papers produced by CONCORDIA partners. The table shows the task to which the paper belongs, the partner's institute, the year of publication, the title and whether the paper is available as open access paper.

No	Task	Partner	Year	Title	Open Access
1	T1.1	UNIMIL	2019	Trustworthy IoT: An Evidence Collection Approach… [12]	N
2	T1.1	FORTH	2019	MobileTrust: Secure Knowledge Integration… [58]	Link
3	T1.1	FORTH	2019	Pattern-driven Security, Privacy… [131]	N
4	T1.1	UZH	2019	Scalable Transport Mechanisms for Blockchain IoT Applications [125]	Link
5	T1.1	SnT	2019	Security Assessment of Trustzone-M based Software [61]	Link
6	T1.1	ISI, UP	2019	Rowhammer Attacks in Embedded Systems [48]	N
7	T1.1	ISI, UP	2019	Simeck32/64 Crypto-Core… for IoT Security [77]	Link
8	T1.1	FORTH	2019	End-to-End Semantic Interoperability Mechanisms for IoT [74]	N
9	T1.1	UMIL	2019	Trust assurance for IoT based on human behavior compliance [8]	N
10	T1.1	FORTH	2019	Secure Semantic Interoperability for IoT Apps with Linked Data [57]	N
11	T1.2	FORTH	2019	GPU-accelerated encrypted traffic inspection [103]	Link
12	T1.2	UZH	2019	MENTOR… Protection Services Recommender System [50]	Link
13	T1.2	UM	2019	Fault Tolerance in Permissioned Blockchain Networks [108]	N
14	T1.2	RISE	2019	Machine Learning for Security at the IoT Edge… [142]	N
15	T1.2	UZH	2019	Cooperative Signaling of DDoS Attacks … [118]	N
16	T1.2	UT, SIDN	2019	Analysis of the first ever DNSSEC Root KSK Rollover [88]	Link
17	T1.2	UT	2019	RPKI Deployment and Invalid Route Origins [33]	Link
18	T1.2	UT	2019	DDoS… Effectiveness of a Booter Services Takedown [70]	Link
19	T1.2	CODE	2019	Wireless SDN for Highly Utilized MANETs	N
20	T1.3	TUD	2019	Extracting Safe Thread Schedules… Model Checking Results [84]	Link
21	T1.3	TUD	2019	Parallel testing for C [127]	Link
22	T1.3	TUD	2019	Gyro: A Modular Scale-out Layer… Single-Server DBMSs [121]	Link

Continued on next page…

No	Task	Partner	Year	Title	Open Access
23	T1.3	TUD	2019	Inferring Performance Bug Patterns. . . [146]	Link
24	T1.3	SnT	2019	Auto-encoding Robot State against Sensor Spoofing Attacks [114]	Link
25	T1.3	TUD	2019	Customer side cloud certifiability. . . [150]	N
26	T1.3	BD	2019	Preventing File-less Attacks with Machine Learning Techniques [25]	N
27	T1.3	BD	2019	Detection of Malicious Office Documents . . . One-Side Classifiers [138]	N
28	T1.4	TUBS	2019	Red-Zone: Towards an Intrusion Response Framework. . . [53]	Link
29	T1.4	TUBS	2019	Secure Data Exchange for Constrained Devices [109]	N
30	T1.4	TUBS	2019	Efficient Monitoring of Library Call Invocation [134]	N
31	T1.4	UI	2019	Confidential Discovery of IoT Devices. . . [119]	Link
32	T1.4	UP, FORTH	2019	Big Data Analysis of Private Information [75]	N
33	T1.4	UZH	2019	The Swiss Postal Voting Process. . . [68]	Link
34	T1.4	TELENOR, OsloMET	2019	Enhancing Security of Cellular IoT with Identity Federation [124]	N
35	T1.4	TELENOR, OsloMET	2019	Bringing 5G into User's Smart Home [38]	N
36	T1.4	TELENOR, OsloMET	2019	Towards Achieving a Secure Authentication . . . [36]	N
37	T1.4	UMIL	2019	Cost-effective deployment of certified cloud composite services [10]	Link
38	T1.4	UMIL	2019	A Continuous Certification Methodology for DevOps [9]	N
39	T1.4	CODE	2019	AIMED: Evolving Malware. . . to Evade Detection [30]	N
40	T1.4	UZH	2019	Security Management and Visualization. . . Collaborative Defense [65]	N
41	T1.4	CODE	2019	Gradient attacks for Attacking Malware Classifiers [73]	N
42	T1.4	UM	2019	Students' Background and Conceptual Database Modelling [91]	N
43	T1.4	OsloMET, TELENOR	2019	Thunderbolt-3 Backbone for Augmented 5G Network Slicing. . . [40]	N
44	T1.5	TID, FORTH	2019	Cookie synchronization: Everything you always wanted to know. . . [104]	N
45	T1.5	TID, FORTH	2019	Measuring ad value without bankrupting user privacy. . . [99]	Link

