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 D1.3: 3rd Year Report on Designing and Developing an European Secure, Resilient and Trusted Ecosystem (ESRTE)

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

Contractual Date of Delivery	M36
Actual Date of Delivery	2021-12-31
Deliverable Dissemination Level	Public
Editors	Vassilis Prevelakis (TUBS) Mattijs Jonker (UT) Jean-Yves Marion (UL) Aiko Pras (UT) Jürgen Schönwälder (JUB) Nikos Salamanos (CUT) Michael Sirivianos (CUT)
Contributors	All partners involved in WP1
Quality Assurance	Jan Kohlrausch (DFN-CERT) Nicolas Kourtellis (Telefonica I+D) Nicola Bena (University of Milan) Emil Lupu (Imperial College London)

[†]This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 830927.

The CONCORDIA Consortium

UniBW/CODE	University Bundeswehr Munich / 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
JACOBSUNI	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
~~TUDA~~	~~Technical University of Darmstadt~~	~~Germany~~
MU	Masaryk University	Czech Republic
BGU	Ben-Gurion University	Israel
OsloMET	Oslo Metropolitan University	Norway
Imperial	Imperial College London	UK
UMIL	University of Milan	Italy
BADW-LRZ	Leibniz Supercomputing Centre	Germany
EIT DIGITAL	EIT DIGITAL	Belgium
TELENOR ASA	Telenor ASA	Norway
AirbusCS-GE	Airbus Cybersecurity GmbH	Germany
SECUNET	secunet Security Networks AG	Germany
IFAG	Infineon Technologies AG	Germany
SIDN	Stichting Internet Domeinregistratie Nederland	Netherlands
SURF	SURF BV	Netherlands
CYBER-DETECT	Cyber-Detect	France
TID	Telefonica I+D SA	Spain
RUAG	RUAG AG (as a replacement for RUAG Schweiz AG)	Switzerland
BITDEFENDER	Bitdefender SRL	Romania
ATOS	Atos Spain S.A.	Spain
SAG	Siemens AG	Germany
Flowmon	Flowmon Networks AS	Czech Republic
TÜV TRUST IT	TUV TRUST IT GmbH	Germany
TI	Telecom Italia SPA	Italy
Efacec	EFACEC Electric Mobility SA (as a replacement for EFACEC Energia)	Portugal
ARTHUR'S LEGAL	Arthur's Legal B.V.	Netherlands
eesy-inno	eesy-innovation GmbH	Germany
DFN-CERT	DFN-CERT Services GmbH	Germany
CAIXABANK SA	CaixaBank SA	Spain
~~BMW Group~~	~~Bayerische Motoren Werke AG~~	~~Germany~~
NCSA	Ministry of Digital Governance - National Cyber Security Authority	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	ATHINA-ISI	Greece
UNI PASSAU	University of Passau	Germany
RUB	Ruhr University Bochum	Germany
CRF	Centro Ricerche Fiat	Italy
ELTE	EOTVOS LORAND TUDOMANYEGYETEM	Hungary
Utimaco	Utimaco managment GmbH	Germany
FER	University of Zagreb, Faculty of Electrical Engineering and Computing	Croatia

Document Revisions & Quality Assurance

Internal Reviewers

1. Jan Kohlrausch (DFN-CERT)
2. Nicolas Kourtellis (Telefonica I+D)
3. Nicola Bena (University of Milan)
4. Emil Lupu (Imperial College London)

Revisions

Ver.	Date	By	Overview
0.1	2021-10-07	Vassilis Prevelakis (TUBS)	Initial template
0.2	2021-10-25	Mattijs Jonker (UT)	UT input
0.3	2021-10-26	Despoina Antonakaki (FORTH)	FORTH input
0.4	2021-10-28	Yair Meidan (BGU)	BGU input
0.5	2021-11-01	Vassilis Prevelakis (TUBS)	TUBS input
0.6	2021-11-03	Jürgen Schönwälder (JUB)	JUB input
0.7	2021-11-04	Jean-Yves Marion (UL)	UL input
0.8	2021-11-08	Nicola Bena (UNIMI)	UNIMI input
0.8	2021-11-08	Salman Manzoor (ULANC)	ULANC input
0.8	2021-11-08	Eder Scheid (UZH)	UZH input
0.9	2021-11-11	Nicolas Kourtellis (TID)	TID input
0.9	2021-11-14	Nikos Salamanos (CUT)	CUT input
0.9	2021-11-15	Vassilis Prevelakis (TUBS)	TUBS input
1.0	2021-11-22	Vassilis Prevelakis (TUBS)	Final input and release for review
2.0	2022-06-22	Jürgen Schönwälder (JUB)	Ready for publication

Executive Summary

All effort within CONCORDIA tries to build 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. As stated in the Description of Action, the objectives for CONCORDIA's research in Work Package 1 (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.

In the previous years (as detailed in deliverables D1.1 and D1.2), CONCORDIA researchers have managed to publish more than 170 workshop, conference and journal papers. In this third year, we are proud to report that research has been carried forward even further, an achievement in itself, especially if we take into account the continued effects of the COVID-19 pandemic. More than 75 additional papers were published; see Appendix A for details. In fact, the total number of papers entered into the EU-ECAS system is even higher, since white papers and papers without peer-review are not counted in this deliverable. Although the objective regarding the number of papers, as stated in the Description of Action, has already been reached, research never ceases, so in the remaining years CONCORDIA members will continue to publish their work in top venues (USENIX Security, ACM-IMC, NDSS and IEEE INFOCOM, etc.). In 2021, CONCORDIA researchers contributed to the organization of 13 scientific conferences (see Appendix B) and acted as Technical Program Committee member for more than 45 different conferences (see Appendix C). Since the beginning, CONCORDIA has also established a liaison with three other EU pilot projects (ECHO, SPARTA and Cybersec4Europe), that led to the organization of the *4th International NeCS PhD Winter School (Trento)*, the *IFIP Summer School on Privacy and Identity Management* (virtual conference), the *Early Stage PhD workshop* and the *CODE research day*. Continuing to take novel initiative, the project organized this year as well, the *CONCORDIA Open Door Event* in October which was held virtually because of restrictions concerning COVID-19. From the beginning of the project, women have been playing a pivotal role, as is evident from their participation in the organization of these events

Furthermore, CONCORDIA researchers played a leading role in the organization of the scientific community by acting as chairs of the IFIP Technical Committee 6 (Communication Systems) and the IFIP Working Group 6.6 (Management of Network and Distributed Systems), member of the Scientific Council of ANSSI and president of the Scientific Council of CNRS. Additionally, they were involved as editorial board members in 15 journals and transactions including ACM, IEEE, Springer and Wiley (see Appendix D).

Beyond the traditional academic venues, continued effort in 2021 led CONCORDIA researchers to contribute to the "societal impact" of the work presented in this deliverable. Concrete examples are the research papers by T1.5 that study the impact of Twitter and YouTube on the US elections, aggression propagation and diffusion, change in behavior around COVID-19 lockdowns and detection of disinformation cascades.

The fruitful collaboration among WP1 tasks and the project pilots has continued during this third year, and resulted in joint work which materialized in more than 10 publications (see Appendix A). A concrete example is the strong collaboration that has been established between application-centric security research (T1.4) and the Threat Intelligence for the Telco Sector pilot (T2.1), where T1.4's efforts in investigating application security have been instrumental in making sure that requirements set by T2.1 are met.

Last year, in Deliverable D1.2, we introduced task-specific "considerations" for a security roadmap, in order to give a sense of direction for future research, aiming to strengthen Europe's digital sovereignty. In this third version of the Deliverable, we have included a similar dedicated section as well, which has been introduced as a consolidated section. This also applies to the section about the impact of COVID-19 on research activities.

In general, WP1 has excelled in its scientific output concerning publications, continuing to produce high-quality research that has long exceeded the objectives as described in the Description of Action. Collaboration between research in WP1 and the pilots in WP2 and WP3 was strengthened as well. Despite COVID-19 still having significant effects in everyday life and resulting in a number of activities being canceled, CONCORDIA partners managed to disseminate the project in various ways and represent it in several venues, by participating virtually.

Contents

1 Introduction

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

- Perform excellent academic research to build an European Secured, Resilient and Trusted Networked Ecosystem, papers for scientific journals, conferences and workshops
- Organize scientific events in the area of cybersecurity, including a dedicated annual European cybersecurity conference
- Take a leading role in the organization of the scientific community, outreach to different target audiences, including public media and the general public
- Contribute to standardization, open research data and code, shared via systems such as 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 depicts the research organization within WP1. It is divided into five tasks, each focusing on a particular aspect of cybersecurity (bottom to top):

- 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

1.1 Structure of the document

This deliverable is structured as follows. Section 2 outlines key achievements of the third year. Section 3, discusses the actions taken after the interim review, earlier in 2021, and details how each of the review comments has been addressed.

The research activities of the various tasks are then discussed in detail in the following sections. In some cases, where the papers are discussing similar topics, subjects, or issues, and in order to avoid repetition, activities are merged/grouped together into a common section. An overview of research related to device-centric security is provided in Section 4. Section 5 outlines network-centric security research activities. Section 6 discusses research concerning software and system-centric security aspects. Data and application specific security aspects are the fo-

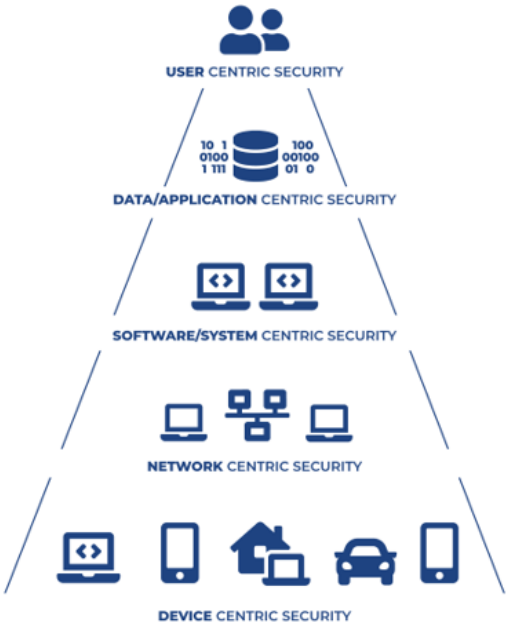

Figure 1: Research organization within WP1

cus of Section 7. Finally, Section 8 covers efforts related to user security and privacy. Each Section starts by shortly summarizing the corresponding task's goals, followed by an overview of its main topics of research. Then, the relations and collaboration with the project pilots are enumerated, followed by a detailed list of all the research work undertaken in the third year of the project.

The deliverable concludes with a summary of the organization of scientific events and the scientific community, in Section 9. Section 10 discusses the contributions to standards and open research data/tools. Conclusions from the activities described are drawn in Section 11. The appendices detail publications, conferences, technical program committee (TPC) memberships as well as editorial boards.

1.2 Security Roadmap Considerations

In last year's Deliverable D1.2, we included a section for each task concerning considerations for a security roadmap. The objective is to identify the research challenges that can act as enablers for the European industry and academia to build the most secure products in the world, offer excellent quality services, carry out cutting-edge research, and support Computer Emergency Response Teams (CERTs), security teams/vendors, and Secure Operations Centers (SOCs) in identifying and defending against attacks, among others. Consequently, since little has

changed in terms of directions to which future research should move towards since last year, we refer the reader to D1.2 [1] and most importantly to the preliminary version of deliverable D4.4 (Cybersecurity Roadmap for Europe) [2], which is dedicated to this subject.

However, in order to align research activities in the current deliverable with the identified targets detailed in Section 4 of D4.4 (Roadmap for Research and Innovation) and to show how this research addresses/puts forward challenges as introduced in the roadmap, we list here the connections between the two:

- Device-Centric Layer
  - Section 4.3.1, 4.3.2 correspond to Section 4.1.5 “Embedded operating systems utilizing hardware security features” of D4.4
  - Sections 4.3.3, 4.3.4, 4.3.5 correspond to Section 4.1.8 “Postquantum cryptography schemes on constrained devices” of D4.4.
  - Sections 4.3.6, 4.3.7, 4.3.8, 4.3.9, 4.3.10, 4.3.11, 4.3.12, 4.3.13, 4.3.14, 4.3.15, 4.3.16, 4.3.17, 4.3.18 correspond to Section 4.1.4 “Device identification and assessment mechanisms” of D4.4.
- Network-Centric Layer
  - Section 5.3.13, correspond to Section 4.2.1 “Open Networking: The Responsible Internet” of D4.4.
  - Sections 5.3.1, 5.3.2 and 5.3.3 correspond to Section 4.2.2 “Trustworthy DNS Resolver Infrastructures” of D4.4.
  - Sections 5.3.4, 5.3.5, 5.3.6 and 5.3.7 correspond to Section 4.2.3 “DDoS Protection Services” of D4.4.
  - Sections 5.3.8, 5.3.9, 5.3.10, 5.3.11 and 5.3.12 correspond to Section 4.2.4 “Monitoring and Data Collection Infrastructure (Data Lakes)” of D4.4.
  - Sections 5.3.14 and 5.3.15 correspond to Section 4.2.5 “Network Assurance & Certification” of D4.4.
- Software/System-Centric Layer
  - Sections 6.3.1 and 6.3.2 are related to Section 4.4.1 on HPC and EU-controlled Cloud Infrastructure of D4.4.
  - Sections 6.3.3, 6.3.4 6.3.5 and also 6.3.6 correspond to Section 4.3.2 on ”Malware detection and analysis” of D4.4. Section 6.3.5 is also related to Section 4.1.6 on ”Microkernel isolation and virtualization mechanisms” of D4.4.

  - Section 6.3.7 corresponds to Section 4.4.2 “Smart Technologies” of D4.4.

- Data/Application-Centric Layer
  - Sections 7.3.12, 7.3.14 correspond to Section 4.4.1 “EU-controlled Cloud Infrastructure” of D4.4.
  - Sections 7.3.5, 7.3.6, 7.3.7, 7.3.8, 7.3.9 correspond to Section 4.4.2 “Smart Technologies” of D4.4.
  - Sections 7.3.1, 7.3.2, 7.3.3, 7.3.4, 7.3.13 correspond to Section 4.4.3 “Securing data/software in distributed computing environments” of D4.4.
  - Sections 7.3.10, 7.3.11, 7.3.15, 7.3.16 correspond to Section 4.4.4 “Inter-networking in the future” of D4.4.
- User-Centric Layer
  - Sections 8.3.17, 8.3.18, 8.3.21, 8.3.22, 8.3.24 correspond to Section 4.5.1 “Fighting disinformation in Europe” of D4.4.
  - Sections 8.3.1, 8.3.2, 8.3.3, 8.3.4, 8.3.11, 8.3.13 correspond to Section 4.5.2 “Data Ownership and Data Privacy” of D4.4.
  - Sections 8.3.8, 8.3.15, 8.3.16 correspond to Section 4.5.3 “Dynamic Attribute-Based Trusted Digital Identity Management” of D4.4.

1.3 Impact of COVID-19 on Research Activities

Originally, we had a number of plans for 2021 with respect to activities concerning the research domain:

- Organize physical meeting(s) of the Consortium
- Organize visit(s) to other CONCORDIA partners
- Summer and Winter Schools, together with other pilots
- Organize event(s) for PhD researchers

However, the outbreak of the COVID-19 pandemic put a stop to our plans. The restrictions applied on gatherings and international travel, both by local governments and by organizations alike, made physical meetings difficult – if not impossible – to organize for quite a while.

The impact of the COVID-19 pandemic turned out to be much higher than could be anticipated from the initial stages of the pandemic, to both PhD and senior researchers. Most PhD researchers are well into their third year by now and COVID-19 has had a negative impact on their mental well-being and productivity. In order

to avoid feeling isolated, some of them decided to move temporarily back home to their families, in many cases across different time zones, which limited the possibility of organizing virtual events, making new contacts, or developing new ideas or activities, making cross-institutional collaboration and event organization more difficult.

Senior researchers have also had limited opportunities to socialize and exchange ideas at informal coffee breaks or social events at physical conferences, since most of these events moved to online presentations. Additionally, there was limited enthusiasm to organize new physical conferences or events due to low attendance, even after the introduction of vaccinations.

Furthermore, COVID-19 has also had serious impact in terms of significantly increasing the burden of work for all involved in the project. Most organizations and institutions have had to design and implement new procedures, new modes of teaching, and new modes of interaction to cope with the effects of the pandemic. This has significantly increased the workload for everyone.

However, despite the lower number of papers and organized events compared to the start of the project, WP1 has still met and even exceeded its objectives.

2 Key Achievements

As mentioned in Section 1, the Description of Action defines four objectives for WP1.

Excellent Academic Research

The main objective of WP1 is to stimulate the publication of scientific results in the broad field of cybersecurity (see Section 1). Deliverables D1.1 and D1.2 in the previous years, report that more than 170 papers have already been published in workshops, conferences and journals during years one and two of the CONCORDIA project. Although the SMART objective has therefore already been met, CONCORDIA partners still write high quality papers and will continue to do so in the project's future. During the third year, more than 75 additional papers were published (see Appendix A for the details).

journals (SJR)			conferences (CORE)		
> 1	$(1, 0.9]$	$(0.9, 0.8]$	A*	A	B
8	4	2	4	4	18

Table 2: Journal ranking (SJR) and conference rankings (CORE)

Table 2 provides an overview of the rankings of the journals and conferences in which publications have been published in 2021. Eight papers appeared in journals with a SCImago Journal Rank (SJR) greater than 1. Out of these, two papers were published in journals with an SJR above 5. These are journals that are ranked at position 15 (IEEE Communications Magazine) and 16 (ACM Computing Surveys) of all Computer Science journals. Four papers appeared in conferences ranked A* (USENIX Security, Web Conference, IEEE Conference on Data Engineering, ACM Conference on Computer and Communications Security). Overall, 40 papers were published in journals and conferences with a strong to good scientific standing. The remaining papers often appeared in smaller dedicated workshops or conferences that are not ranked by CORE or they were published as book chapters that are typically not ranked.

Organization of Scientific Events

In year one, as reported in deliverable D1.1, CONCORDIA researchers were members of the technical program committees (TPC) of 40 different conferences and contributed to the organization of more than 25 scientific events. In year two, as reported in deliverable D1.2, they were members of the TPC of over 70 different conferences and contributed to the organization of almost 15 scientific conferences. During the third year of the project, researchers were members of the TPCs of over

47 different conferences and they contributed to the organization of 13 scientific conferences (additional details are provided in Section 9).

As is evident, there has been a slight decrease in the organization of and participation to scientific events compared to the previous years. This is to be expected, as the COVID-19 pandemic has had serious impact on these events, which led to complete cancellation in the worst case.

Organization of the Scientific Community

Professional organizations play an important role in organizing the scientific community. As in the two previous years of the project, various CONCORDIA researchers held positions in such organizations during the third year, as well. A CONCORDIA researcher chairs the IFIP Technical Committee 6 (TC6). TC6 focuses on *Communication Systems* and is the largest TC within IFIP. The chair of Working Group 6.6 of TC6 is also a CONCORDIA member. WG 6.6 focuses on the management of network and distributed systems. Another CONCORDIA researcher serves as the President of the Scientific Council of CNRS, as well as a member of the Scientific Council of ANSSI.