Continued on next page. . .

No	Task	Partner	Year	Title	Open Access
46	T1.5	UI	2019	Adapting Users' Privacy Preferences in Smart Environments [4]	N
47	T1.5	UI	2019	Helping Users Managing Context-Based Privacy Preferences [5]	N
48	T1.5	UZH	2019	TradeMap: . . . Trading Market Place [93]	Link
49	T1.5	UZH	2019	Threat Management Dashboard. . . [66]	N
50	T1.5	UM	2019	A Model of Perception of Privacy, Trust, and Self-Disclosure . . . [92]	Link
51	T1.5	UM	2019	Complexity Analysis of Decentralized Application Development . . . [113]	N
52	T1.5	UM	2019	An Example of a Cybersecurity Education Model [135]	N
53	T1.5	UMIL	2019	A Trust Assurance Technique . . . Human Behavior Compliance [8]	N
54	T1.5	SnT	2019	Blockchain. . . Automation of KYC Result Sharing [94]	N
55	T1.5	SnT	2019	BlockZoom: Large-Scale Blockchain Testbed [129]	N
56	T1.5	CUT	2019	Disinformation Warfare: Understanding State-Sponsored Trolls. . . [147]	Link
57	T1.5	CUT	2019	Understanding State-Sponsored Trolls. . . [148]	Link
58	T1.5	TID	2019	Beyond content analysis: Detecting targeted ads via distributed . . . [62]	Link
59	T1.5	TID, FORTH	2019	TALON:. . . Cross-Device Tracking Detection [130]	Link

B Organization of Conferences

The table on the next pages shows all conferences that are or have been organized by CONCORDIA partners. Organization implies one of the following roles: General (co)chair, TPC (co)chair, Program (co)chair, Local chair or Organizing Committee (co)chair. The table is organized in alphabetic order of the first name.

Name	Partner	Scientific Event	Abbr.	Where	When	URL
Anna Sperotto	UT	ACM Internet Measurement Conference 2019	IMC 2019	Amsterdam, Netherlands	21-Oct-19	https://conferences.sigcomm.org/imc/2019/
Anna Sperotto	UT	Passive and Active Measurement (PAM) conference	PAM 2020	Eugene, Oregon	30-Mar-20	https://pam2020.cs.uoregon.edu/
Barbara Carminati	UI	ACM Conference on Data and Application Security and Privacy 2020	CODASPY 2020	New Orleans, LA, USA	16-Mar-20	http://www.codaspy.org/2020/
Barbara Carminati	UI	36th IEEE International Conference on Data Engineering	ICDE 2020	Dallas, Texas, USA	20-Apr-20	https://www.utdallas.edu/icde/
Barbara Carminati	UI	IEEE SERVICES Doctoral Symposium 2019	SERVICES 2019	Milan, Italy	08-Jul-19	https://conferences.computer.org/services/2019/symposia/doctoralsymposium.html
Barbara Carminati	UI	Manifesto for Today - workshop co-located at ACM WomENcourage 2019	Manifesto for Today 2019	Rome, Italy	08-Jul-19	https://conferences.computer.org/iciot/2019/
Claudio Ardagna	UMIL	International Symposium on Secure Virtual Infrastructures Cloud and Trusted Computing 2019	C&TC 2019	Rhodes, Greece	21-Oct-19	http://www.otmconferences.org/index.php/conferences/ctc-2019
Claudio Ardagna	UMIL	11th IEEE International Conference on Cloud Computing	IEEE CLOUD 2019	Milan, Italy	08-Jul-19	https://conferences.computer.org/cloud/2019/
Cristian Hesselman	SIDN	ACM Internet Measurement Conference 2019	IMC 2019	Amsterdam, Netherlands	21-Oct-19	https://conferences.sigcomm.org/imc/2019/
Elena Ferrari	UI	11th IEEE International Conference on Web Services	IEEE ICWS 2019	Milan, Italy	08-Jul-19	https://conferences.computer.org/icws/2019/

Continued on next page…

Name	Partner	Scientific Event	Abbr.	Where	When	URL
Elena Ferrari	UI	IEEE International Conference on Trust, Privacy and Security in Intelligent Systems, and Applications	IEEE TPS 2019	Los Angeles, CA, US	12-Dec-19	http://www.sis.pitt.edu/lersais/tps/2019/
Elena Ferrari	UI	36th IEEE International Conference on Data Engineering	ICDE 2020	Dallas, Texas, USA	20-Apr-20	https://www.utdallas.edu/icde/
Elena Ferrari	UI	35th IEEE International Conference on Data Engineering	ICDE 2019	Macau, SAR, China	08-Apr-19	http://conferences.cis.umac.mo/icde2019/
Ernesto Damiani	UMIL	IEEE 2019 IEEE World Congress on Services	IEEE Services 2019	Milan, Italy	08-Jul-19	https://conferences.computer.org/services/2019
Ernesto Damiani	UMIL	2019 IEEE International Congress on Big Data	IEEE BigData 2019	Milan, Italy	08-Jul-19	https://conferences.computer.org/bigdatacongress/2019
Ernesto Damiani	UMIL	15th International Conference on Signal Image Technology & Internet based Systems	SITIS 2019	Sorrento, Italy	26-Nov-19	http://www.sitis-conf.org
Ernesto Damiani	UMIL	1st Model-Driven Simulation and Training Environments for Cybersecurity	MSTEC 2019	Luxembourg	26-Sep-19	https://www.threat-arrest.eu/html/mstec/
Ernesto Damiani	UMIL	OnTheMove Federated Conferences & Workshops	OTM 2019	Rhodes, Greece	21-Oct-19	http://www.otmconferences.org/index.php
Ernesto Damiani	UMIL	14th International Conference on Evaluation of Novel Approaches to Software Engineering	ENASE 2019	Heraklion, Greece	04-May-19	http://www.enase.org/Home.aspx?y=2019
Lili Nemec Zlatolas	UM	European Conference on Advances in Databases and Information Systems	ADBIS 2019	Bled, Slovenia	08-Sep-19	https://adbis2019.um.si/

Continued on next page...

Name	Partner	Scientific Event	Abbr.	Where	When	URL
Lili Nemec Zlatolas	UM	International Conference on ICT Systems Security and Privacy Protection	IFIP SEC 2020	Maribor, Slovenia	26-May-20	https://sec2020.um.si/
Olivier Festor	UL	IEEE Conference on Network Softwarization	NetSoft 2019	Paris, France	24-Jun-19	https://netsoft2019.ieee-netsoft.org
Rémi Badonnel	UL	IFIP/IEEE/In Assoc. with ACM SIGCOMM International Conference on Network and Service Management 2019	CNSM 2019	Halifax, Canada	21-Oct-19	http://www.cnsm-conf.org/2019/committees.html
Rémi Badonnel	UL	IFIP/IEEE International Symposium on Integrated Network Management	IM 2019	Washington DC, USA	08-Apr-19	https://im2019.ieee-im.org/
Tatjana Welzer	UM	European Conference on Advances in Databases and Information Systems	ADBIS 2019	Bled, Slovenia	08-Sep-19	https://adbis2019.um.si/
Tatjana Welzer	UM	International Conference on ICT Systems Security and Privacy Protection	IFIP SEC 2020	Maribor, Slovenia	26-May-20	https://sec2020.um.si/