A number of CONCORDIA researchers have also served on journal editorial boards during year three and are members of steering committees. A total of 15 journals by publishers such as ACM, Springer and IEEE are involved. Further details are given in Section 9.

Contributions to Standardization and Open Research Data

The fourth objective of WP1 is to contribute to standardization efforts, open research data, and tools/software. Notably, during the third year of CONCORDIA, the OpenINTEL measurement project continued to grow and gather additional information about the state of the Domain Name System (DNS). More details can be found in Section 10.

Noteworthy here is that CONCORDIA has two tasks (T5.3 and T6.4) specifically related to standardization and open data. Consequently, various CONCORDIA efforts related to these activities are not reported on in this deliverable. A more in-depth view can be found in the deliverables of WP5 (D5.2, D5.3, D5.4) and WP6 (D6.4, D6.5, D6.6).

3 Recommendations from the review

This Section discusses the recommendations provided by the reviewers in the third (February 2021) review report, as far as these comments relate to WP1.

Feedback: *There is a lack of coordination between WP1 and T2.3, i.e. taking into account the results of research papers in WP1, such as the one published in Cordis "Safer wireless charging for connected electric vehicles".*

Answer: In 2021, a number of activities were undertaken to strengthen collaboration between WP1 and the new T2.3 leaders from CRF. At the beginning of the year, we organized several meetings to explore synergies between the WP1 tasks and T2.3. Unfortunately, the COVID crisis had more impact on T2.3 than envisaged, causing some delay. In the second half of 2021, CRF started collaboration with Politecnico di Torino, which joined the CONCORDIA project in October 2021. For administrative reasons Politecnico di Torino joined T2.3 (instead of WP1), even though it is a research institute and as such contributes to the research output of CONCORDIA. Since their research has a strong focus on blockchain technology, the plan for 2022 is to strengthen collaboration between Politecnico di Torino and the WP1 activities on blockchain (most notably by the partners UZH, UM, UMIL, UI, SnT). In that way, the collaboration between WP1 and the T2.3 activity 'Monitoring and certification of inbound logistics for the EV battery pack' will be strengthened.

Feedback: *Increase the coordination between WP1 and the industrial pilots. For example, improve discussions with partners of WP1 to define the SoA for Cybersecurity in electric vehicles and how to go beyond the state of the art.*

Answer: In 2021, the collaboration between WP1 and the industrial pilot-specific tasks of WP2 has intensified and, in general, runs very well:

- T1.1 has strong collaboration with the Telecom, Finance, and Communication Sectors (T2.1, T2.2 and T2.5 respectively). For more details, see Section 4.2.
- T1.2 is closely collaborating with the Telecom and Finance Sectors (T2.1 and T2.2 respectively). For more details, see Section 5.2.
- T1.3 has improved coordination with the Telecom, e-Heath, and Communication Sectors (T2.1, T2.4 and T2.5 respectively). For more details, see Section 6.2.
- T1.4 has been increasingly coordinating with the Telecom and Finance Sectors (T2.1 and T2.2 respectively). For more details, see Section 7.2.
- T1.5 has an ongoing collaboration with the Telecom, Finance, e-Heath, and Commuication Sectors (T2.1, T2.2, T2.4 and T2.5 respectively). For more details, see Section 8.2.

Regarding specifically the state of the art for Cybersecurity in electric vehicles, as mentioned in the previous paragraph, real collaboration between T2.3 (Transport E-Mobility Sector) and WP1 is still in progress.

To conclude, throughout the project there has been strong collaboration among the industrial and academic partners. Each of the WP1 tasks has contributed to at least 2, usually 3 or more pilots. Figure 2 shows the overview of the relationships among the several tasks.

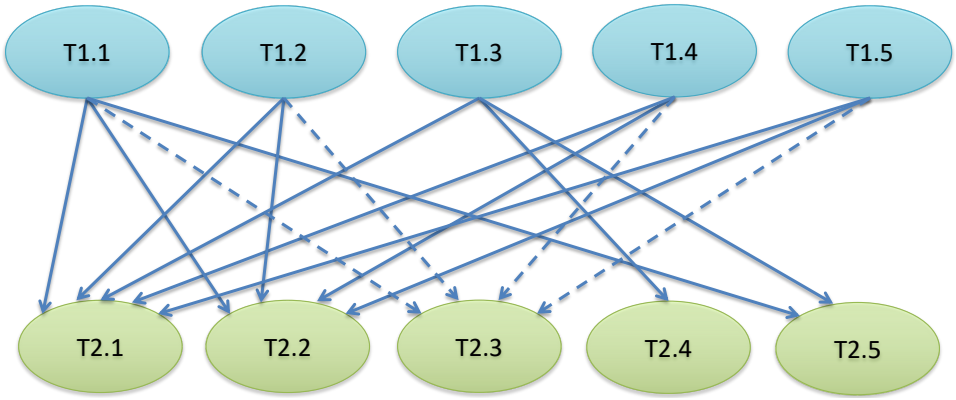

Figure 2: Links among WP1 tasks and the CONCORDIA industrial pilots. Dashed lines indicate collaborations to be explored and strengthened further, in 2022.

4 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.

4.1 Overview

The research undertaken in 2021 can be grouped into five clusters hardware, cryptography, Internet of Things, data analytics, and blockchains, as described in the following.

Hardware

The first cluster deals with research that is focused on the hardware level. The work described in Section 4.3.1 discusses how hardware-based control-flow integrity mechanisms can be implemented. An FPGA implementation of a SPARC system on a chip is used for evaluation purposes. Section 4.3.2 describes research that aims at finding mechanisms that isolates low-level unsafe library code that may be used as part of memory and type-safe runtime systems.

Cryptography

Cryptography plays a major role in securing devices and hence research related to cryptographic primitives forms the second cluster. Cryptographic primitives have to be correct, efficient, and they need to be safe against side-channel attacks. The work summarized in Section 4.3.3 investigates how homomorphic encryption libraries can be tested efficiently. A case study using the PALISADE homomorphic encryption library revealed a number of bugs that have been reported to the PALISADE project. Section 4.3.4 describes research investigating hardware countermeasures to secure systems against side-channel attacks utilizing machine learning techniques, which is, particular, relevant for cryptographic algorithms. The research outlined in Section 4.3.5 concerns the acceleration of post-quantum cryptography operations using GPUs and FPGAs and number theoretic transforms.

Internet of Things

The third cluster consists of research efforts related to devices that are data sources or actuators in the Internet of Things. The work described in Section 4.3.6 focuses

on the trustworthiness of data in IoT systems. An argument is made that the view should focus more towards trusting data instead of necessarily trusting devices. The work summarized in Section 4.3.7 develops architectures for blockchain and IoT integration with a specific focus on constraints imposed by wireless technologies such as LoRaWAN and IEEE 802.15.4 with very small frame sizes and very low data rates. Closely related to this effort is the work described in Section 4.3.8, which aims at implementing the IPFIX protocol on IoT devices that communicate over IEEE 802.4 link layers.

MQTT is a protocol used to exchange data between IoT devices and their cloud-based servers. The research described in Section 4.3.9 concerns an access control framework for MQTT interactions. Performance improvements have been realized that reduce the overhead during policy retrieval. The work described in Section 4.3.10 focuses on device authentication aspects. The goal is to provide a secure comprehensive authentication framework for users, as well as for IoT devices for verifying their identity to remote services.

Data Analytics

The fourth cluster integrates research applying data analysis techniques to discover or analyze the context a device is operating in. The work summarized in Section 4.3.11 aims at detecting anomalies and cyber attacks on industrial control systems using statistical and machine learning techniques. The research described in Section 4.3.12 uses neural networks to discover domain names generated to contact command and control servers of botnets. The system has been deployed and data and code have been made publicly available. A data stream temporal clustering algorithm called DeepStream is described in Section 4.3.13. It detects sequential and overlapping clusters in IoT datasets. Section 4.3.14 investigates distributed learning techniques for anomaly detectors in 5G IoT scenarios. A novel peer-to-peer algorithm has been defined, which avoids a central orchestrator.

Blockchains

The fifth cluster concerns research using and improving blockchain technology in the context of device security. Section 4.3.15 investigates how IoT devices can sign blockchain transactions so that even if an edge node is compromised, the private key of a device is not leaked.

Data stored on blockchains is by design immutable. However, there are often requirements to support the deletion of data or to update data on blockchains. Section 4.3.16 describes ways to support such data updates.

Section 4.3.17 describes an access management system utilizing artificial intelligence and blockchain technology in an IoT context. The system uses face recognition technology executed on an IoT device. It leaves a log of access decisions in an immutable blockchain. Finally, the investigation summarized in Section 4.3.18

concerns the usage of blockchains to improve the fairness and robustness of bug-bounty programs.

4.2 Link between T1.1 work and CONCORDIA pilots

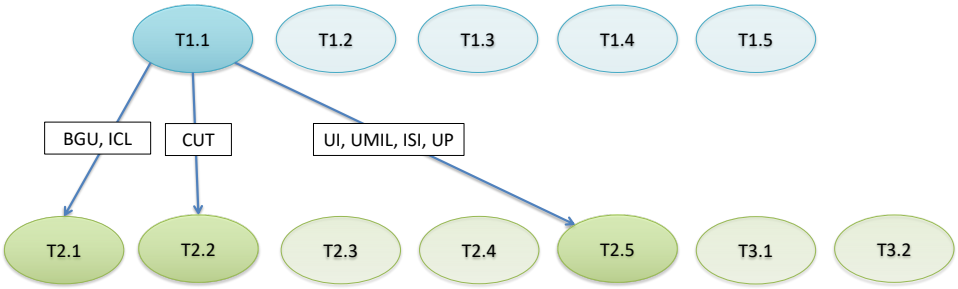

Figure 3: Relation between Task 1.1 and the CONCORDIA pilots

Figure 3 shows how the efforts undertaken in T1.1 relate to the pilots of the CONCORDIA project.

1. BGU has been collaborating with Ericsson in the context of the Telcommunications Pilot (T2.1). Ericsson's research focused on cellular IoT botnet detection by utilizing AI/ML-models for anomaly detection. A fundamental assumption of this research is that equipment will be monitoring the traffic and reporting flows in a format based on the IPFIX protocol. To enable experiments, BGU researchers shared with Ericsson's research team a large set of network traffic flow traces and associated meta-data collected from a variety of real-world high-end IoT devices deployed at BGU. The set includes both benign and malicious flows, such that Ericsson's botnet detection approach can be trained, tested, and compared with existing approaches.

   In addition, ICL has been cooperating with RISE in the Telcommunications Pilot on a peer-to-peer distributed learning algorithm for anomaly detection in 5G IoT scenarios.

2. CUT has been collaborating with CAIXA in the context of the Finance Sector Pilot (Task 2.2). The joint work is centered around the idea of attribute-based device-centric authentication, further described in Section 4.3.10.

3. ACS, UI, and UMIL collaborated in the context of the Communication Pilot (T2.5), where the security of unoccupied aerial systems must be protected. The work initially focused on identifying reference scenarios and requirements of interest, which resulted in a paper published in 2020. The pandemic has delayed the collaboration in 2021 but the partners hope to resume the joint work in 2022. Initial steps towards this direction have been done in late 2021 to identify mechanisms for dynamic trust estimation in an aeronautical context.

4. The work by ISI and UP described in Section 4.3.5 is linked to the efforts of Communication Pilot (T2.5) as a possible part of the secure authentication use case.

4.3 Summaries of T1.1 Research Activities

The following sections briefly summarize the research activities carried out during 2021.

4.3.1 Hardware-based Control-Flow Integrity for Embedded Devices

Contact: George Christou (FORTH), Sotiris Ioannidis (FORTH)

The diversification of computing systems and the wide adoption of IoT devices that pervade our lives have grown the security and safety concerns in home appliances, enterprise infrastructure and control systems. Typical examples range from traditional IoT environments where data are collected and processed in back-end cloud systems, to more sophisticated, edge-based scenarios where part of processing also occurs in end-devices. Protecting against such cases using software-only solutions is not sufficient, since advanced attacks can modify even the security software itself, thus bypassing any restrictions posed. In addition, the performance overheads of software-based solutions is non-negligible in certain cases. The use of hardware-backed solutions can vitally improve the security of embedded devices, even though this is still challenging due to their limited resources and their intrinsic budget of performance and memory.

At the same time, the exploitation techniques are constantly evolving. More than a decade ago, exploiting software was as easy as just simply smashing the stack. An attacker could simply inject code into a vulnerable buffer in the stack and overwrite the return address (of the current stack frame) to point back to their code. Today, this is not possible due to data execution prevention (DEP) mechanisms, however attackers can still exploit software in other ways e.g. code-reuse attacks. A way of stopping such exploits is to prevent the execution of any new functionality, by employing Control-Flow Integrity (CFI) techniques. An attacker cannot inject code or introduce any new functionality that is not part of the legitimate control-flow graph (CFG). Unfortunately, the majority of existing CFI proposals have still many open issues (related to accuracy and performance), that hinder their applicability.

We explore the implementation of a full-featured CFI enabled Instruction Set Architecture (ISA) on actual hardware [3]. Our new instructions provide the finest possible granularity for both intra-function and inter-function Control-Flow Integrity. We implement hardware-based CFI (HCFI) by modifying a SPARC SoC and evaluate the prototype on an FPGA board by running SPECInt benchmarks instrumented with a fine-grained CFI policy. HCFI can effectively protect applications from code-reuse attacks, while adding less than 1% average runtime and

2% power consumption overhead, making it particularly suitable for embedded systems.

HCFI is a hardware design that offers a CFI solution that is (i) complete, since it protects both forward and backward edges, (ii) fast, since the experienced overhead is, on average, less than 1%, and (iii) more accurate, since it employs a full-functional shadow stack implemented inside the processor core. Furthermore, we argue that HCFI is the most complete hardware implementation of CFI so far, supporting many problematic cases (such as setjmp/longjmp, recursion, fall-through functions and indirect jumps within functions).

4.3.2 Blue pill JS[1]

Contact: George Christou (FORTH), Sotiris Ioannidis (FORTH)

Modern applications written in high-level memory-managed programming languages enjoy the security benefits of memory and type safety. Unfortunately, even a single low-level memory- and type-unsafe library can wreak havoc on the rest of an otherwise safe application, by bypassing all the safety and security guarantees offered by the semantics of the high-level language and implemented by the language runtime. We propose a new hybrid permission model aimed at protecting a library binary core. This is work we started recently and which is still in the initial stages.

4.3.3 Applying Metamorphic Testing to Homomorphic Cryptography

Contact: Melvin Wolf (JUB), Jürgen Schönwälder (JUB)

In software testing, a test oracle is a mechanism for determining whether a test has passed or failed. One compares the output of the system under test, for a given test-case input, to the output that the oracle determines to be correct. The oracle problem refers to the problem of determining the correct output for a given input (and a set of program or system states). It comes up when it is impossible (or prohibitively expensive) to know whether a returned mapping from the input to the output space is correct. Metamorphic testing alleviates this issue by instead verifying that multiple different mappings from the input space to the output space are consistent within themselves. It checks whether a system behaves according to a certain set of properties called metamorphic relations which are relationships between multiple input and output pairs. Our work investigates whether metamorphic testing can be applied to homomorphic encryption, where operations like addition and multiplication can be carried out on encrypted data without having to reveal the cleartext values.

The PALISADE homomorphic encryption software library implements two somewhat homomorphic encryption (SWHE) schemes over integers, Brakerski-Fan-

[1] Javascript

Vercauteren (BFV) (including two variants, BFVrns and BFVrnsB) and Brakerski-Gentry-Vaikuntanathan (BGV). It also implements a SWHE scheme for real numbers, the Cheon-Kim-Kim-Song (CKKS) scheme. Both additive and multiplicative homomorphisms are implemented. In order to analyze homomorphic encryption schemes, we defined five metamorphic relations covering the encrypt-decrypt test, identity operations and testing for overflows since PALISADE uses its own implementation of mathematical operations. 50 test cases were run for every metamorphic relation on each of the homomorphic encryption schemes. The results revealed multiple faults which have since been confirmed by the PALISADE team to be bugs in their implementation. The details of this work have been published in [4].

4.3.4 Machine Learning Attacks and Countermeasures on Hardware Accelerators

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

Continuing on the research activities defined in the previous year, the University of Patras (UP) and the Industrial Systems Institute (ISI)/RC ATHENA have been collaborating during 2021 in assessing hardware countermeasures against machine learning attacks in a side channel analysis context. Machine Learning techniques have proven effective in Side Channel Analysis (SCA), enabling multiple improvements over the already-established profiling process of Template Attacks. Focusing on the need to mitigate their impact on embedded devices, a design model and strategy is proposed [5] that can effectively be used as a backbone for introducing SCA countermeasures on Elliptic Curve Cryptography (ECC) scalar multipliers. The proposed design strategy is based on the decomposition of the round calculations of the Montgomery Power Ladder (MPL) algorithm and the Scalar Multiplication (SM) algorithm into the underlined finite field operations, and their restructuring into parallel-processed operation sets.

The parallelization of the various field operations required during an MPL round is achieved by utilizing 3 $GF(2^k)$ Adders and 3 $GF(2^k)$ Multipliers. The operations are equally distributed through the modules in a timely and synchronized manner and with minimal idle time. Having as a basis this proposed parallel design strategy, we showcase how advanced SCA countermeasures can be easily introduced, focusing on randomizing the projective coordinates of the MPL round's ECC point results.

Samples of the aforementioned power traces for either bit-0 or bit-1 are presented in Figure 4. The inclusion of both an unprotected and a protected implementation is clearly depicted in the number of rounds visible in the power traces, where only two additional rounds are needed as overhead for the countermeasures of randomization and MPL parallelization. As a next step, in order to evaluate the design approach and its SCA countermeasures, several simple ML-based SCAs are per-

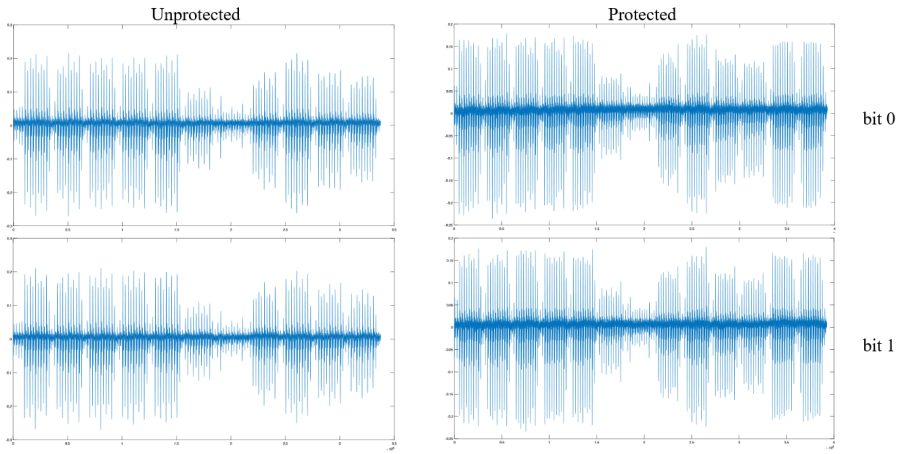

Figure 4: Trace samples of bit-0 or bit-1 for both implementations

formed, and an attack roadmap is provided. The proposed roadmap assumes attackers that do not have access to a huge number of leakage traces, and that have limited resources with which to mount Deep Learning attacks. The trained models' performance reveals a high level of resistance against ML-based SCAs when including SCA countermeasures in the proposed design strategy. This work complements the work done in [6] and reported in previous year's deliverable D1.2. It elaborates on the the importance of accessing hardware accelerators in security applications regarding their side channel attack resistance under the presence of side channel attack countermeasures and provides insight and details on performing such assessment with real measurements. Given that Europe has a long tradition on design and implementation of independent and secure SoCs as well as on leakage assessment and certification of such solutions, the above research work contributes in the aforementioned European efforts.