C Technical Program Committee Membership

The table on the next pages shows all conferences for which CONCORDIA partners are or have been member of the Technical Program Committee (TPC). The table is organized in alphabetic order of the first name.

Name	Partner	Scientific Event	Abbr.	Where	When	URL
Anna Sperotto	UT	Passive and Active Measurement (PAM) conference	PAM 2019	Puerto Varas, Chile	27-Mar-19	http://pam2019.niclabs.cl/
Anna Sperotto	UT	15th International Conference on Network and Service Management	CNSM 2019	Halifax, Canada	21-Oct-19	http://www.cnsm-conf.org/2019/
Anna Sperotto	UT	IEEE/IFIP Network Operations and Management Symposium	NOMS 2020	Budapest, Hungary	20-Apr-20	https://noms2020.ieee-noms.org/
Anna Sperotto	UT	4th International Workshop on Traffic Measurements for Cybersecurity	WTMC 2019	San Francisco, California	23-May-19	http://wtmc.info/index.html
Antonio Ken Iannillo	SnT	IEEE 30th International Symposium on Software Reliability Engineering	ISSRE 2019	Berlin, Germany	28-Oct-19	http://2019.issre.net/
Antonio Ken Iannillo	SnT	4th IEEE International Workshop on Reliability and Security Data Analysis	RSDA 2019	Berlin, Germany	28-Oct-19	http://dessert.ddns.net/RSDA2019/
Barbara Carminati	UI	ACM Conference on Data and Application Security and Privacy 2019	CODASPY 2019	Dallas, Texas, USA	25-Mar-19	http://www.codaspy.org/2019/index.html
Barbara Carminati	UI	COMPSAC 2019: Data Driven Intelligence for a Smarter World	COMPSAC 2019	Milwaukee, Wisconsin, USA	15-Jul-19	https://ieeecompsac.computer.org/2019/
Barbara Carminati	UI	IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining	ASONAM 2019	Calgary, Canada	20-Aug-19	http://asonam.cpsc.ucalgary.ca/
Claudio Ardagna	UMIL	IEEE 2019 International Conference on Cognitive Computing	IEEE ICCC 2019	Milan, Italy	08-Jul-19	https://conferences.computer.org/iccc/2019/

Continued on next page...

Name	Partner	Scientific Event	Abbr.	Where	When	URL
Claudio Ardagna	UMIL	15th IEEE International Conference on Services Computing	IEEE SCC 2019	Milan, Italy	08-Jul-19	https://conferences.computer.org/scc/2019/
Claudio Ardagna	UMIL	18th IEEE International Conference on Trust, Security and Privacy in Computing and Communications	TrustCom 2019	New Zealand	05-Aug-19	https://crow.org.nz/Trustcom2019
Claudio Ardagna	UMIL	2019 IEEE Global Communications Conference: Communication & Information System Security	GLOBECOM 2019	Waikoloa, HI, USA	09-Dec-19	https://globecom2019.ieee-globecom.org/
Claudio Ardagna	UMIL	11th International Symposium on Cyberspace Safety and Security	CSS 2019	Guangzhou, China	01-Dec-19	http://nsclab.org/css2019/
Claudio Ardagna	UMIL	2nd IFIP International Conference on Machine Learning for Networking	MLN 2019	Paris, France	03-Dec-19	http://www.adda-association.org/mln-2019/
Claudio Ardagna	UMIL	11th IEEE International Conference on Cloud Computing Technology and Science	CloudCom 2019	Sydney, Australia	11-Dec-19	http://2019.cloudcom.org/
Claudio Ardagna	UMIL	3rd International Conference on Security with Intelligent Computing and Big-data Services	SICBS 2019	New Taipei City, Taiwan	04-Dec-19	http://www.sicbs2019.info
Claudio Ardagna	UMIL	13th International Conference on Information Security Theory and Practice	WISTP 2019	Paris, France	11-Dec-19	http://www.wistp.org/
Daniel Tovarňák	MUNI	IEEE/IFIP Network Operations and Management Symposium	NOMS 2020	Budapest, Hungary	20-Apr-20	https://noms2020.ieee-noms.org/

Continued on next page...