4.3.5 Optimizations on Number Theoretic Transform (NTT)

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

In this research domain, we have explored the Number Theoretic Transform (NTT) that, as the most computationally intensive building block of several post-quantum cryptography schemes, needs to be considerably optimized to achieve a high level of efficiency. The research works that have been done in the third year of the project related to this activity are focused on parallelizing NTT internal operations using the OpenCL programming language, so as to create optimal software designs for use in GPUs and FPGAs. Furthermore, they are aimed at restructuring the NTT algorithm C code implementation, so that it can optimally be realized for High Level Synthesis (HLS) tools that generate hardware NTT implementations.

The main performance bottlenecks of lattice-based cryptography is the polynomial multiplication taking place in the LBC lattices and the noise sampler function. Multiplication can be optimally performed using the Number Theoretic Transform (NTT): the polynomials to be multiplied are converted into the spectral domain, reducing the polynomial multiplication to a simple point-wise multiplication of the two polynomials. This way, they drastically convert the complexity from $O(n^2)$ to $O(n \cdot \log(n)))$.

As mentioned above, the first approach that has been followed is a study [7] on the feasibility of using OpenCL, a portable framework for parallel programming of hardware accelerators.

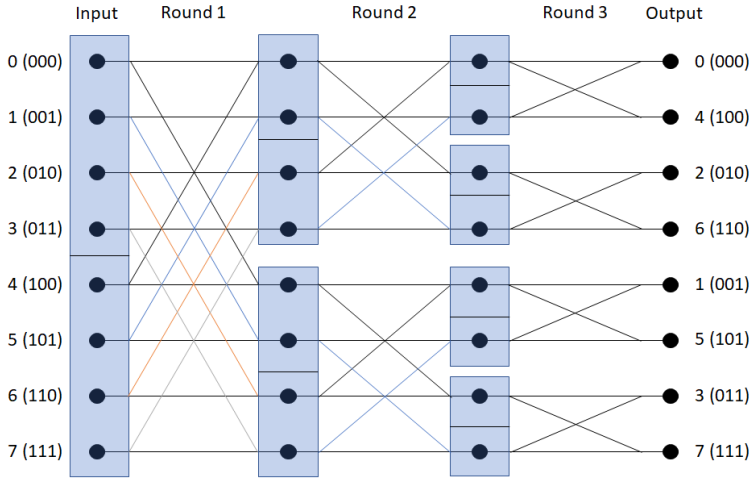

Figure 5: NTT computations without initial bit reversal

With its help, a parallelized implementation that has been realized can be seen in Figure 5. Both NTT and Inverse NTT algorithms using CT-FFT (Cooley-Tukey Fast Fourier Transform) three for-loops are used, that take an input of a polynomial in the form of a vector v with n values being the coefficients of the polynomial. The outer for-loop counts the rounds of the algorithm and the number of rounds is equal to $log_2(n)$. We measure the performance of our implementation on a GPU and evaluate when and where such a deployment is beneficial. Our results showed that the proposed parallel implementation is a viable acceleration approach for these algorithms for lattice-based cryptography solutions. Our results showed that there is a potential speed up against the CPU implementations, with a maximum of speed up around 7.5x times depending on the number of input coefficients. We found that for an average number of 1664 coefficients the GPU execution time is on par with the CPU and after that there is significant gain on using the GPU. As many algorithms for the post-quantum cryptography of the NIST Post-Quantum standardization competition have high degree polynomials and also a large number of them, our approach can potentially be applied to them, since they make use of

the NTT. The overall results can be seen in the following tables for NTT (Table 3) and Inverse NTT respectively (Table 4).

Table 3: NTT C/OpenCL speedup

Batch Size	Number of coefficients					
	64	128	256	512	1024	2048
1	x0.03	x0.08	x0.18	x0.33	x0.61	x1.26
2	x0.07	x0.15	x0.31	x0.63	x1.01	x2.16
4	x0.15	x0.3	x0.61	x1.2	x1.93	x4.04
8	x0.3	x0.59	x1.2	x2.33	x3.25	x5.51
16	x0.6	x1.16	x2.31	x4.08	x4.39	x6.93
32	x1.17	x2.23	x3.91	x4.93	x4.48	x7.54
64	x2.17	x3.62	x4.58	x5.73	x5.11	x6.75

Table 4: Inverse NTT C/OpenCL speedup

Batch Size	Number of coefficients					
	64	128	256	512	1024	2048
1	x0.04	x0.09	x0.19	x0.38	x0.49	x1
2	x0.08	x0.19	x0.4	x0.78	x0.95	x1.8
4	x0.17	x0.37	x0.78	x1.14	x1.85	x4.09
8	x0.34	x0.74	x1.54	x2.28	x3.44	x5.8
16	x0.69	x1.45	x2.9	x4.13	x4.69	x6.91
32	x1.35	x2.7	x4.99	x5.13	x4.88	x7.49
64	x2.42	x4.71	x6.13	x6.25	x5.5	x7.74

Further studying the NTT optimization problem, an NTT design approach for HLS (High Level Synthesis) tools is proposed in [8]. It manages to minimize the dependency between NTT processing elements while increasing the access of these elements to memory by forcing the HLS tool to use dual port RAMs. Initially, we tested the configuration capabilities of OpenCL NTT-code (using the highly constrained number of OpenCL attributes), which however did not manage to produce efficient results. We introduce in the paper a reformulation of C/C++ NTT code and then further optimizing this approach by adapting a FFT-based technique in the NTT domain. As a result of this effort, two architectures are presented in this paper, the initial and an optimized/proposed architecture. The latter architecture is able to increase the NTT processing elements access rate to memory and manages to minimize the dependency between memory read and memory write loops within the algorithm which is used to derive the proposed optimized architecture. We also propose the use of the DEPENDENCE pragma at appropriate points within the code, in order to effectively eliminate the dependency in the optimized architecture. By using as a case study the NTT version used in the PQC Dilithium Digital Signature, it can be shown that the proposed approaches lead to 20–50% latency

improvement using only two parallel butterfly units compared to other existing HLS-based NTT solutions in the relevant literature.

The above research approaches, as their results indicate, provide significant proposals on how to generate efficient implementations of the NTT operation used in post-quantum cryptography. Given that the need for post-quantum schemes is high, as the quantum computers become a reality, replacing traditional public key cryptography schemes with post-quantum ones must be achieved with minimum efficiency compromises. Thus, the proposal of techniques and mechanisms to provide this type of efficiency is of considerable value. This becomes more important considering that European research teams have significant presence in the postquantum cryptography standardization efforts (manifested in the NIST postquatnum cryptography standard competition) and providing efficient implementations for the European postquantum cryptography standard candidates (eg. Dilithium algorithm) can further benefit such candidates in the NIST contest.

4.3.6 Security Assurance for Emerging Systems

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

Current distributed systems increasingly rely on hybrid architectures built on top of IoT, edge, and cloud, backed by dynamically configurable networking technologies like 5G and content-centric networking. These networking technologies promise to be game-changers, and are already being applied in scenarios of critical safety and of large-scale content distribution, where security, safety and privacy requirements are crucial.

First, we tackled one of the main hurdles towards trustworthy IoT systems, that is, data trustworthiness. In [9], we considered the IoT scenario of billions of devices collecting data from the physical environment, which are pre-processed at the edge and then forwarded to processing services at the core of the infrastructure, on top of which cloud-based applications are built and provided to mobile end users. IoT comes with important advantages in terms of applications and added value for its users, making their world smarter and simpler. These advantages, however, are mitigated by the difficulty of guaranteeing IoT trustworthiness, which is still in its infancy. IoT trustworthiness is a must especially in critical domains (e.g., health, transportation) where humans become new components of an IoT system and their life is put at risk by system malfunctioning or breaches. We put forward the idea that trust in IoT can be boosted if and only if its automation and adaptation processes are based on trustworthy data. We departed from a scenario that considers the quality of a single decision as the main goal of an IoT system and consider the trustworthiness of collected data as a fundamental requirement at the basis of a trustworthy IoT environment. We therefore defined a methodology for data collection that filters untrusted data out according to trust rules evaluating the status

of the devices collecting data and the collected data themselves. Our approach is based on blockchain and smart contracts and collects data whose trustworthiness and integrity are proven over time. The methodology balances trustworthiness and privacy, and has been experimentally evaluated in real-world and simulated scenarios using Hyperledger fabric blockchain.

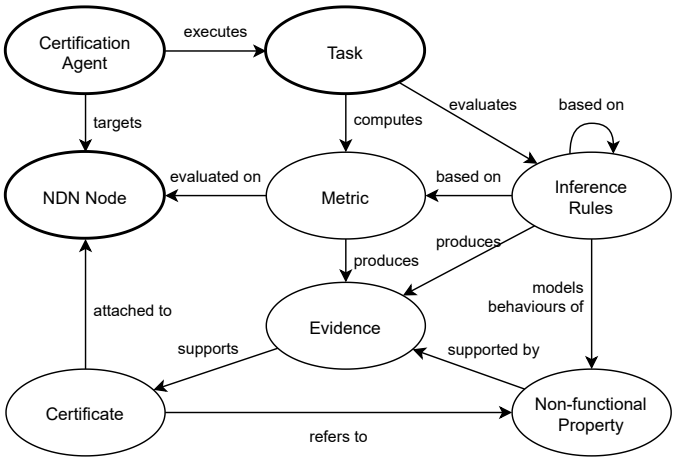

Figure 6: Certification methodology for content-centric networks

Second, we considered content-centric networks. They are based on the strong assumption of being able to access genuine content from genuine nodes, which is however unrealistic and could open the door to disruptive attacks. Network node misbehavior, either due to poisoning attacks or malfunctioning, can act as a persistent threat that goes unnoticed and causes dangerous consequences. In [10], we proposed a novel certification methodology for content-centric networks that improves transparency and increases trustworthiness of the network and its nodes. The proposed approach, shown in Figure 6, builds on behavioral analysis and implements a continuous certification process that collects evidence from the network nodes and verifies their non-functional properties using a rule-based inference model. We experimentally evaluated utility, performance, and soundness of our approach in a simulated Named Data Networking (NDN) network targeting properties availability, integrity, and non-repudiation.

Finally, in this complex environment, traditional security governance solutions cannot provide the holistic view that is needed to manage these systems in an effective and efficient way. In [11], we proposed a security assurance framework for edge and IoT systems based on an advanced architecture capable to deal with 5G-native applications. The framework is based on a cloud/edge architecture implemented on Kubernetes, as shown in Figure 7.

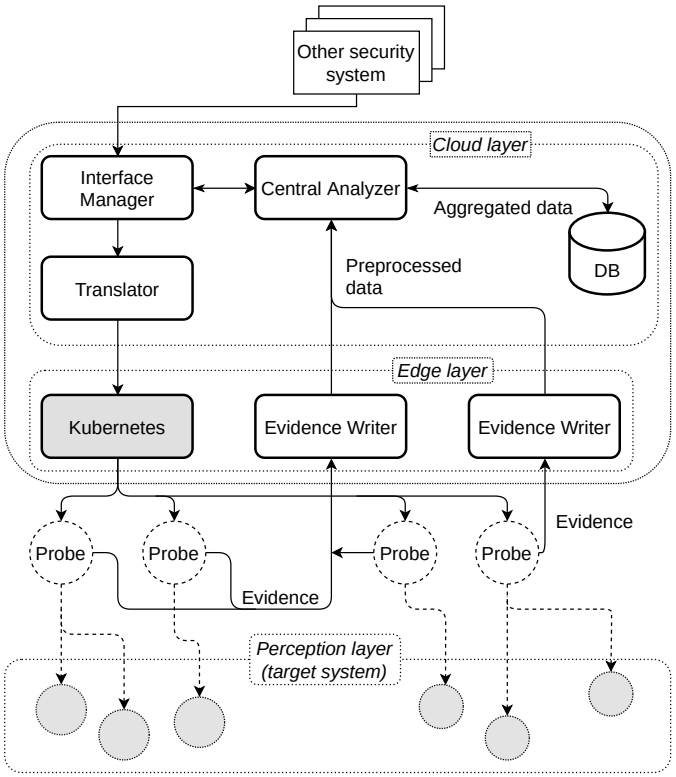

Figure 7: Architecture of the assurance framework for edge and IoT systems

4.3.7 Architectures for Blockchain-IoT Integration

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

Blockchains (BC) serve as a chain of transactions persisted as backward-linked lists, while being created and maintained within a network of distributed nodes. Potential advances with BCs have reached various application areas beyond FinTech-oriented use cases. Since Internet-of-Things (IoT) based use cases are an important part of them, this work [12] focuses specifically on defining and determining measures and criteria to be met proactively for an efficient BC and IoT integration, termed BIoT.

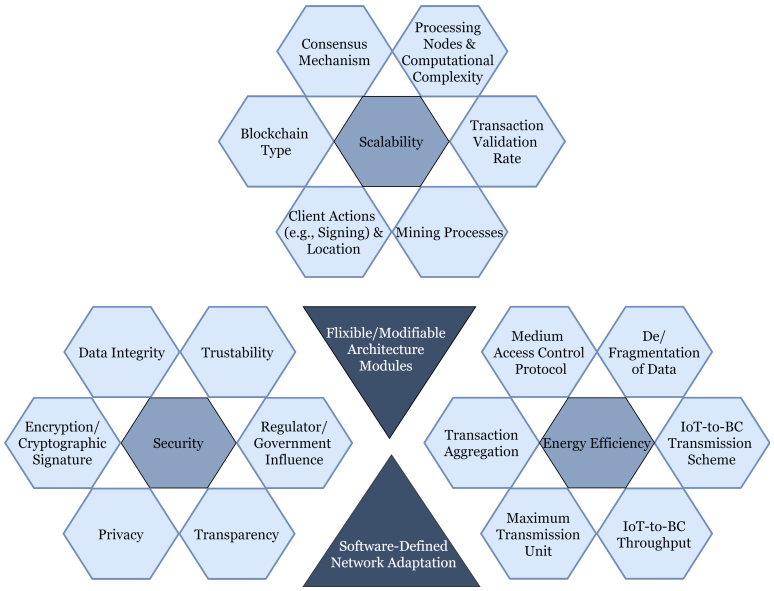

Figure 8: Blockchain-IoT Integration (BIoT) Efficiency Affecting Metrics

This work resembles BIoT potentials and incentives and collects practical BIoT challenges as shown in Figure 8. In the next step, driven by state-of-the-art BIoT architectures, the new architecture BIIT 1.0 (as shown in Figure 9) is proposed by [12] to pave the path toward practical and efficient BIoT architectures. A BIIT IoT-to-BC Transaction (TX) adaptation scheme based on Software Defined Network (SDN) management modules proposes on-the-fly adaptation of the transmission scheme for higher throughout.

BIIT is composed of a set of (a) IoT, BC, and networking components, (b) a set of management and configuration components, and (c) a set of instructions on what shall be considered before adapting and while using an IoT-BC communication. BIIT especially considers limits faced (or forced) by the underlying networking protocols, e.g., for a LoRaWAN network with a Maximum Transmission Unit (MTU) of 256 Bytes, and air time of 30 s per day, BIIT enhances the network

efficiency by employing a Listen Before Talk (LBT) media access mechanism instead of using the default configuration of LoRa, i.e., the Duty Cycle Enforcement (DCE), which performs weakly (like Aloha). Furthermore, BIIT enables Acknowledgments (Ack) via Automated Repeat reQuests (ARQ), and facilitates TX fragmentation approach such as in IEEE 802.15.4, to guarantee a high TX throughout, transmission reliability, and energy efficiency.

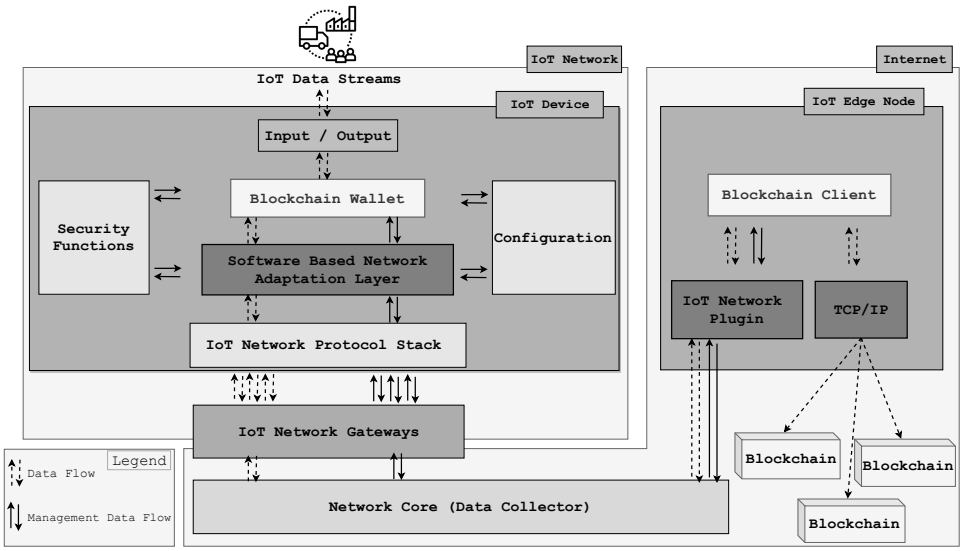

Figure 9: BIIT 1.0, The Blockchain-IoT Integration (BIoT) Architecture

4.3.8 Python-Based TinyIPFIX in Wireless Sensor Networks

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

While Wireless Sensor Networks (WSN) offer potential, their limited programmability and energy limitations determine operational challenges. Thus, a TinyIPFIX-based system was designed [13], such that this application layer protocol is now used to exchange data in WSNs efficiently. The system implementation in MicroPython is efficient and straightforward compared to a lower-level programming language while displaying valuable properties in terms of overhead and power efficiency. Furthermore, it demonstrates that MicroPython may pave the way towards Network Function Virtualization (NFV) on Internet-of-Things (IoT) devices by providing highly portable software functions implemented in a high-level programming language.

In this work, a TinyIPFIX platform was implemented using Espressif ESP-WROOM-32D devices. Two ESP-WROOM-32D devices were chosen as the hardware platform: Espressif ESP32 DevKitC V4 and Macchina SuperB. These devices were prepared for programming using MicroPython. Then each ESP device

was equipped with a Digi XBee board, which features the IEEE 802.15.4 standard allowing for low power communication among devices in the WSN.

TinyIPFIX maintains a more negligible data overhead than the regular Type-Length-Value (TLV) data transfer in selected application scenarios. Furthermore, it has been demonstrated experimentally that the Python-based TinyIPFIX works well in a home-based IEEE 802.15.4 network providing almost a 100% delivery ratio. The solution offers a uniform Python-based implementation spanning multiple elements of the system. It includes TinyIPFIX functional entities such as End Devices, Concentrators, and the Collector. Furthermore, it brings a ZMQ broker and applications on the Application Server. The Python-based environment provides much faster value creation than older systems depending on low-level programming languages. Moreover, the energy consumption of those devices running TinyIPFIX has been evaluated. The primary method of reducing the energy consumption in the WSN is to leave devices as long as possible in the deep sleep mode (*i.e.,* ESP32 and XBee devices). The high deep sleep current currently experienced on Micropython-programmed ESP32 devices (*e.g.,* 20 mW on the SuperB device) has to be decreased in the future, allowing for an extensive life-span of the network.

4.3.9 MQTT-based Privacy Preferences Enforcement

Contact: Elena Ferrari (UI), Pietro Colombo (UI)

During 2021, the ABAC framework for MQTT environments originally proposed in [14] has been extended to regulate IoT device communication in both ordinary and emergency situations. An emergency typically requires granting exceptional privileges to subjects, which in an ordinary situation would not be permitted. Therefore, we have proposed an approach to regulate device communication on the basis of the possible emergency situations where users administering IoT devices connected to an MQTT environment are involved. Our framework allows one to i) model the events that trigger an emergency, ii) bind events to MQTT messages, iii) specify emergency situations along with their possible evolution, and iv) specify ordinary and emergency policies. In addition, our framework allows i) the detection of occurrences of modeled events starting from the analysis of MQTT control packets exchanged in a monitored application, and ii) the identification of the possible evolution of emergency situations, and iii) the retrieval and enforcement of the applicable ordinary and emergency policies. The feasibility of the proposed framework has been studied by employing it in a case study of pseudo realistic complexity related to a MQTT based health monitoring application to be used in a nursing home during COVID-19 pandemic. Our experimental performance evaluations have shown a reasonably low enforcement overhead. The proposed approach, the case study and the experimental results have been presented in a paper which is currently under review.

In addition, since we observed that in our ABAC framework for MQTT environments [14] a large part of access control time overhead is due to policy retrieval, we have studied an approach to speed up ABAC policies selection. The proposed mechanism relies on different index structures which are jointly used to short the retrieval of ABAC policies that apply to an access request. Initial experimental results show a significant optimization of the selection performances in scenarios where sets of ABAC policies with different characteristics have been considered.

4.3.10 Attribute-based Device-centric Authentication

Contact: Kostantinos Papadamou (CUT), Nikos Salamanos (CUT), Michael Sirivianos (CUT)

In the context of task T1.1, CUT is working on a password-less authentication solution that combines state-of-art technologies and protocols for replacing the traditional password paradigm with a preserving-privacy attribute-based device-centric authentication solution. CUT is aiming to provide a secure comprehensive identity verification and authentication framework for users and IoT devices. The novelty in our solution is that users will have to prove their identity once and then use our solution as many times as they wish to: 1) prove their identity and that they are who they claim to be to any interested Service Provider (i.e., banks); and 2) authenticate with any Service Provider using their mobile devices leveraging the FIDO 2.0 and the OpenID Connect protocols. The value of our solution for the users is that they only have to create and maintain only one account in our solution and prove once their identity to third-party Service Providers using our identity verification solution. At the same time, Service Providers do not have to employ their own KYC infrastructure to verify the identity of their users neither they have to maintain any user accounts or sensitive personal information. In the context of the CAIXA Bank pilot, the developed solution will be integrated and used by CAIXA bank as their KYC solution to verify the identity of their customers.

In another context, the developed solution can also be used by IoT devices to verify their identity to remote services. More precisely, CUT is working on leveraging the developed solution to enable IoT devices to securely verify their footprint (identity) to a given server (IoT Controller) providing assurances to the IoT controller that it is not communicating with a malicious device. In addition, users who have access rights to specific IoT devices can use the developed solution to authenticate with them securely using strong authenticators (such as their fingerprint). In general, with the developed solution we will be able to eliminate attacks on IoT devices and infrastructures like large Denial-of-service (DoS) attacks, as well as other types of attacks since IoT devices will have to verify their footprint/identity to communicate with a given server. Last, CUT is also currently working on how to handle failed authentication attempts differentiating errors from malicious attempts, as well as on a secure fallback authentication framework for users who have lost their device.

4.3.11 Poisoning Attacks on Cyber Attack Detectors for Industrial Control Systems

Contact: Asaf Shabtai (BGU)

Industrial control systems (ICSs) combine distributed computing with physical process monitoring and control. Many ICSs are safety-critical, and an attack interfering with their functionality can cause substantial financial and environmental harm, and endanger people. The importance of ICSs makes them an attractive target for cyber attacks. In this research [15], BGU researchers propose a method for detecting anomalies and cyber attacks in physical-level ICS data using 1D CNNs, shallow undercomplete autoencoders (UAEs), variational autoencoders (VAEs), and PCA. This method (outlined in Figure 10) improves upon the methods presented in previous researches, allowing arbitrary length sequence prediction and an arbitrary prediction horizon, adding a max-based method for threshold detection, and formalizing the detection hyperparameter criteria. In addition, BGU researchers propose a feature selection approach using the Kolmogorov-Smirnov test and transform time domain signals into frequency representation using short-time Fourier transform and energy binning, and model the system in both the time and frequency domains. The method was evaluated on three popular public datasets representing both real-world and simulated data (SWaT, BATADAL, WADI) and achieved better detection performance than previously published research in this area. In addition, this evaluation demonstrates the effectiveness of the proposed feature selection method and its generalizability. Finally, BGU researchers evaluated the robustness of the proposed method to adversarial evasion attacks under a threat model of a white-box attacker that has gained control of the sensor data. The results show the method's resilience: to evade detection, the attacker must abandon his/her goal of physically impacting the system.

4.3.12 DGA Domain Embeddings for Tracking and Exploring Botnets

Contact: Asaf Shabtai (BGU)

Botnets (both PC and IoT) have been using domain generation algorithms (DGA) for over a decade to covertly and robustly identify the domain name of their command and control servers (C&C). Recent advancements in DGA detection have motivated botnet owners to rapidly alter the C&C domain and use adversarial techniques to evade detection. As a result, it has become increasingly difficult to track botnets in DNS traffic. In this research [16, 17], BGU researchers present Helix, a method for tracking and exploring botnets. Helix uses a spatio-temporal deep neural network autoencoder to convert domain names into numerical vectors (embeddings) which capture the DGA and the seed used to create the domain. This is made possible by leveraging both convolutional (spatial) and recurrent (temporal) layers, and by using techniques such as attention mechanisms and highways. Furthermore, by using an autoencoder architecture, the network can be trained in

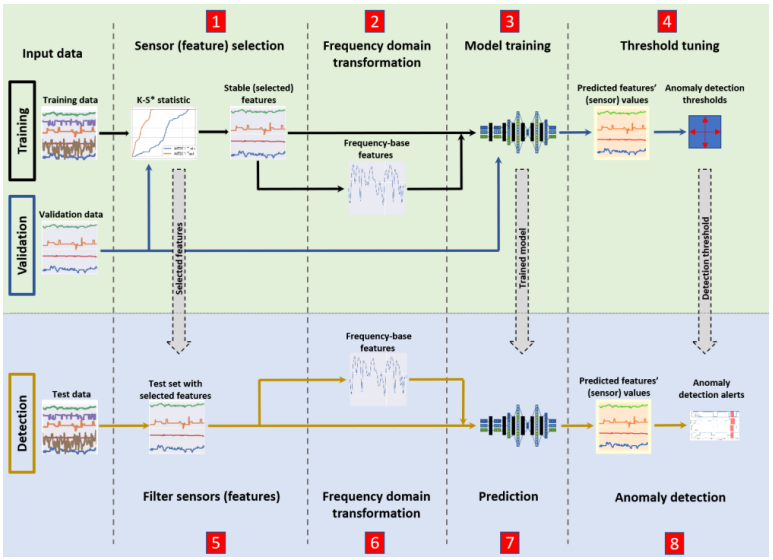

Figure 10: Overview of the proposed method for cyber attack detection in industrial control systems

an unsupervised manner (no labeling of data) which makes the system practical for real world deployments. In their evaluation, BGU researchers found that Helix can track botnet campaigns, distinguish between DGA families and seeds, and can identify domains generated using the latest adversarial machine learning techniques. Helix is currently being used to monitor botnets in one of the world's largest ISPs, and the related code and DGA datasets are now publicly available online for reproducibility.

4.3.13 Autoencoder-based Stream Temporal Clustering and Anomaly Detection

Contact: Asaf Shabtai (BGU), Yair Meidan (BGU)

The increasing number of IoT devices in smart environments, such as homes, offices, and cities, produces seemingly endless data streams and drive many daily decisions. Consequently, there is growing interest in identifying contextual information from sensor data to facilitate the performance of various tasks, e.g., traffic management, cyber attack detection, and healthcare monitoring. The correct identification of contexts in data streams is helpful for many tasks, for example, it can assist in providing high-quality recommendations to end users and in reporting anomalous behavior based on the detection of unusual contexts. In this research [18, 19] BGU researchers present DeepStream, a novel data stream temporal clustering algorithm that dynamically detects sequential and overlapping clusters. DeepStream is tuned to classify contextual information in real time and is capable of coping with a high-dimensional feature space. DeepStream utilizes stacked au-

toencoders to reduce the dimensionality of unbounded data streams and for cluster representation. This method detects contextual behavior and captures nonlinear relations of the input data, giving it an advantage over existing methods that rely on PCA. DeepStream was evaluated empirically using four sensors and IoT datasets and was compared to five state-of-the-art stream clustering algorithms. The evaluation shows that DeepStream outperforms all of these algorithms. The evaluation also demonstrates how DeepStream's improved clustering performance results in improved detection of anomalous data.

4.3.14 Non-IID Data Re-Balancing at IoT Edge with Peer-to-Peer Federated Learning for Anomaly Detection

Contact: Luis Muñoz-González (ICL) Muhammad Zaid Hameed (ICL)

In 5G IoT networks, ML-based anomaly detectors enable the learning and prediction of normal and anomalous patterns. Due to the distributed nature of the deployment, distributed learning techniques allow the training of an anomaly detection model without collecting the data from all the sensors. In contrast to traditional federated learning approaches, which rely on an orchestrator to coordinate the training process, the distribution of the training can be done in a peer-to-peer fashion. However, two problems still subsist: firstly, the data that each sensor has is different and we are thus in a non-IID[2] scenario and secondly, each sensor sees very few anomalies by comparison with the size of the benign data. Addressing these problems is therefore fundamental in order to apply distributed anomaly detection approaches in IoT deployments. We have collaborated with RISE in the context of the Telecom Pilot to develop a solution to these challenges. The result has been a novel peer-to-peer algorithm, P2PK-SMOTE, which adaptively applies a data re-balancing approach where peers that participate in the FL share complex synthetic points that are artificial generated by linear interpolation from several nearest neighbours. The work was published in [20] and further details are available in deliverable D2.3:Telecom Pilot.

4.3.15 IoT Device to Blockchain Interoperability Transaction

Contact: Eder J. Scheid (UZH), Muriel Franco (UZH), Christian Killer (UZH), Bruno Rodrigues (UZH), Burkhard Stiller (UZH)

One crucial aspect in the integration and security of Blockchain (BC) and Internet-of-Things (IoT) is the management of private keys and signing BC transactions. Once a BC transaction is signed using the private key of the BC address, it cannot be altered without the possession of the corresponding private key. This scheme provides (i) non-repudiation, (ii) authentication, and (iii) integrity to the transaction, which originates from the device that holds the private key. However, if the

[2] IID stands for independent and identically distributed

device's private key is not secured or leaked while signing a BC transaction, a malicious actor can generate or modify transactions, rendering the IoT devices' measurements untrusted. Therefore, the approach Edge2BC [21] was proposed, which directly signs Elliptic Curve Digital Signature Algorithm (ECDSA)-based BC transactions in constrained Class 0 IoT devices without revealing the private key to external actors. The transaction signing was implemented in the IoT device's Hardware Security Module (HSM) and exposes an interface for interaction commands. Edge2BC relies on an edge node to *(a)* retrieve specific BC information, such as transaction fees, and *(b)* interact with a BC interoperability Application Programming Interface (API). Such an API sends the signed raw transactions to Ethereum, HyperLedger, or any Remote Protocol Enabled (RPC)-enabled BC. A Proof-of-Concept (PoC) of Edge2BC was implemented, and evaluations were performed, showing the feasibility of Edge2BC and that by relying on Edge2BC, even if the edge node is compromised, the private key of the device is not leaked, since the key never leaves the device. If the IoT device is stolen, the private key of that device can be included in a deny-list in the BC-enabled application that specifies that transactions signed with that private key are not trusted.

Furthermore, it was performed extensions on the design and development of the BC interoperability API (named Bifröst) employed by Edge2BC. In [22], it is presented the following Bifröst extensions. The first is the definition of an interaction standard for Notary-based BC interoperability APIs based on a JavaScript Object Notation (JSON) format, which can be used by different IoT devices that cannot sign a BC transaction directly in the device. The second is the design and implementation of an encryption scheme that abstracts technical details (*e.g.,* salt generation and encryption algorithm) from users or IoT devices to provide privacy to publicly stored data (*i.e.,* data stored in public BCs). Evaluations on the extended Bifröst prototype show an increase in the data size due to the employed encryption algorithm (not format-preserving), but a constant performance overhead. In addition, discussions on the standard proposed in [22] support the argument that it is flexible and generic; thus, it can be well-employed in different Notary-based BC interoperability solutions besides Bifröst. Hence, such extension of adding an encryption mechanism to Bifröst allow IoT devices to protect the privacy of data sent to the BC using a simple API endpoint on Bifröst without implementing resource-intensive encryption functions.

4.3.16 On-Chain IoT Data Modification in Blockchains

Contact: Sina Rafati Niya (UZH), Burkhard Stiller (UZH)

Blockchain IoT Integration (BIoT) applications face many challenges to comply with the European General Data Protection Regulation (GDPR) i.e., enabling users to hold on to their rights for deleting or modifying their data stored on publicly accessible and immutable BCs. In this regard, this work [23] identifies the requirements of BCs for being GDPR compliant in BIoT use cases. Accordingly,

an on-chain solution is proposed that allows fine-grained modification (update and erasure) operations on Transaction (TX) data fields within a BC. The proposed solution is based on a cryptographic primitive called Chameleon Hashing. The novelty of this approach is manifold. In this approach, users have the authority to update their data TX, which are addressed at the TX level and not in block level with no side-effects on chain. By performing and storing the data updates, all on-chain, traceability and verifiability of the BC are preserved. Moreover, the compatibility with TX aggregation mechanisms introduced in [24] that allow the compression of the BC size is maintained.

4.3.17 IoT-based Access Management Supported by AI and Blockchains

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

There is a growing number of research initiatives where Artificial Intelligence (AI), Blockchains (BC), and Internet-of-Things (IoT) are jointly used to solve problems in various application domains. The main contribution of this work [25] is a strong use case in the area of IoT-based access management, which encompasses AI, BC, and IoT, where a user needs to present her face in front of a camera installed on an IoT device to access a resource.

The second contribution is the hardware architecture. The system is composed of an IoT device, an IoT Gateway (GW), and infrastructure supporting a BC. The image capturing and face recognition are handled by a Tensilica Xtensa LX6-based Esp Eye IoT device, which contains a double-core architecture supporting the 240 MHz CPU frequency, equipped with a 2-megapixel OV2640 camera and an IEEE 802.11 network adapter. The IoT GW, equipped with an IEEE 802.11 network adapter as well, serves as the middle man, which waits for data (*i.e.*, images and metadata) coming from Esp Eye devices and inserts received data into the BC.

The third contribution is the software architecture. The system takes an image of the person requesting access, performs face detection and recognition, and checks whether this given user has the right to access this given resource. Face detection and recognition are performed directly on the IoT device to increase the responsiveness of the system. A Multi-Task Cascaded Convolutional Network (MT-CNN) with MobileNets (MN) were deployed to detect faces. Furthermore, face recognition is based upon a Convolutional Neural Network (FRMN). The images taken by the sensor and AI decisions on access rights are stored in immutable, tamper-resistant storage, allowing for BC-based backend communication or auditing and establishing a good level of transparency. The system's performance was evaluated at an excellent level, where a 5.3 s end-to-end delay is reached, from the moment the sensor takes the image until the transaction reaches all peers spanning a BC.

4.3.18 Blockchain-based Bug Bounty Framework

Contact: Asaf Shabtai (BGU)

Bug bounty programs are a popular solution for security researchers to disclose software vulnerabilities in exchange for compensation. They suffer, however, from two main drawbacks that limit their effectiveness: (i) they use a trusted intermediary that charges hefty commission fees and may have a conflict of interest with the software vendor, and (ii) they may mistreat security researchers by compensating less than guaranteed and no means to appeal against it. In this research [26], BGU researchers propose a permissioned Blockchain-based framework that addresses the drawbacks of existing bug bounty programs. The framework allows a confidential exchange of vulnerabilities and compensations using smart contracts. In cases of policy violation, security researchers can appeal to a trusted group of security experts called arbitrators, that can force the software vendors to compensate the security researchers fairly. A formal evaluation of the proposed framework using TLA+ specification supports the viability of the proposal. A Hyperledger Fabric-based prototype is implemented to simulate the proposed framework. The analysis of the framework uses a game-theoretic notion to argue that if the majority of arbitrators behave honestly, then the rational strategy of software vendors is to compensate security researchers that disclose vulnerabilities accurately. Similarly, rational security researchers do not gain any financial profit by playing unfairly.

5 Network-Centric Security (T1.2)

The main theme of Task 1.2 of the CONCORDIA project is network-centric security. Within this theme, the task identifies three broader areas of research that are at the intersection of networking and security. The *first* relates to making network infrastructure and services more resilient against (Distributed) Denial-of-Service (DDoS) attacks and more specifically by strengthening the DNS infrastructure. The *second* area involves methods to analyze encrypted network traffic, specifically towards detecting network-based threats. The *third* and final area involves the use of Software-Defined Networking (SDN) to advance networking, and particularly to the end of improved stability and resilience.

Within T1.2, the intent is to:

- Investigate proactive, coordinated and distributed strategies to defend against network-based attacks and DDoS attacks in particular, stemming from the DNS infrastructure;
- Develop techniques to analyze encrypted network traffic for security-related purposes such as monitoring, threat detection, attack mitigation and attribution; and
- Investigate SDN as a means to form a trusted and resilient Internet for Europe.