Name	Partner	Scientific Event	Abbr.	Where	When	URL
Elena Ferrari	UI	ACM Symposium on Access Control Model Technologies	SACMAT 2019	Toronto Canada	04-Jun-19	http://www.sacmat.org/2019/index.php
Elena Ferrari	UI	5th IEEE International Conference on Collaboration and Internet Computing	CIC 2019	Los Angeles, CA, US	12-Dec-19	http://www.sis.pitt.edu/lersais/cic/2019/
Elena Ferrari	UI	10th ACM Conference on Data and Application Security and Privacy	CODASPY 2020	New Orleans, US	16-Mar-20	http://www.codaspy.org/2020/
Elena Ferrari	UI	9th ACM Conference on Data and Application Security and Privacy 2019	CODASPY 2019	Dallas, Texas, USA	25-Mar-19	http://www.codaspy.org/2019/index.html
Elena Ferrari	UI	28th ACM International Conference on Information and Knowledge Management)	CIKM 2019	Bejing, China	03-Nov-19	http://www.cikm2019.net/
Elena Ferrari	UI	IEEE Conference on Secure and Dependable Computing	DSC 2019	HangZhou, China	18-Nov-19	https://conference.cs.cityu.edu.hk/dsc2019/
Elena Ferrari	UI	The Second International Conference on Machine Learning for Cyber Security	ML4CS 2019	Xian, China	19-Sep-19	https://www.springer.com/gp/book/9783030306182
Elena Ferrari	UI	IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining	ASONAM 2019	Calgary, Canada	27-Aug-19	http://asonam.cpsc.ucalgary.ca/2019/
Elena Ferrari	UI	15th IEEE International Conference on Services Computing	IEEE SCC 2019	Milan, Italy	08-Jul-19	https://conferences.computer.org/scc/2019/
Elena Ferrari	UI	10th IEEE International Conference on Information Reuse and Integration for Data Science	IRI 2019	Los Angeles, CA, US	31-Jul-19	http://www.sis.pitt.edu/lersais/iri/2019/

Continued on next page...

Name	Partner	Scientific Event	Abbr.	Where	When	URL
Elena Ferrari	UI	IEEE International Congress on Internet of Things	ICIOT 2019	Milan, Italy	08-Jul-19	https://conferences.computer.org/iciot/2019/
Elena Ferrari	UI	11th IEEE International Congress on Cloud Computing	IEEE CLOUD 2019	Milan, Italy	08-Jul-19	https://conferences.computer.org/cloud/2019/
Elena Ferrari	UI	International Conference on Extending Database Technology	EDBT 2019	Lisbon Portugal	26-Mar-19	http://edbticdt2019.inesc-id.pt/
Elena Ferrari	UI	39th IEEE International Conference on Distributed Computing Systems	ICDCS 2019	Dallas, Texas, US	07-Jul-19	https://theory.utdallas.edu/ICDCS2019/
Ernesto Damiani	UMIL	11th IEEE International Conference on Cloud Computing	IEEE CLOUD 2019	Milan, Italy	08-Jul-19	https://conferences.computer.org/cloud/2019/
Ernesto Damiani	UMIL	20th International Conference on Intelligent Data Engineering and Automated Learning	IDEAL 2019	Manchester, UK	14-Nov-19	http://www.confercare.manchester.ac.uk/events/ideal2019/
Ernesto Damiani	UMIL	11th IEEE International Workshop on Security Aspects for Process and Services Engineering	SAPSE 2019	Milwaukee, Wisconsin, USA	15-Jul-19	https://ieeecompsac.computer.org/2019/sapse/
Ernesto Damiani	UMIL	27th ACM Conference On User Modeling, Adaptation And Personalization	UMAP 2019	Larnaca, Cyprus	09-Jun-19	http://www.cyprusconferences.org/umap2019/index.html
Ernesto Damiani	UMIL	KES Smart Digital Futures 2019 Multiconferences	Smart Digital Futures 2019	Malta	17-Jun-19	http://sdf-19.kesinternational.org/