5.1 Overview

In the paragraphs hereafter, we will provide a summary description of research efforts undertaken within the Task 1.2 context during year 3 of the project. Parts of the summary will be accompanied by a reference to a later, more detailed discussion, which are placed under Section 5.3. We will also outline how the research efforts relate to the pilot projects of the CONCORDIA project. Where applicable, we also consider how continued efforts relate to their counterparts in the previous years of the project.

Below we provide a Bird's-Eye view on year 3 efforts. The various categories of research focus are marked in bold.

Distributed Denial-of-Service (DDoS) attacks

In year 2, we started looking at the potential role of the DNS in DDoS. That is, situations in which the DNS itself is not necessarily on the receiving end of attacks but rather on the facilitating end. As a prominent example, consider that open DNS resolvers can be misused to bring about reflection and amplification DDoS attacks. The DNS plays a crucial role in connecting services and users on the Internet. As such, its availability and stability is of paramount importance to the Internet as a whole. We strengthened related efforts during year 3, by looking in more detail at

ways to identify more harmful infrastructure and problematic configuration of DNS zones, by further investigating the resilience properties of the DNS, and also by looking at the role of cloud service providers, which typically provide potentially harmful infrastructure. New to year 3, we have also looked at future-proofing the DNS in the face of quantum computing and identified other challenges in the DNS ecosystem.

In Section 5.3.1 we describe efforts to create a comprehensive and all-encompassing overview of the modern DNS, its ongoing development, and open challenges.

Among open challenges is the post quantum computing landscape of the security extensions for the DNS (DNSSEC), which calls for, among others, the replacement of to-be insecure signature algorithms. In Section 5.3.2 we describe efforts towards: (1) identifying problems that hinder the replacement of algorithms; (2) developing and testing tools to simplify the deployment of secure algorithms; and (3) assessing the suitability of cryptographic algorithms for DNSSEC that cannot be broken by a quantum computer.

IP anycast is a technology to make networked services more resilient against attacks. Year 2 already put IP anycast in focus in T1.2. Partners developed and presented a methodology to infer whether a given Internet host is anycast. In year 3 we built on these efforts to assess DNS properties. Specifically, in Section 5.3.3, we describe recent work to assess the resilience of DNS infrastructure, which in part requires knowledge of anycast deployment.

In Section 5.3.4 we discuss our investigation efforts into configuration diversity among open resolvers. The aim was to characterize open DNS resolvers in terms of their ability to bring about varying attack strengths, which can help prioritize efforts by operators to root the most-threatening resolvers out first.

Cloud providers generally are well-provisioned in terms of connectivity. As a result, inadvertently open DNS resolvers hosted in cloud infrastructure by customers, but also the DNS infrastructure of cloud providers themselves, could make for potentially harmful, attack-enabling infrastructure. In Section 5.3.5 we discuss efforts to investigate the role of cloud providers in attacks.

During year 2 we (also) observed that the DNS configuration of some domain names lends itself well for amplification attacks, particularly in the case of zones with many or far above average-sized resource records. In year 3 we investigated this in more depth. Section 5.3.6 describes partner efforts towards identifying and classifying problematic domain name configurations, and looking at the potential benefits of disabling more problematic DNS query types.

At the intersection of DNS and DDoS, Section 5.3.7 outlines *DNSAttackStream*, which is a novel system that fuses indicators of DDoS activity with large-scale

DNS data to provide real-time situational awareness, in particular, about the impact of DDoS attacks on DNS infrastructure and dependent domain names.

Analyzing encrypted network traffic

From the start of the project, T1.2 partners have been concerned with overcoming barriers to analyze and safeguard encrypted traffic. Threat intelligence and intrusion detection are key to safeguarding network security, but the widespread application of encryption (and thus confidentiality of payload) poses challenges for monitoring solutions. Year 3 has seen efforts build on foundations developed in earlier years.

Section 5.3.8 summarizes a state-of-the-art survey on analyzing encrypted connections, which involves a breakdown of techniques and methods that appear in related works, as well as their limitations.

In Section 5.3.9 we introduce HeaderHunter, which is a fast intrusion detection system that uses signatures of network traffic. These signatures can be built even when traffic is encrypted.

QUIC is a fairly new protocol, which provides confidentiality through encryption, and thus is also challenging from a monitoring perspective. In year 2, we saw the initial steps towards processing QUIC packets for flow export, which is key input to flow-based network monitoring solutions. These efforts continued in year 3. Section 5.3.10 details work on dealing with newer QUIC versions, as well as ongoing, novel-in-scale measurement efforts to gain insights into QUIC use in the wild.

Network traffic data lost in encryption may be supplemented from other sources that are commonly being monitored. Section 5.3.11 describes our research in enriching encrypted IP flows with data gathered by host-based monitoring.

Section 5.3.12 reports on the cyber threat identification using publicly available data about cyber threats and data from network monitoring in the encrypted network environment.

Software-Defined Networking (SDN)

SDN, which is the third research area of Task 1.2, has, relatively speaking, seen the least research focus over the past years when compared to, e.g., DDoS and encrypted traffic analyses. Year 3 has however – in the SDN space – led to some promising results that aid SDN deployments. In Section 5.3.13 we describe work on getting a better understanding of distributed policy conflicts in SDN-enabled networks. These conflicts can result from conflicting information from multiple applications across nodes. The discussed study promotes approaches to conflict resolution and interpretation, which can improve SDN deployments.

Next to a strong focus on DNS and DDoS, and some work in the SDN area, year 3 has, similarly to the previous year, also seen efforts in the blockchain space. In year 2, T1.2 researchers tackled scalability issues of blockchain networks. In this year, the focus shifted slightly. Section 5.3.14 describes Policy-based Network Management (PBNM) efforts to aid network (security) management (e.g., firewall configuration) using blockchain techniques.

5.2 Link between T1.2 efforts and the CONCORDIA pilots

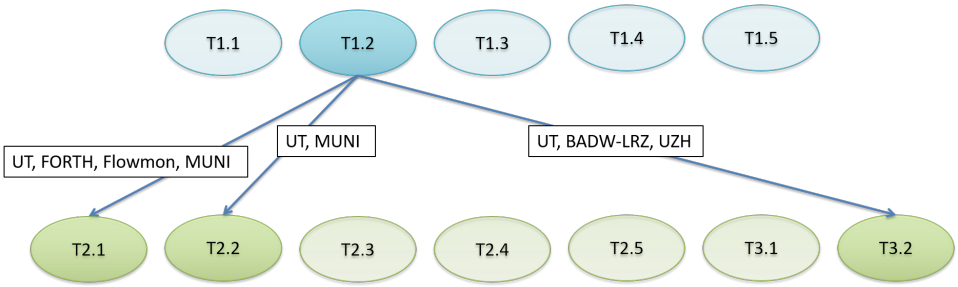

Figure 11: Relation between Task 1.2 and the CONCORDIA pilots

Figure 11 shows in which way the efforts undertaken in T1.2 relate to the pilots of the CONCORDIA project.

1. Task 1.2 continues to see a strong research focus on topics that related to the DNS and security. Given the critical role of proper functioning of the DNS for most networked services, efforts towards increasing stability, security, and resilience benefit virtually any application that depends on the DNS. This thus benefts any CONCORDIA task in which applications are being developed that depend, in some way, on the correct functioning and continued availability of the DNS;

2. The extensive work on encrypted network traffic analysis within T1.2, along with the tools and software (prototypes) that are being and have been developed in the process, benefit CONCORDIA tasks in which threats need to be learned from otherwise confidential connections. As examples, consider Tasks T2.1 (Telco threat intel), T2.2 (Finance threat intel) and T3.2 (Clearing house for Europe);

3. Task 1.2 involves a number of efforts that enable or (further) develop data sources. These data sources provide valuable input to tasks in which situational awareness (e.g., threat visibility) is key. The link with T3.2 continues to be a prominent example in this respect. The T1.2 work on DDoS attacks and DNS infrastructure characterization have been, in part, integrated into the clearing house pilot in the form of *fingerprints*. Research from year 3 has led to additional type of fingerprinting information;

4. Our work on inferring and characterizing IP anycast deployments and – new from year 2 – assessing for specific services how much benefit anycast brings, favors tasks that involve anycast deployment or fingerprinting anycast services (e.g., T3.2);
5. The progress that is being made on blockchain-related topics can prove useful towards tasks that are considering to leverage blockchain for information sharing.

5.3 Summaries of T1.2 Research Activities

The following sections contain more detailed write-ups of the research efforts that were introduced at the beginning of this section.

5.3.1 Addressing the Challenges of Modern DNS: A Comprehensive Tutorial

Contact: Olivier van der Toorn (UT) Moritz Müller (UT)

Since its first specification, DNS has been extended in numerous documents to keep it fit for today's challenges and demands. And these challenges are many. Revelations of snooping on DNS traffic led to changes to guarantee confidentiality of DNS queries. Attacks to forge DNS traffic led to changes to shore up the integrity of the DNS. Finally, Denial-of-Service attacks on DNS operations have led to new DNS operations architectures.

All of these developments make DNS a highly interesting, but also highly challenging research topic. Our work towards a tutorial – aimed at graduate students and early-career researchers – provides a overview of the modern DNS, its ongoing development and its open challenges. The tutorial has four major contributions. We first provide a comprehensive overview of the DNS protocol. Then, we explain how DNS is deployed in practice. This lays the foundation for the third contribution: a review of the biggest challenges the modern DNS faces today and how they can be addressed. These challenges are (i) protecting the confidentiality and (ii) guaranteeing the integrity of the information provided in the DNS, (iii) ensuring the availability of the DNS infrastructure, and (iv) detecting and preventing attacks that make use of the DNS. Last, we discuss which challenges remain open, pointing the reader towards new research areas.

Our tutorial paper [27] shows many challenges which accompany the DNS. There are many areas of overlap with the contexts of CONCORDIA, making our tutorial paper a good way of getting up to speed with some aspects of CONCORDIA that relate to DNS

5.3.2 Making DNSSEC Future Proof

Contact: Moritz Müller (SIDN, UT)

The DNS Security Extensions (DNSSEC) play an important role in securing the naming system of the Internet, by adding authenticity and integrity. DNSSEC, however, can only fulfil this role reliably when operators can easily adopt modern cryptographic algorithms and deprecate algorithms that are considered insecure or inefficient. This becomes even more important with quantum computers appearing on the horizon. An attacker with a quantum computer could be able to break all cryptographic algorithms currently used in DNSSEC in polynomial time. This would render DNSSEC useless and would severely undermine the trust in the DNS and the Internet.

We contributed significantly to DNSSEC's ability to keep the DNS secure in the future, by: (1) identifying problems that hinder the replacement of old, insecure algorithms, with new ones. (2) Developing and testing tools to simplify the deployment of secure cryptographic algorithms. (3) Assessing cryptographic algorithms that cannot be broken by a quantum computer, and their suitability for DNSSEC [28].

Our contributions are based on large scale active and passive measurements of the whole DNS ecosystem over a period of several years. This allowed us to learn from algorithm replacements in the past, from the perspectives of standardization bodies, the domain registration channel, domain name operators, and resolver operators. Also, we studied unique events in the DNS history as the first, namely the first ever Root KSK rollover. This rollover did not only come with great risks for many Internet users but also laid the foundation for future algorithm replacements at the most important component of the DNS name space. Our measurements highlighted several problems during the rollover and were followed closely by the DNS community. At the same time, our study also demonstrated that replacing the algorithm at the DNS root is feasible, paving the way for a future-proof DNSSEC.

5.3.3 Characterizing Anycast-Based Resilience of DNS Authoritative Name Servers

Contact: Raffaele Sommese (UT), Mattijs jonker (UT)

Over the past decades, IP anycast has proven to be an effective mechanism to enhance resilience in the DNS. Widely used among root nameservers, anycast has also started to be adopted at lower levels of DNS hierarchy, such as top-level domain (TLD) and second-level domain (SLD) nameservers. Using efforts and tooling developed in Y2 of CONCORDIA[3], we mapped and quantified the adoption of anycast to support authoritative domain nameservers for TLDs and SLDs [29]. Moreover, using a publicly available anycast census dataset of 2017, we were able to analyze the evolution of the adoption of anycast services. The results obtained show that adoption of anycast reached 97% of TLDs and 50% of SLDs. Finally, we discussed the implication of anycast adoption for other resilience mechanisms

[3]This work was presented in Section 5.4.4 of Deliverable D1.2.

of DNS showing a trend towards domain hosting concentration in anycast-based services.

5.3.4 A Matter of Degree: Characterizing the Amplification Power of Open DNS Resolvers

Contact: Ramin Yazdani (UT)

Reflection and Amplification (R&A) DDoS attacks continue to be a major source of disruption for networks. Open DNS resolvers are among widely misused reflectors to bring about R&A DDoS attacks. While the general position is that unchecked open resolvers should be rooted out, indiscriminate efforts to address the issue and take down resolvers have had limited effect, and millions of open resolvers remain available on today's Internet. This brings forward the question if we should not instead focus on eradicating the most problematic resolvers, rather than all open resolvers indiscriminately.

Contrary to existing studies, which focus on quantifying the existence of open resolvers, in our research in progress we investigated configuration diversity among open resolvers and aimed to characterize open DNS resolvers in terms of their ability to bring about varying attack strengths. Such a characterization brings nuances to the problem of open resolvers and their role in amplification attacks, as it allows for more problematic resolvers to be identified. Our findings show that the population of open resolvers lies above 2.6 M range over our one-year measurement period. On the positive side, we observe that the majority of identified open resolvers cut out when dealing with bulky and DNSSEC-related queries, thereby limiting their potential as amplifiers. We show, for example, that 59% of open resolvers lack DNSSEC support. On the downside, we see that a non- negligible number of open resolvers facilitate large responses to ANY and TXT queries (8.1% and 3.4% on average, respectively), which stands to benefit attackers. This work has been submitted and is currently under review

5.3.5 Mirrors in the Sky: On the Potential of Clouds in Reflection and Amplification DDoS

Contact: Ramin Yazdani (UT)

As a complement to efforts to characterize resolvers, we studied prospective reflectors in cloud provider networks and datacenters, postulating that the likely well-provisioned state of such infrastructure should guide more selective mitigation efforts. As part of our study, we identify and formalize six attack models under which the infrastructure of cloud providers can be misused to bring about reflection attacks. We assess the feasibility of these attacks on the Internet by conducting a proof of concept study on 19 public and top cloud providers. Our results reveal that a handful of providers expose various parts of their DNS infrastructure to outsiders, which attackers can use to attack provider infrastructure, its customers, and hosts

in external networks. We engaged coordinated vulnerability disclosure procedures with these providers. Our findings also reveal that, on average, roughly 12% of open DNS resolvers are hosted in cloud or datacenter networks, which makes them particularly worrisome. Furthermore, we confirm (and bolster) earlier findings that the vast majority are running in mobile and fixed ISP networks.

Our ongoing research can be employed by operators of well-provisioned networks (cloud providers) to minimize their exposure to be misused in R&A DDoS attacks, thus limiting the global power of R&A DDoS attacks.

5.3.6 ANYway: Measuring the Amplification DDoS Potential of Domains

Contact: Olivier van der Toorn (UT), Mattijs jonker (UT)

DDoS attacks threaten Internet security and stability, with attacks reaching the Tbps range. A popular approach involves DNS-based reflection and amplification, a type of attack in which a domain name, known to return a large answer, is queried using spoofed requests to reach high attack volumes. Do the chosen names offer the largest amplification, however, or have we yet to see the full amplification potential? And while operational countermeasures are proposed, chiefly limiting responses to 'ANY' queries, up to what point will these countermeasures be effective?

In this work [30] we make three main contributions. First, we propose and validate a scalable method to estimate the amplification potential of a domain name, based on the expected ANY response size. Second, we create estimates for hundreds of millions of domain names and rank them by their amplification potential. By comparing the overall ranking to the set of domains observed in actual attacks in honeypot data, we show whether attackers are using the most-potent domains for their attacks, or if we may expect larger attacks in the future. Finally, we evaluate the effectiveness of blocking ANY queries, as proposed by the IETF, to limit DNS-based DDoS attacks, by estimating the decrease in attack volume when switching from ANY to other query types.

Our results show that by blocking ANY, the response size of domains observed in attacks can be reduced by 57%, and the size of most-potent domains decreases by 69%. However, we also show that dropping ANY is not an absolute solution to DNS-based DDoS, as a small but potent portion of domains remain leading to an expected response size of over 2,048 bytes to queries other than ANY.

This work has direct impact on CONCORDIA, as we investigated the DDoS problem from the DNS side. By combining active DNS measurement data with DDoS honeypot data we show how the situation looks like on the DNS front.

5.3.7 DNSAttackStream: from DNS Attack Detection to Proactively Assessing Impact on Infrastructure

Contact: Raffaele Sommese (UT), Mattijs Jonker (UT)

Distributed Denial of Service (DDoS) attacks are one of the most disruptive attacks on today's Internet. With their rising firepower, they can cause severe issues to DNS infrastructure and the global Internet. Obtaining a real-time insight into these DNS DDoS attacks and responding to them with reactive measurements can provide opportunities to understand the practical consequences and impact of these attacks.

DNSAttackStream (Figure 12) represents the first step in integrating different data streaming sources to respond to attacks. We started with RS-DoS (Reflected Spoofed Denial of Service) attack information, collected by the UCSD STARDUST project, and joined it with OpenINTEL live measurements. DNSAttackStream merges the obtained information of IP addresses inferred to be under attack based on the STARDUST (UCSD Network Telescope) data every 5 minutes with the list of IP addresses of authoritative nameservers measured by OpenINTEL. This mechanism allows us to provide insights into the number of authoritative nameservers and related Second Level Domains (SLDs) affected by attacks. Moreover, we build a DNS reactive measurement platform to respond to attack events with additional measurements. These measurements provide us insights into the impact of attacks and will help identify working strategies for improving DNS resilience.

Analyzing a selection of attacks on DNS operators (e.g., large registrars), we see evidence of increased round trip time and failure rate. In particular, during two large DDoS attacks against a European registrar, we have seen a tenfold increase in resolution time. In our ongoing research, we are developing a methodology to systematically characterize these attacks and their impact based on the collected data.

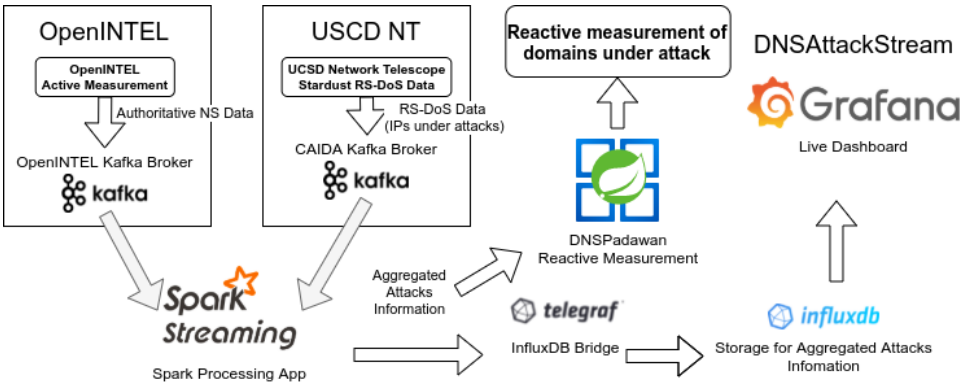

Figure 12: DNS Attack Stream

5.3.8 Encrypted network traffic analysis and inspection

Contact: Eva Papadogiannaki (FORTH), Sotiris Ioannidis (FORTH)

The adoption of network traffic encryption is continually growing. Popular applications use encryption protocols to secure communications and protect the privacy of users. In addition, a large portion of malware is spread through the network traffic and takes advantage of encryption protocols to hide presence and activity. Entering into the era of completely encrypted communications over the Internet, we must rapidly start reviewing the state-of-the-art in the wide domain of network traffic analysis and inspection, to conclude if traditional traffic processing systems will be able to seamlessly adapt to the upcoming full adoption of network encryption. In year 3 we produced a survey [31], in which we examine the literature that deals with network traffic analysis and inspection after the increased use of encryption in communication channels. We notice that the research community has already started proposing solutions on how to perform inspection even when the network traffic is encrypted and we demonstrate and review these works. As part of our literature survey, we also present the techniques and methods that existing works use, as well as their limitations. Finally, we examine the countermeasures that have been proposed in the literature in order to circumvent traffic analysis techniques that aim to harm user privacy.

Our ongoing focus is on encrypted network traffic analysis and inspection and specifically we aim to generate signatures for intrusion detection even in encrypted networks. As part of ongoing work, we are also investigating network packet processing acceleration using commodity hardware (e.g., GPUs).

5.3.9 Acceleration of Intrusion Detection in Encrypted Network Traffic Using Heterogeneous Hardware

Contact: Eva Papadogiannaki (FORTH), Sotiris Ioannidis (FORTH)

More than 75% of Internet traffic is now encrypted, and this percentage is constantly increasing. The majority of communications are secured using common encryption protocols such as SSL/TLS and IPsec to ensure security and protect the privacy of Internet users. As previously explained in Section 5.3.8, encryption can be exploited to hide malicious activities, camouflaged into normal network traffic. Traditionally, network traffic inspection is based on techniques like deep packet inspection (DPI). Common applications for 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 payload content are becoming less effective, demanding the development of more sophisticated techniques in order to adapt to current network encryption trends.

To start addressing the aforementioned demand, we presented HeaderHunter [32]. HeadHunter is a fast, signature-based intrusion detection system, that works even with encrypted network traffic. HeadHunter involves signatures that are generated only on the basis of network packet metadata, which can be extracted from packet headers. We validated HeadHunter on different heterogeneous hardware architectures, examining the processing acceleration of the intrusion detection engine.

5.3.10 Encrypted Traffic Analysis – NetFlow QUIC Plugin

Contact: Martin Holkovič (Flowmon), Tomáš Plesník (MUNI)

We aim to increase the capabilities of encrypted network traffic analysis. This year we have been focusing on the protocol QUIC. The main purpose of the QUIC protocol is to improve the user web experience by reducing web page load times compared to most commonly used TCP+TLS+HTTP/2 protocols. QUIC is based on the UDP protocol and uses TLS to encrypt transferred application data. Because the performance advantages of this new protocol are significant, the protocol is being gradually deployed by more Web applications. Our goal for this year was to analyze the protocol from the network monitoring perspective and implement a new NetFlow processing plugin for Flowmon monitoring solution based on this result. The plugin's purpose will be to increase the visibility into the encrypted traffic by extracting metadata about encrypted connections.

In the last year, we analyzed the specification of the QUIC protocol versions Q039 and Q040, and we created a plugin for those versions. However, the QUIC standard was not finished at that time yet, and several new versions were released since. Changes in the newer QUIC versions created incompatibility challenges with the previously implemented plugin. Unfortunately, we noticed that many QUIC versions are used simultaneously, which brings several difficulties related to network monitoring. To be able to properly monitor the protocol, all used versions should be supported. Instead of implementing a new version of the plugin that would support all new protocol versions, we have started measuring the occurrence of each version in real network traffic. The reason was to pick only versions that make sense to support.

We started our measurement efforts on January 1, 2021 in the Masaryk University network. These measurement efforts are still ongoing. As a large organization with more than 40,000 users, including more than 35,000 students and 6,000 employees, Masaryk University has a vast computer network with more than 25,000 communicating IP addresses every day. The university is connected to the CESNET2 academic network using two connection points. The monitoring point for measuring the QUIC protocol usage is located in only one connection point. Because of various anti-COVID restrictions and the school semester, the number of students and employees in the Masaryk University network was changing distinctly

over time. However, we do not expect that these changes are significantly affecting the results of our monitoring.

During the monitoring period, a final version of the QUIC protocol was released[4]. We decided to continue with the monitoring activity to see how the network applications will migrate to the final version of the protocol. The new version started to be used very soon. However, other versions of the protocol are still present in the network. Therefore, it is necessary to support the older versions as well. The most frequently observed versions of the QUIC protocol were 0x01 (RFC9000) and Q050.

We authored a report that describes the results from monitoring activity between January 1 and July 31, 2021. The report was published in the form of a blog post[5]. The blog post aims at QUIC protocol from the monitoring perspective, and therefore its focus is more generic than just counting the protocol's versions.

Because of the changes in the protocol specification, we had to reimplement our processing plugin. The current implementation can extract the version, connection ID, and SNI from the encrypted communications. These values are sent in the form of IPFIX records and sent to the collector (ipfixcol2[6]). What is happening with the data then depends on the collector. In the case of the Flowmon Networks solution, the network administrator can identify individual QUIC connections, visualize the usage of the protocol, create periodical reports, or detect malware that is trying to hide its activity by encrypting the transferred data.

By analyzing the QUIC protocol, we were able to detect and devise a list of mostly used versions, for which we have implemented a network processing plugin. The plugin improves the visibility into encrypted traffic by extracting basic information about the QUIC connections. We have also published a blog post that describes QUIC protocol from the monitoring perspective. With this contribution, we enable operators and consortium members to improve the security of their networks by extending their knowledge about encrypted connections. In the following year, we will continue monitoring the usage of different QUIC versions to react to new specifications. We will also continue testing the plugin to make it more stable and ready for customers deployment.

5.3.11 Enriching Encrypted IP Flows via Host-Based Monitoring

Contact: Stanislav Špaček (MUNI), Pavel Čeleda (MUNI)

The current approach to encrypted traffic analysis focuses on remaining unencrypted parameters, analyzing statistical parameters such as interpacket arrival times, and pattern matching. However, insight into encrypted traffic is also pos-

[4] https://datatracker.ietf.org/doc/html/rfc9000

[5] https://www.concordia-h2020.eu/blog-post/quic-protocol-from-the-monitoring-perspective/

[6] https://github.com/CESNET/ipfixcol2

sible by replacing the data lost in encryption from another source. With this motivation, we propose event-flow correlation; an approach correlating encrypted IP flows with server events captured by host-based monitoring. The server events have the advantage that they are always easily accessible to the service provider, while no one else has access to them. Consequently, no private user data are compromised during transmission, and at the same time, the service provider gains an opportunity to enrich IP flow monitoring with new data from events.

In the initial steps of our research, we focused on relations between DNS IP flows and logs of DNS resolvers, as the DNS protocol is a prime candidate for encryption in the near future. In particular, we examined the impact of encryption on DNS IP flow monitoring and its possible mitigation by correlating logged events to specific DNS flows. We set up a time-synchronized monitoring environment and collected a dataset of DNS events and flows. Then we designed and evaluated a correlation method using common features and a time window to identify relations between events and flows in the dataset. Our evaluation showed that it is possible to match related DNS events and flows with high accuracy, even when relying exclusively on features available in unencrypted DNS traffic. The research results were summarized in the paper Enriching DNS Flows with Host-Based Events that was submitted to and presented during the IFIP SEC 2021 conference [33].

We continued our research by collecting and analyzing a dataset containing HTTPS encrypted web traffic and web server logs. We then applied the same event-flow correlation method based on common features and time-window to investigate relations between HTTPS IP flows and events captured on the web servers. The method was redesigned to accommodate for HTTPS web traffic profile and evaluated on the HTTPS dataset. Our results showed that event-flow correlation is possible for HTTPS web traffic but also that it requires monitoring of custom features that are not present in default configurations of current web server applications. We summarize our results and lessons learned in the paper HTTPS Event-Flow Correlation: Improving Situational Awareness in Encrypted Web Traffic that is currently under review. The HTTPS dataset used in our research is currently being prepared for publishing in an open-access journal.

5.3.12 Identification of Cyber Threats in Encrypted Networks using Public Intelligence

Contact: Lukáš Sadlek (MUNI), Pavel Čeleda (MUNI)

The current approach for governance of cyber threats is based on sharing public intelligence that describes cyber threats (see T2.1 - Telecom Sector: Threat Intelligence for the Telecom Sector). Existing enumerations and knowledge bases (e.g., CVE, MITRE ATT&CK, CAPEC, CWE, and CPE) provide information about vulnerabilities, threats, and assets in a standardized form. On the contrary, current asset discovery tools analyze encrypted network traffic and provide information

about discovered assets. However, it is essential to provide more sophisticated methods for threat identification than the simple correlation of indicators (e.g., IP addresses). These low-level indicators can be easily changed by the attacker.

TTPs (Tactics, Techniques, and Procedures) are indicators of attacker's behavior. The attack techniques can be collected from network focused MITRE ATT&CK matrix. Modeling cyber threats using attack techniques allows identifying complex chains of attacker's actions and their concrete materialization within the protected network. The well-known concept for modeling multi-step cyber threats are attack graphs.

We introduced a novel threat model of a kill chain attack graph. The kill chain attack graph can depict paths consisting of the attacker's actions and represent steps of possible attacks mapped to kill chain phases. The methodology uses STRIDE as a taxonomy for MITRE ATT&CK techniques. We are preparing a conference research paper. Supplementary materials will contain open-source proof-of-concept implementation with the kill chain attack graph generator. Additional contribution is a custom ruleset containing a subset of MITRE ATT&CK techniques.

5.3.13 Experimental Examination of Distributed Conflicts in Software Defined Networks

Contact: Reinhard Gloger (BADW-LRZ)

Software Defined Networks (SDN) are a promising paradigm in computer networking and facilitate network policy enforcement by multiple entities through a centralized interface. The concept entails new possibilities for policy enforcement but also new problems. Policies by two applications on multiple nodes in a network can contradict and can result in connection failure or compromised security. Security in SDN depends on a global view of data plane state on all network devices, real-time processing of policy insertions as well as conflict detection and automated resolution. Related work has shown that firewall applications are subject to new challenges due to the novel properties of SDN.

Our work, which is currently ongoing, explores policy conflicts in SDN between multiple applications and across multiple nodes in a network. The research questions comprise a concept for reproducible experimental infrastructures to examine conflicts and the development of new applications that enforce policies. In experiments, the new implementations are deployed in combination with existing applications in order to discover new conflict classes. The experiments include implementations for load balancing, traffic engineering and filtering. While the implementations are based on Openflow as a technical framework, the formal definitions for conflict classes and presented detection algorithms apply to SDN in general.

For any discovered conflict classes, detection mechanisms are required. The role of presented implementations while taking into account security in SDN is evaluated. We examine existing work and the technical background to the presented approach in order to answer these topics. A software architecture for experiment automation is devised, and a set of designed and random network topologies serve as input for the experimental process. The results include new conflict classes based on the tested networks and applications, detection algorithms for the conflicts and an experimental architecture that facilitates further research.

Three approaches to generate random networks that display realistic properties and three new control applications are under investigation. Insights on the validity of discovered distributed conflicts and the soundness of presented detection methods as well as experiment automation are gained by a second set of experiments on additional, untested networks. The results advance the understanding of distributed conflicts in SDN and promote future approaches to conflict resolution and interpretation.

5.3.14 Employing Policy-based Network Management to Blockchain Selection

Contact: Eder J. Scheid (UZH), Bruno Rodrigues (UZH), Burkhard Stiller (UZH)

Computer networks face challenges regarding the management of different device configurations (*e.g.,* firewall rules) and the interoperability between them. Policy-based Network Management (PBNM) offers a means for flexible network management. This flexibility extends to security management as well. We investigated and applied PBNM in our work for this reason. We address various challenges in security management with the employment of Policy-based Management (PBM) applied in the management of physical networking devices (*i.e.,* PBNM), and to manage network functions (virtualized or not) decoupled from data-plane devices and concentrated in a central point as Software-Defined Networks (SDN) and Network Function Virtualization (NFV). Therefore, the use of rules (*i.e.,* policies) on user requirements outlines a viable approach to managing not only such interactions but also data stored in different Blockchains (BC), exploiting the BC's immutability and decentralization and selecting the most suitable BC to secure data. Hence, [34], developed within this task context, *(a)* details a standard for the refinement of high-level policies (named *intents*) to actions in BC nodes (*e.g.,* broadcast of a BC transaction) based on the traditional Policy Continuum; and *(b)* describes and comparing in detail the state-of-the-art in BC selection. Further, [34] *(c)* discusses the abstraction of technical details from BC selection policies, *(d)* presents directions on the employment of Machine Learning (ML) in the selection process, and *(e)* details key points on the interoperability of BC solutions given the growing number of BC platforms.

5.3.15 Stranger VPNs: Investigating the Geo-Unblocking Capabilities of Commercial VPN Providers

Contact: Olivier van der Toorn (UT)

Limiting content on the Internet, based on physical borders, does not always work, because technologies such as VPNs can artificially place users into a different country. This is the core business of commercial VPN providers. These providers have seen a large growth of their business and will very likely continue to grow. Allied Market Research for example has estimated that the total amount of sales in 2019 was around US $25.41 billion and that it will continue to increase to US $75.59 billion, by 2027.

Multimedia providers, who must abide by their licensing agreements, have started cracking down on the use of VPNs to circumvent these restrictions. The result is that a subset of commercial VPN providers has started to deploy new methods to avoid detection.

Our work is the first to investigate these new methods, by firstly identifying their exact mechanisms, and secondly measuring the use and scale of these methods over a longer period.

We have found that multimedia providers make use of commercial IP metadata lists to allow or deny access to their sites. The commercial VPN providers on the other hand try to not get included in these lists, to offer their geo-unblocking. Specifically, they are making use of three distinct methods to avoid inclusion:

- The use of “obscure” hosting providers. Big providers, such as for example Amazon Web Services, Google Cloud or Microsoft Azure are included in the metadata lists, but some very small local providers are not;
- The creation of “fake” (residential) ISPs, by acting as a LIR, or have the WHOIS information of an IP network thusly changed to resemble a regular local ISP;
- The use of so-called residential proxies. These proxies are often either compromised end-user (or IoT) devices, or the users willingly participate in return for some form of compensation.

We are in the process of preparing our results for submission to a peer-reviewed security venue.

6 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:

- security by design
- malware analysis
- system security validation and zero-days
- UAV Resilience

6.1 Overview

The T1.3 research undertaken in 2021 about system-centric security can be classified into four clusters, in relation to the main research topics listed above:

Security by design

In Section 6.3.1, the main result of thesecurity by design: adaptive software and OSs is related to failure diagnosis in High performance Computing Systems (HPC). In this work conducted by ULANC, a novel approach dubbed IFADE has been devised and distributed in open-source. In Section 6.3.2 on Detecting Service dependencies, studies focused on assessing operating system dependability in monolithic systems, which is quite a challenging question. A novel approach has been developed by ULANC combining static and dynamic analysis to preform efficient Software fault injections (SFI). ULANC's approach has been evaluated on Linux file systems successfully.

Malware analysis

In Sections 6.3.3 and 6.3.4 on malware analysis, JUB and UL teams concentrate upon dynamic analysis of malware. Indeed, malware is today heavily protected using in particular *packers* and as a result static analysis is very challenging. On the contrary dynamic analysis allows executing malware into a secure sandbox in order to "see" malware behaviors. That said, there are a couple of issues. The first issue is to recover the payload, which is the aim of UL works and the development of the tool API-Xray. The second issue is to trigger several malware executions by providing different sandbox environments, which is essential to "see" malicious behaviors.

System security validation and zero-days

On System security validation and zero-days, the works split in two directions. The first one in Section 6.3.5, led by ULANC, consists in demonstrating that current Programmable Logic Controllers (PLC) programming practices facilitates new

vulnerabilities. A proof of concept has been built, named PCaaD, to validate the approach to a large number of PLC. The second one in Section 6.3.6, led by UL, consists of attempting to identify the compilation toolchain of compiled libraries such as COTS (commercial off-the-shelf libraries). The mid-term goal is then to be able to identify vulnerable libraries.

UAV Resilience

Lastly in Section 6.3.7, the work on the resilience of UAV missions in collaboration with ACS, introduces methods to not only make systems more robust to cyber-attacks but also evaluate the impact of cyber-attacks on the system taking into account system dependencies and roles of system components and plan deployments accordingly. It provides methods to validate whether the goals of the system can be attained when impacted by cyber-attacks.

6.2 Link between T1.3 work and CONCORDIA pilots

The research performed within Task 1.3 is linked to the CONCORDIA pilots as follows.

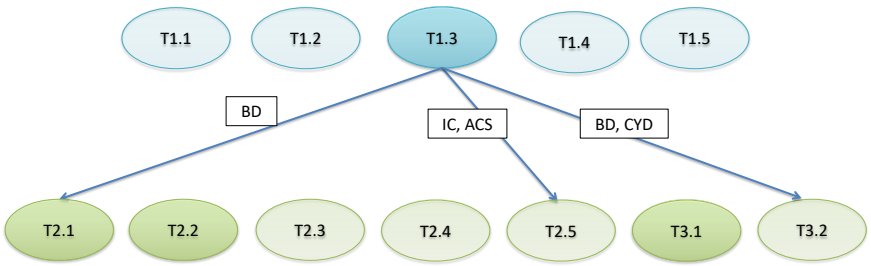

Figure 13: Links between task T1.3 and the CONCORDIA pilots

1. BD contributed to the telecom pilots T2.1 with the collaboration of Telecom Italia. BD contribution was on the data validation and augmentation with detection information. The collaboration is still active and it will continue to be visible through the MISP information enrichment.

2. Together, IC and ACS has pursued their works on analyzing the Viability of UAV Missions Facing Cyber Attacks with a fairly comprehensive report included in the T2.5 deliverable.

3. BD and CYD contributes also to T3.2 Thread intelligence bricks.

6.3 Summaries of T1.3 Research Activities

The sections below contain an account of the diverse research activities of task T1.3.