Continued on next page...

Name	Partner	Scientific Event	Abbr.	Where	When	URL
Jürgen Schönwälder	JUB	IFIP/IEEE International Symposium on Integrated Network Management	IM 2019	Washington DC, USA	08-Apr-19	https://im2019.ieee-im.org/
Jürgen Schönwälder	JUB	IFIP/IEEE/In Assoc. with ACM SIGCOMM International Conference on Network and Service Management 2019	CNSM 2019	Halifax, Canada	15-Oct-19	http://www.cnsm-conf.org/2019/
Jürgen Schönwälder	JUB	IEEE/IFIP Network Operations and Management Symposium	NOMS 2020	Budapest, Hungary	20-Apr-20	https://noms2020.ieee-noms.org/
Martin Drašar	MUNI	IEEE/IFIP Network Operations and Management Symposium	NOMS 2020	Budapest, Hungary	20-Apr-20	https://noms2020.ieee-noms.org/
Muhamed Turkanović	UM	Special Session on Blockchain Technology and its applications at the 27th International Conference on Information, Communication and Automation Technologies	ICAT (BCTA) 2019	Sarajevo, Bosnia and Herzegovina	20-Oct-19	http://icat.etf.unsa.ba/icat-2019/cms/round-table/
Muhamed Turkanović	UM	Blockchain Forum 2019 Co-located with the 17th Int. Conference on Business Process Management (BPM)	BPM 2019	Vienna, Austria	03-Sep-19	https://bpm2019.ai.wu.ac.at/ http://www.wikicfp.com/cfp/servlet/event.showcfp?eventid=86050
Muhamed Turkanović and Tatjana Welzer	UM (6)	24th conference on Advanced information technologies and services	OTS 2019	Maribor, Slovenia	18-Jun-19	https://www.ots.si/language/en/
Pavel Čeleda	MUNI	IEEE/IFIP Network Operations and Management Symposium	NOMS 2020	Budapest, Hungary	20-Apr-20	https://noms2020.ieee-noms.org/

Continued on next page...

Name	Partner	Scientific Event	Abbr.	Where	When	URL
Radu State	SnT	IEEE International Conference on Blockchain and Cryptocurrency	ICBC 2019	Seoul, South Korea	14-May-19	https://icbc2019.ieee-icbc.org/
Radu State	SnT	IFIP/IEEE International Symposium on Integrated Network Management	IM 2019	Washington DC, USA	08-Apr-19	https://im2019.ieee-im.org/
Michael Sirivianos	CUT	13th International AAAI Conference on Web and Social Media	ICWSM 2019	Munich, Germany	11-Jun-19	https://www.icwsm.org/2019/
Michael Sirivianos	CUT	14th International AAAI Conference on Web and Social Media	ICWSM 2020	Atlanta, Georgia, USA	08-Jun-20	https://www.icwsm.org/2020/
Michael Sirivianos	CUT	ACM SIGMETRICS 2020	SIGMETRICS 2020	Boston, Massachusetts, USA	08-Jun-20	https://www.sigmetrics.org/sigmetrics2020/

D Editors of Journals

The table on the next pages shows all journals for which CONCORDIA members act as editors. Editor roles can be: Editor in Chief, Series Editor, Associate Editor, Area Editor, Guest Editor, Editorial Board Member or Editorial Advisory Board Member.