6.3.1 Security by design: adaptive software and OSs

Failure Diagnosis for Cluster Systems using Partial Correlations

Contact: Neeraj Suri (ULANC)

Failures have expensive implications in HPC (High Performance Computing) systems. Consequently, effective diagnosis of system failures is desired to help improve system reliability from both a remedial and preventive perspective. As HPC systems conduct extensive logging of resource usage and system events, parsing this data is an often advocated basis for failure diagnosis. However, the high levels of concurrency that exist in HPC systems cause system events to frequently interleave in time and, as such, certain interactions appear or become indirect which will be missed by current failure diagnostics techniques. To help uncover such indirect interactions, in this paper [35], we develop a novel approach that leverages the concept of partial correlation. The novel failure diagnostics workflow - called *IFADE* - extracts partial correlation of resource use counters and partial correlation of system errors. As part of our contributions, we (a) compare our diagnostics approach with current ones, (b) identify two previously unknown causes of system failures, validated by system designers and (c) provide insights into Lustre (filesystem) I/O and segmentation faults. IFADE has been put in the public domain to support system administrators in failure diagnosis.

6.3.2 Detecting Service dependencies

Fast Kernel Error Propagation Analysis in Virtualized Environments

Contact: Neeraj Suri (ULANC)

Assessing operating system dependability remains a challenging problem, particularly in monolithic systems. Component interfaces are not well-defined and boundaries are not enforced at runtime. This allows faults in individual components to arbitrarily affect other parts of the system. Software fault injection (SFI) can be used to experimentally assess the resilience of such systems in the presence of faulty components. However, applying SFI to complex, monolithic operating systems

poses challenges due to long test latencies and the difficulty of detecting corruptions in the internal state of the operating system. In this paper [36], we present a novel approach that leverages static and dynamic analysis alongside modern operating system and virtual machine features to reduce SFI test latencies for operating system kernel components while enabling efficient and accurate detection of internal state corruptions. We demonstrate the feasibility of our approach by applying it to multiple widely used Linux file systems.

6.3.3 Dynamic analysis and recovering an executable malware program

Contact: Jean-Yves Marion (UL)

As malware's APIs (Application Programming Interface) provide rich information about malicious behavior, one common anti-analysis strategy is API obfuscation, which removes the metadata of imported APIs from malware's header and complicates API name resolution from API callsites. In this way, even when security analysts obtain the unpacked code, a disassembler still fails to recognize imported API names, and the unpacked code cannot be successfully executed.

The goal is to reconstruct an executable payload from an obfuscated and packed binary.

Based on the process memory when the original entry point (OEP) is reached, we develop [37] a hardware-assisted tool, API-Xray, to reconstruct import tables. Import table reconstruction is challenging enough in its own right. Our core technique, API Micro Execution, explores all possible API callsites and executes them without knowing API argument values. At the same time, we take advantage of hardware tracing via Intel Branch Trace Store and NX bit to resolve API names and finally rebuild import tables. Compared with the previous work, API-Xray has a better resistance against various API obfuscation schemes and more coverage on resolved Windows API names. Since July 2019, we have tested API-Xray in practice to assist security professionals in malware analysis: we have successfully rebuilt 155,811 executable malware programs and substantially improved the detection rate for 7,514 unknown or new malware variants. This is a work joint with The University of Texas at Arlington, Hubei Normal University and Wuhan University.

6.3.4 A Comparison of Upper Confidence Bounds For Exploration in Active Malware Analysis

Contact: Abhilash HOTA (JUB)

Malware analysis is the process of determining the behavior of a program in order to decide whether it is malicious or not. Dynamic malware analysis seeks to map the behavior of a previously unknown program by executing it in a restricted environment and collecting execution traces. Current dynamic analysis approaches are

however limited, in that only a single program execution is observed. For many recent malware samples, certain malicious actions are only triggered under specific conditions like the presence or absence of a specific file or the execution of a specific command by a user. Thus dynamic analysis tools can be insufficient when it comes to exploring multiple possible execution paths. Active Malware Analysis (AMA) aims to develop dynamic analysis systems that perform actions in order to trigger different behaviors of the malware being analyzed. Payload deployment can often be hidden behind conditional requirements that look for specific user actions, especially in mobile malware. Initial approaches to AMA have selected user inputs based on existing data about user behavior patterns or simply used a pseudo-random selection of possible actions. Further work has investigated the use of reinforcement learning based agents that learn to select triggering actions and develop models of the malware samples being investigated.

This work discusses various Upper Confidence Bound (UCB) algorithms and compares their effectiveness in a Monte-Carlo Tree Search (MCTS) based approach to Active Malware Analysis. Malware analysis is modeled here as a stochastic game. The dataset is a collection of Android malware samples. The analysing agent is based on MCTS and learns to select trigger actions during the analysis process in order to gather information about multiple execution paths. A malware model is developed in the form of an API call graph for each malware sample. We implement an analysis framework consisting of an Android emulator running an Android 10.0 image, an observation module that can record the malware behaviour and hook API calls and the analyser that selects trigger actions. We use the Frida framework to hook the API calls and monitor the malware behaviour when triggered by the analyser. The analyser is based of different implementations of MCTS using different UCB algorithms in the selection stage for each implementation. We implement UCB1, UCB-Tuned, UCB-Normal [38] and Kullback-Leibler UCB (KL-UCB) [39] and compare execution times during the selection stage of MCTS. We have trained the models on a dataset consisting of 15000 Android malware samples and results obtained show that KL-UCB minimizes the number of trigger actions required but UCT-Tuned manages a comparable learning rate with less overall execution time. UCT-Normal is particularly affected by the reward distribution and as such cannot manage a comparable performance given that the rewards are not normally distributed in this application. The findings are being submitted for publication. The eventual goal is to integrate the trained agents in a dynamic analysis workflow. The agents generate API call graphs for the malwares being analysed, which can then be used for training further models to identify malware on Android devices. Malware analysis in its current form is largely a manual process. Automating this part of the malware analysis workflow would help deal with the increased volume of new malware being released in the market.

6.3.5 PCaaD: Towards Automated Determination and Exploitation of Industrial Systems

Contact: Neeraj Suri (ULANC)

Over the last decade, Programmable Logic Controllers (PLCs) have been increasingly targeted by attackers to obtain control over industrial processes that support critical services. Such targeted attacks typically require detailed knowledge of system-specific attributes, including hardware configurations, adopted protocols, and PLC control-logic, i.e., process comprehension. The consensus from both academics and practitioners suggests that stealthy process comprehension obtained from a PLC alone, to execute targeted attacks, is impractical. In contrast, we assert that current PLC [40] programming practices open the door to a new vulnerability class, affording attackers an increased level of process comprehension. To support this, we propose the concept of Process Comprehension at a Distance (PCaaD), as a novel methodological and automatable approach towards the system-agnostic identification of PLC library functions. This leads to the targeted exfiltration of operational data, manipulation of control-logic behavior, and establishment of covert command and control channels through unused memory. We validate PCaaD on widely used PLCs through its practical application.

6.3.6 Recovering the compiling chain used to generate a stripped binary code

Contact: Jean-Yves Marion (UL) Tristan Benoit(UL)

Identifying the toolchain provenance, i.e. the compiler family (e.g. Visual Studio or GCC), the compiler version (e.g. 10.0, 12.0) and its optimization options, that have been used to produce a given stripped binary code is an important problem in at least two scenarios:

- Determination of security flaws inside binary codes. Applications are often built by linking together commercial off-the-shelf libraries (COTS). While allowing faster development cycles, developers do not have the source code of these COTS and do not know the compiling chain used to generate them. This is an important issue in software maintenance and long-term support as compilers may inject vulnerabilities that are discovered after the COTS released and after the deployment of the applications that used them.
- Identification of known functions. Library function identification in a binary code is another primary issue for software maintenance and security, such as clone detection and function similarities.

We present [41] a Graph Neural Network framework at the binary level to solve this problem, with the idea to take into account the shallow semantics provided by the binary code's structured control flow graph (CFG). We introduce a Graph Neural Network, called Site Neural Network (SNN), dedicated to this problem. To

attain scalability at the binary level, feature extraction is simplified by forgetting almost everything in a CFG except transfer control instructions and performing a parametric graph reduction. Our experiments show that our method recovers the compiler family with a very high F1-Score of 0.9950 while the optimization level is recovered with a moderately high F1-Score of 0.7517. A comparison with a previous work demonstrates the accuracy and performance of this framework.

6.3.7 Resilience to Cyber Attacks of UAV Missions

Contact: Emil Lupu (ICL)

Analysing the resilience of a system when it faces cyber-attacks requires new methodologies that take into account not only the progression of the attack but also the impact that the attack has on the functionality of the system. Such methodologies are at the moment lacking in the literature. In collaboration with T2.5, we have focussed on developing such a methodology within the context of missions of UAVs, and we have illustrated our methodology using a scenario based on a firefighting mission. The main principle of the method developed is to simultaneously analyse the progression of a (cyber-)attack and its impact on the vehicles tasked with a mission in order to determine: a) the degree to which the mission can be completed, and thus if it is still viable, b) if alternative configurations of the mission roles UAVs fulfil can be adopted to increase the degree of completion of the mission and c) how much help additional (e.g. standby) resources can provide towards completing the mission. Our method can be used both in the planning of the mission and for decision making during mission operation. Our approach is based on modelling attack progression and analysing the impact of the attacks on the mission using Petri nets and more specifically Stochastic Well-Formed Nets, analysing the model with the GreatSPN tool and then combining the outcome with the analysis with a system performance model of the mission's degree of completion and viability. Although the method was developed within the context of UAVs, it has broader applicability to systems with reusable resources playing different roles in the operation of the system that is subjected to cyber-attacks. A more comprehensive report of this activity is included in D.2.3 and the work was published in [42].

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

Task 1.4 (T1.4) of the CONCORDIA project deals with data and application-centric security, in the context of cloud computing. Such kind of resources are highly exposed to security attacks, since they are typically distributed over several cloud providers, and may be shared among different clients. Furthermore, the migration of data and applications to cloud infrastructure may also permit access to and alteration of them by unauthorized parties due to improper access control. In order to organize the research efforts in the task, as well as to track any possible collaboration efforts with the different tasks in the other WPs, the task identifies three areas of research:

- How to protect (big) data before and after storing it in the cloud.
- How to protect the cloud services themselves.
- How to perform behavioral analysis on the applications to detect suspicious behaviors or attacks.

7.1 Overview

In the following paragraphs, we report on the research activities performed within T1.4 for year 3 of the project. The activities are categorized based on the research areas mentioned before. Although some work has been developed for different domains (e.g. vehicular, cellular, etc.), it can be adapted to fit the needs of a cloud infrastructure.

It is worth noting that 5G networks are virtualized networks consisting of virtual Network Functions (vNFs) distributed at various clouds from edge to regional and national clouds. Consequently, 5G network security includes in the one hand the protection of the mobile network itself and on the other hand the protection of the cloud infrastructure. With 5G networks, organisations can monitor malicious actions and react in time to stop the cyber attacks even before they happen. 5G and Cloud Computing together would make the cyber technology space more agile and robust like never before.

Application Behavioral Analysis

Section 7.3.1 presents an updated approach of a previous work, that maps private memory pages in a running application on-the-fly that contain privileged information, accessible only by a specific part of the program. In Section 7.3.2, there are details about a custom Linux kernel that introduces security policies during runtime of an application. Section 7.3.3 mentions an Intrusion Response Framework for the vehicular domain. In Section 7.3.9, the authors present and compare different machine learning algorithms for anomaly detection in the cellular IoT domain. Section 7.3.11 is a survey about techniques for mining blockchain data, with spe-

cial attention towards anomaly/fraud detection. Section 7.3.13 describes an automated orchestration methodology for security chains in order to secure connected devices and their applications, by establishing behavioral models and inferring security constraints from them.

Protection of Cloud Services

In Section 7.3.5, the authors present their elderly care at home approach that uses millimeter wave sensors and 5G network slicing by obtaining data from sensors and have it relayed through an isolated network slice and analysed by the designated center for healthcare support. In Section 7.3.6, they detail a dedicated 5G network slice for at home healthcare systems for the elderly, focusing on network isolation and data privacy and security. Section 7.3.7 deals with the inherent networking stack performance issues related with 5G environments. Section 7.3.8 proposes an approach to enhance the isolation of network slices, by employing the Enhanced VPN+ technology, while tackling DDoS and flooding attacks on the communication between the 4G/5G core networks and the Cloud-Radio Access Network. Section 7.3.10 concerns a discussion about the problem of representing cardinality constraints in graph databases. Section 7.3.12 deals with the automation of security enhancements of cloud composite services during the migration of their elementary resources. In Section 7.3.14, the authors detail their works on security assurance: a multi-dimensional certification scheme for cloud services, a methodology integrating risk management and security assurance, and a security assurance framework based on API and VPN for the assessment of hybrid systems. In Section 7.3.15, the authors analyze reputation-based systems, which are widely used in (cloud) service provisions, identifying challenges and solutions in implementations based on blockchains.

Protection of (big) Data

In Section 7.3.4, the authors present a Publish/Subscribe model for supporting Vehicle to Cloud (V2C) communication which enables encryption and access control of messages/data published by vehicles. Section 7.3.16 presents a scenario where individuals express their preferences on how their data have to be consumed, when leveraging blockchain to manage a secure execution of collaborative processes among organizations.

7.2 Link between T1.4 work and CONCORDIA pilots

Collaboration between T1.4 and some of the project pilots has been ongoing and has yielded some interesting results. In this sense, we consider it one of the highlights of the task for the third year. Figure 14 indicates these relationships.

Strong collaboration has been developed with the telecommunications sector (T2.1). As WP1 deals with the scientific research within CONCORDIA which is

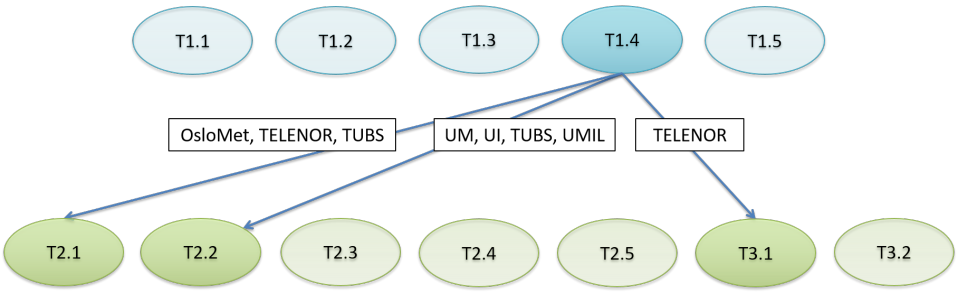

Figure 14: Links between task T1.4 and the CONCORDIA pilots

mostly performed by the academic partners, specifically in the context of T1.4 they are able to contribute to T2.1 by researching and developing software that adheres to the requirements set by the Telecom sector partners. Derived implementations can also be integrated into the Threat Intelligence Platform, with the option to propose them for integration directly into the MISP open source project. Collaboration between the two tasks has been strengthened further, by taking part in virtual meetings where issues raised by the respective partners as well as possible solutions are being discussed.

Additionally, collaboration with the finance sector (T2.2) has been ongoing. Technologies proposed/developed by academic partners (secure data exchanging protocol, big data analysis based on blockchain, etc.) have been taken into consideration for the Financial Threat Intelligence Platform.

Lastly, as we have mentioned before, all the collected information about the IoCs are shared with the project's Threat Intelligence Platform (T3.1). TELENOR and TIM have established and support the collaboration between T3.1 and T1.4 partners.

7.3 Summaries of T1.4 Research Activities

In this section we present details concerning the research efforts undertaken within the T1.4 context.

7.3.1 Securing Runtime Memory via MMU manipulation

Contact: Marinos Tsantekidis (ICS-FORTH), Vassilis Prevelakis (AEGIS GmbH)

It is often useful for a code component (e.g., a library) to be able to maintain information that is hidden from the rest of the program (e.g., private keys used for signing, or usage counters used for behavioral monitoring of the program). In this paper [43], we present an extension to a previously developed mechanism for controlling access to libraries, in order to implement a scheme that allows each library

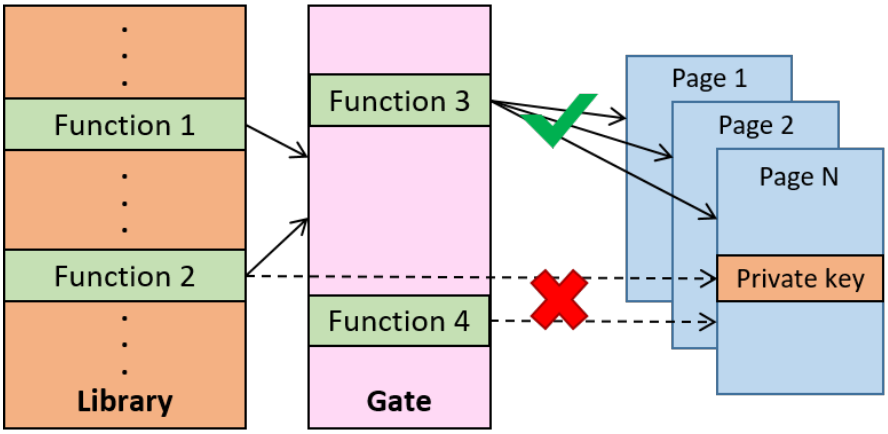

Figure 15: Secure memory mapping

to have its own private storage space (Figure 15). When running code outside the address space of a given library, the pages containing the private memory of that library are not mapped into the program's address space, hence are not accessible to the rest of the program. Finally, we present an API that allows library developers to utilize private storage.

7.3.2 MMU-based Access Control for Libraries

Contact: Marinos Tsantekidis (ICS-FORTH), Vassilis Prevelakis (AEGIS GmbH)

Code Reuse Attacks can trick the CPU into performing some actions not originally intended by the running program. This is due to the fact that the execution can move anywhere within a process's executable memory area, as well as the absence of policy checks when a transfer is performed. In our effort to defend against this type of attacks, in an earlier paper we present a Proof-of-Concept mitigation technique based on a modified Linux kernel where each library - either dynamically or statically linked - constitutes a separate code region. The idea behind this technique is to compartmentalize memory in order to control access to the different memory segments, through a gate. Taking our previous work one step further, in this paper [44] we present an updated version of our kernel-side technique, where we implement security policies in order to identify suspicious behavior and take some action accordingly.

7.3.3 Intrusion Response System for Vehicles: Challenges and Vision

Contact: Mohammad Hamad, Marinos Tsantekidis, Vassilis Prevelakis (TUBS)

Recently, significant developments were introduced within the vehicular domain, making the modern vehicle a network of a multitude of embedded systems communicating with each other, while adhering to safety-critical and secure systems specifications. Many technologies have been integrated within modern vehicles to give them the capability to interact with the outside world. These advances have significantly enlarged the attack surface. We already have numerous instances of successful penetration of vehicular networks both from inside the vehicle and from the outside. To face these attacks, many intrusion prevention and detection mechanisms were implemented inside a vehicular system. Nonetheless, even if all security mitigation is adopted, an attack still can happen. In critical-safety environments, such as the vehicle, the response to the attack is as essential as detecting the attack itself. Although Intrusion Response Systems (IRSs) have been adopted in other domains to add an extra layer of security, there is a lack of such systems in the vehicular field. In this work [45], we investigate the challenges and identify the requirements for integrating such a mechanism within the vehicle system. Besides, we present an IRS framework, which meets the identified requirements. Also, we discuss the integration of IRS through the vehicle system development and the different aspects which support such a process. Finally, we use the automated obstacle avoidance system to explain how we could develop intrusion response strategies and to measure the overhead of such security system.