Name	Partner	Role	Description	Publisher	URL
Anna Sperotto	UT	Area Editor	SIGCOMM Computer Communication Review (CCR)	ACM	https://ccronline.sigcomm.org/editorial-board/
Elena Ferrari	UI	Associate Editor	Transactions on Data Science	ACM	https://tds.acm.org/editorial.cfm
Claudio Ardagna	UMIL	Associate Editor	Access	IEEE	https://ieeeaccess.ieee.org/
Jürgen Schönwälder	JUB	Series Editor	Communications Magazine - Network and Service Management series	IEEE	https://www.comsoc.org/publications/magazines/ieee-communications-magazine/editorial-board
Elena Ferrari	UI	Associate Editor in chief	Computing	IEEE	https://www.computer.org/csdl/magazine/ic/about/15624?title=Editorial%20Board&periodical=IEEE%20Internet%20Computing
Jürgen Schönwälder	JUB	Associate Editor	Transactions on Network and Service Management	IEEE	https://www.comsoc.org/publications/journals/ieee-tnsm/ieee-transactions-network-and-service-management-editorial-board
Elena Ferrari	UI	Associate Editor	Transactions on Service Computing	IEEE	https://www.computer.org/csdl/journal/sc/misc/14407?title=About&periodical=IEEE%20Transactions%20on%20Services%20Computing
Elena Ferrari	UI	Associate Editor	Data science and engineering (DSEJ)	Springer	https://www.springer.com/journal/41019
Anna Sperotto	UT	Associate Editor	Journal of Network and Systems Management	Springer	https://www.springer.com/journal/10922/editors
Jürgen Schönwälder	JUB	Associate Editor	Journal of Network and Service Management	Springer	https://www.springer.com/journal/10922/editors

Continued on next page. . .

Name	Partner	Role	Description	Publisher	URL
Rémi Badonnel	UL	Associate Editor	Journal on Network and Systems Management	Springer	https://www.springer.com/journal/10922/editors
Olivier Festor	UL	Editorial Advisory Board	Journal on Network and Systems Management	Springer	https://www.springer.com/journal/10922/editors
Aiko Pras	UT	Editorial Advisory Board	Journal on Network and Systems Management	Springer	https://www.springer.com/journal/10922/editors
Anna Sperotto	UT	Associate Editor	International Journal of Network Management	Wiley	https://onlinelibrary.wiley.com/page/journal/10991190/homepage/editorialboard.html
Burkhard Stiller	UZH	Associate Editor	International Journal of Network Management	Wiley	https://onlinelibrary.wiley.com/page/journal/10991190/homepage/editorialboard.html
Jürgen Schönwälder	JUB	Associate Editor	International Journal of Network Management	Wiley	https://onlinelibrary.wiley.com/page/journal/10991190/homepage/editorialboard.html
Gabi Dreo Rodosek	CODE	Associate Editor	International Journal of Network Management	Wiley	https://onlinelibrary.wiley.com/page/journal/10991190/homepage/editorialboard.html
Rémi Badonnel	UL	Associate Editor	International Journal on Network Management	Wiley	https://onlinelibrary.wiley.com/page/journal/10991190/homepage/editorialboard.html
Claudio Ardagna	UMIL	Editorial board member	Mobile Information Systems	Hindawi	https://www.hindawi.com/journals/misy/
Claudio Ardagna	UMIL	Review Editor	Cybersecurity and Privacy of Frontiers in Big Data	Frontiers Media SA	https://www.frontiersin.org/journals/big-data/sections/cybersecurity-and-privacy#

Continued on next page. . .

Name	Partner	Role	Description	Publisher	URL
Elena Ferrari	UI	Specialty Chief Editor	Cybersecurity and Privacy of Frontiers in Big Data	Frontiers Media SA	https://www.frontiersin.org/journals/big-data/sections/cybersecurity-and-privacy
Elena Ferrari	UI	Associate Editor	International Journal of Co-operative Information Systems	World Scientific	https://www.worldscientific.com/page/ijcis/editorial-board