7.3.4 SPPS: Secure Policy-based Publish/Subscribe System for V2C Communication

Contact: Mohammad Hamad (TU Munich), Emanuel Regnath (TU Munich), Jan Lauinger (TU Munich), Vassilis Prevelakis (TUBS), Sebastian Steinhorst (TU Munich)

The Publish/Subscribe (Pub/Sub) pattern is an attractive paradigm for supporting Vehicle to Cloud (V2C) communication. However, the security threats on confidentiality, integrity, and access control of the published data challenge the adoption of the Pub/Sub model. To address that, our paper [46] proposes a secure policy-based Pub/Sub model for V2C communication, which allows to encrypt and control the access to messages published by vehicles. A vehicle encrypts messages with a symmetric key while saving the key in distributed shares on semi-honest services, called KeyStores, using the concept of secret sharing. The security policy, generated by the same vehicle, authorizes certain cloud services to obtain the shares from the KeyStores. Here, granting access rights takes place without violating the decoupling requirement of the Pub/Sub model. The implementation of this protocol is based on the well-know Message Queuing Telemetry Transport (MQTT) protocol [7]. Experimental results show that, besides the end-to-end security protection, our proposed system introduces significantly less overhead (almost 70% less) than the state-of-the-art approach SSL when reestablishing connections,

[7] https://mqtt.org/

which is a common scenario in the V2C context due to unreliable network connection.

7.3.5 A Secure 5G Eldercare Solution Using Millimeterwave Sensors

Contact: Boning Feng (OsloMet), Akihiro Kajiwara (University of Kitakyushu), Van Thuan Do (OsloMet), Niels Jacot (Wolffia AS), Bernardo Santos (OsloMet), Bruno Dzogovic (OsloMet), Thanh van Do (Telenor, OsloMet)

The world is ageing fast and the need of efficient digital solution enabling elderly people to age at home is getting urgent. Unfortunately there is so far no such a solution which is sufficiently efficient, customizable, secure and reliable. This paper [47] presents a solution called Ageing@home which is efficient, easy to deploy, privacy preserving, unobtrusive and customizable by making use of millimeter wave sensors, 5G network slicing and open unifying IoT platform. The paper provides a detailed description of the advantageous use of the millimeter wave sensors and present the proposed 5G network slicing alternative.

An open unifying IoT platform capable of bridging diverse heterogeneous IoT devices from different vendors is also introduced [47]. To ensure security and privacy of the elderlies an isolated 5G network slice is established for the connectivity of all the sensors and equipment and most importantly the caregivers' devices.

7.3.6 Ageing@home: A secure 5G welfare technology solution for elderlies

Contact: Thanh van Do (Telenor, OsloMet), Boning Feng (OsloMet), Birgitta Langhammer (OsloMet), Van Thuan Do (Wolffia AS), Niels Jacot (Wolffia AS), Bruno Dzogovic (OsloMet), Bernardo Santos (OsloMet), Per Jonny Nesse (Telenor),

The world population is ageing at a fast pace and to enable elderly to age at home can become a viable solution both economically and socially speaking, leading also to the overall improvement of the elderly's well-being and comfort. There are currently a few AAL (Ambient Assisted Living) systems which although operational are not yet optimal in terms of efficiency and security. This paper [48] proposes a welfare technology solution called Ageing@home which aims at enabling newly hospitalized elderlies to come home earlier by making use of a dedicated 5G network slice for health care system. Such an isolated logical network will provide adequate security, privacy and reliability for the selected welfare technologies and services deployed at the elderly home. The proposed solution allows the selection and customization of needed welfare technologies and services and promotes the re-allocation and re-use of equipment. Validation methods and a business plan have been presented as well as a thorough description of a proof-of-concept implementation.

7.3.7 Optimizing 5G VPN+ Transport Networks with Vector Packet Processing and FPGA Cryptographic Offloading

Contact: Bruno Dzogovic (OsloMet), Bernardo Santos (OsloMet), Boning Feng (OsloMet), Van Thuan Do (OsloMet), Niels Jacot (Wolffia AS), Thanh van Do (Telenor, OsloMet)

Network slicing is the crucial prerogative that allows end users and industries to thrive from 5G infrastructures, however, such a logical network component can deteriorate from security vulnerabilities that prevail within cloud environments and datacenters. The Quality of Experience in 5G is a metric that takes into consideration sets of factors, which play role in the definition of the end-to-end performance, which is indeed latency, packet processing, utilization of legacy protocols, old hardware, encryption, non-optimized network topologies, routing problems and multitude of other aspects.

This research [49] sheds light on the inherent networking stack performance issues that translate into 5G environments, in a use-case where encrypted VPN tunneling is used to secure the backhaul transport network between the 4G/5G cores and the frontend networks. Unfortunately, using a VPN will incur overhead which reduces the performance. To avoid this, Vector Packet Processing and FPGA Cryptographic Offloading are used and proven to be very efficient.

7.3.8 Advanced 5G Network Slicing Isolation Using Enhanced VPN+ for Healthcare Verticals

Contact: Bruno Dzogovic (OsloMet), Tariq Mahmood (OsloMet), Bernardo Santos (OsloMet), Boning Feng (OsloMet), Van Thuan Do (Wolffia AS), Niels Jacot (Wolffia AS), Thanh van Do (Telenor, OsloMet)

Alongside of supporting the human world, 5G aimed towards establishing an all-inclusive ecosystem for Internet of Things to sustain variety of industrial verticals such as e-health, smart home, smart city, etc. With the successful implementation of multitude of sites, it has come to realization that the traditional security approaches incorporated in the 4th generation networks (LTE) may not suffice to protect 5G users and industries from adversaries that develop more advanced attack vectors. This is mostly due to the vulnerabilities brought by softwareization[8] and virtualization of the network which compromise the isolation and protection of the 5G network slices essential for the support of IoT verticals. In this work [49], we propose a progressive approach to enhance the isolation of network slices by

[8] Softwareization of networks, clouds, and internet of things https://onlinelibrary.wiley.com/doi/pdf/10.1002/nem.1967

employing the Enhanced VPN+ technology[9]. Furthermore, we describe a method for tackling DDoS and flooding attacks on the communication between the 4G/5G core networks and the Cloud-Radio Access Network, as well as the radio frontend.

7.3.9 Anomaly Detection in Cellular IoT with Machine Learning

Contact: Bernardo Santos (OsloMet), Imran Qayyum Khan (OsloMet), Bruno Dzogovic (OsloMet), Boning Feng (OsloMet), Van Thuan Do (Wolffia AS), Niels Jacot (Wolffia AS), Thanh van Do (Telenor, OsloMet)

The number of Internet of Things (IoT) devices used in eldercare are increasing day by day and bringing big security challenges especially for healthcare organizations, IoT service providers and most seriously for the elderly users. Attackers launch many attacks using compromised IoT devices such as Distributed Denial of Services (DDoS), among others. To detect and prevent these types of attacks on IoT devices connected to the cellular network, it is essential to have a proper overview of the existing threats and vulnerabilities. The main objective of this work [50] is to present and compare different machine learning algorithms for anomaly detection in the cellular IoT scenario. Five supervised machine learning algorithms, namely KNN, Naïve Bayes, Decision Tree and Logistic Regression are used and evaluated by their performance. We see that, for both normal (using a local test dataset) and attack traffic (CICDDoS2019[10]) datasets, the accuracy and precision of the models are in average above 90%.

7.3.10 Applying k-vertex cardinality constraints on a Neo4j graph database

Contact: Martina Šestak, Marjan Heričko, Tatjana Welzer Družovec, Muhamed Turkanović (UM)

Graph databases are getting popular in security applications, since they offer additional functionalities compared to traditional approaches. As with any other database solution, graph databases also need to be able to implement business rules related to a given application domain. At the moment, aside from integrity constraints, there is a limited number of mechanisms for business rules implementation in Graph Database Management Systems (GDBMSs). The underlying property graph data model does not include any formal notation on how to represent different constraints. Specifically, this paper discusses the problem of representing cardinality constraints in graph databases. We introduce the novel concept of k-vertex cardinality constraints, which enable us to specify the minimum and maximum number of edges between a vertex and a subgraph [51]. We also propose an

[9] IETF TEAS Working Group: A framework for enhanced virtual private networks (VPN+) service. https://tools.ietf.org/html/draft-ietf-teas-enhanced-vpn-06

[10] CICDDoS2019 Dataset: https://www.unb.ca/cic/datasets/ddos-2019.html

approach, which includes the representation of cardinality constraints through the property graph data model, and demonstrate its implementation through a series of stored procedures in Neo4j GDBMS. The proposed approach is then evaluated by performing experiments on synthetic and real datasets to test the influence of checking cardinality constraints on query execution times (QETs) when adding new edges. Additionally, a comparison is performed on synthetic datasets with varying outgoing vertex degrees in order to gain an insight into how increasing the vertex degree affects QETs. In general, the results obtained for each test scenario show that the implemented k-vertex cardinality constraints model does not significantly affect QETs. Also, the results indicate that the model is dependent on the order of the underlying k-vertex cardinality constraints and outgoing vertex degree in the dataset.

7.3.11 Synergy of Blockchain Technology and Data Mining Techniques for Anomaly Detection

Contact: Martina Šestak, Aida Kamišalić (UM), Renata Kramberger (University of Zagreb), Iztok Fister, Jr (UM)

Blockchain and Data Mining are not simply buzzwords, but rather concepts that are playing an important role in the modern Information Technology (IT) revolution. Both technologies can be applied to detect suspicious behaviors or attacks. Blockchain has recently been popularized by the rise of cryptocurrencies, while data mining has already been present in IT for many decades. Data stored in a blockchain can also be considered to be big data, whereas data mining methods can be applied to extract knowledge hidden in the blockchain. In this work, we present a survey of approaches for mining blockchain data, as well as present several real-world applications of the identified approaches [52]. Special attention was paid to anomaly detection and fraud detection, which were identified as the most prolific applications of applying data mining methods on blockchain data. Our review analysed the current trends in exploiting the synergies of blockchain technology and data mining techniques for anomaly detection, while identifying relevant machine learning methods in order to build a taxonomy of those methods used to enhance blockchain technology for specific purposes.

7.3.12 Towards Automating Security Enhancement for Cloud Services

Contact: Mohamed Oulaaffart, Remi Badonnel, Olivier Festor (UL)

Cloud infrastructures provide new facilities (elasticity, load balancing, easy integration) to build and maintain elaborated services built from multiple resources in a flexible manner. The changes that continuously affect these services, in particular the migration of resources amongst such cloud infrastructures, induce configuration changes. These latter may generate configuration vulnerabilities, due to the alteration of a software version or the modification of configuration parameters,

that can compromise the confidentiality, integrity and availability of services. In that context, we have proposed a security framework for automatically supporting the migration of resources in cloud composite services [53]. It bridges the gap between the orchestration language that specifies the cloud composite services, and the vulnerability descriptions that define the configuration states that should be prevented for the migrated resource. We have specified the different building blocks of the framework and their interactions. They operate according to three main phases, namely the projection of the migrated resource, the assessment of the corresponding configuration based on vulnerability datasets, and the selection of adequate security functions when this configuration appears to be vulnerable. We have also formalized and evaluated to what extent the different phases can be supported by SMT[11] solving algorithms in order to both assess and correct vulnerable configurations. As future work, we are interested in evaluating complementary criteria for supporting the selection of security functions, including the different costs that are associated to their activation, such as the additional delays that they may introduce for the operation of the considered services. We are also planning to further investigate the extension of the TOSCA orchestration language with respect to such resource migrations in cloud composite services. In particular, we would like to evaluate its usage for supporting the interactions with multiple cloud providers based on a trusted third party.

7.3.13 Automated Orchestration of Security Chains Driven by Process Learning

Contact: Remi Badonnel (UL), Nicolas Schnepf (Aalborg University), Abdelkader Lahmadi (UL), Stephan Merz

Connected devices, such as smartphones and tablets, are exposed to a large variety of attacks. Their protection is often challenged by their resource constraints in terms of CPU, memory and energy. Security chains, composed of security functions such as firewalls, intrusion detection systems and data leakage prevention mechanisms, offer new perspectives to protect these devices using software-defined networking and network function virtualization. However, the complexity and dynamics of these chains require new automation techniques to orchestrate them. We have introduced an automated orchestration methodology for security chains in order to secure connected devices and their applications [54]. This methodology exploits process learning to establish behavioral models and infer security constraints represented as logical predicates. It then generates and merges a set of chains of security functions on the basis of these predicates. These chains are finally compiled into low-level configuration rules and deployed into the network, optimizing for the underlying topology. The performance of such a methodology combining machine learning and verification techniques have been evaluated by a set of experimental results, and have shown promising results for automating the deployment

[11] Satisfiability modulo theories

of security chains for specific classes of devices, such as Android phones. The flexibility of software-defined networking infrastructures enables synthesizing and deploying security chains that are specific to the networking behavior of individual applications running on smart devices. By construction, the obtained chains ensure certain correctness properties, and specific properties can be formally verified based on SMT solving and model checking. Finally, by applying appropriate optimization methods, the impact of deploying security chains on network performance can be substantially reduced.

7.3.14 Security Assurance for Modern Services

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

Cloud computing, edge and IoT have deeply changed how distributed systems are engineered, leading to the proliferation of ever-evolving and complex environments, where legacy systems, microservices, and nanoservices coexist. These services can severely impact on individuals' security and safety, introducing the need of solutions that properly assess and verify their correct behavior. Security assurance stands out as the way to address such pressing needs, with certification techniques being used to certify that a given target holds some non-functional properties. However, existing techniques fall short in providing a thorough evaluation, missing relevant aspects affecting the non-functional properties under certification.

In this context, we first worked on improving the current state of security assurance, developing a new multi-dimensional certification scheme. In this scheme, the evaluation of non-functional properties goes beyond the dimension on software artifacts and also considers additional dimensions (e.g., design and development) including relevant aspects that significantly contribute to property support (e.g., programming languages and development processes). Our multi-dimensional certification enables a new generation of service selection approaches capable to handle fine-grained, dimension-aware user's requirements. The performance and the quality of our approach are thoroughly evaluated in several experiments. This work has been submitted to IEEE Transactions on Services Computing (TSC), and is currently under review.

Following our research lines of the past years, we then worked on the integration of security assurance techniques within specific verticals. In [55] we defined a novel methodology integrating well-established risk management practices and assurance techniques. Our work is motivated by the fact that benefits in terms of new services and applications of modern distributed systems come at a price of new fundamental risks, and the need of adapting risk management frameworks to properly understand and address them. While research on risk management is an established practice that dates back to the 90s, many of the existing frameworks do not even come close to address the intrinsic complexity and heterogeneity of

modern systems. They rather target static environments and monolithic systems thus undermining their usefulness in real-world use cases. In [55], we presented an assurance-based risk management framework that addresses the requirements of risk management in modern distributed systems. The proposed framework implements a risk management process integrated with assurance techniques. Assurance techniques monitor the correct behavior of the target system, that is, the correct working of the mechanisms implemented by the organization to mitigate the risk. Flow networks compute risk mitigation and retrieve the residual risk for the organization. We evaluated the performance and quality of the framework in a simulated industry 4.0 scenario.

Finally, we extended our previous work, published at SECRYPT 2020, by defining an assurance framework that implements an assurance process evaluating the trustworthiness of hybrid systems [56]. The framework builds on a standard API-based interface supporting full and programmatic access to the functionalities of the framework. The process provides a transparent, non-invasive and automatic solution that does not interfere with the working of the target system. It builds on a Virtual Private Network (VPN)-based solution, to provide a smooth integration with target systems, in particular those mixing public and private clouds and corporate networks. We presented a detailed walkthrough of the process along with a performance evaluation of the framework in a simulated scenario.

7.3.15 Blockchain-Based Reputation Systems: Implementation Challenges and Mitigation

Contact: Claudio A. Ardagna (UMIL)

Reputation expresses the beliefs or opinions about someone or something that are held by an individual or by a community. Reputation Management Systems (RMSs) handle representation, computation, and storage of reputation in some quantitative form, suitable for grounding trust relations among parties. Quantifying reputation is important in scenarios involving interaction between parties who do not know (and potentially distrust) each other, such as online service provision. The basic idea is to let parties rate each other. When a party is considered for interaction, its ratings can be aggregated in order to derive a score for deciding whether to trust it or not. While much valuable research work has been done on reputation-based trust schemes, the problem of establishing collective trust in the reputation management system itself has never been fully solved. Recently, several researchers have put forward the idea of using Distributed Ledger Technology (DLT) as the foundation for implementing trustworthy RMSs. In [57], we identified some of the critical problems that arise when DLTs are used in order to manage evidence about previous interaction and compute reputations. The paper proposes some practical solutions and describes methods to deploy them on top of standard DLT of the Ethereum family.

7.3.16 Blockchain-based Execution of Collaborative Process

Contact: Barbara Carminati (UI), Elena Ferrari (UI)

During Y3, we continued our research to exploit blockchain to manage a secure execution of collaborative processes (workflow) among organizations. In addition to mechanisms to ensure secure data sharing during the deployment of collaborative processes, in Y3, we have also taken into consideration individual privacy. Indeed, following the "privacy by design" principle, also suggested by GDPR, we enhanced the proposed data-centric solution for secure process collaboration. This enhancement considers individuals' privacy preferences during the workflow execution. In this scenario, individuals express their preferences on how their data have to be consumed (e.g., the purpose of usage, retention time, etc.). During the workflow deployment, these preferences have to be verified against the organization's privacy policy involved in the collaborative process.

For this purpose, we have based our research on our proposal for blockchain-based privacy preference compliance [58]. The proposal exploits smart contracts to verify if the providers' privacy policies satisfy individuals' privacy preferences. However, to handle the privacy compliance during the workflow execution, we had to face further challenges. The most relevant one is due to the high number of compliance checks that have to be carried out during the workflow execution. This has required the definition of tailored optimization techniques to improve the compliance process.

To better understand the response of the blockchain and the optimization techniques, the experiments were conducted by varying the complexity of the policies, the number of tasks of a single execution, the privacy preference coverage (i.e. the percentage of privacy preferences present with respect to data tuple) and the selectivity (i.e. the percentage of privacy preferences that pass the compliance check). By varying these parameters, we obtained measures on the execution time of the privacy enforcement smart contract, the throughput of the blockchain in transactions per second, the number of privacy preferences processed per second, the time overhead used to create the optimized data structure, and the memory cost. Each experiment was compared with the naive implementation to validate the results obtained, achieving clear improvements. Execution times improve, as an example, we halve the time of privacy preferences compliance checks, going from 15000 pp/s in the naive case, to 30000 pp/s with the proposed solution.

8 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 at developing techniques for Privacy-Preserving Machine Learning modeling, as well as Personal Identifiable Information (PII) leakage detection to the advertising ecosystem and the general Web.
- **Identity Management:** The objective of this task is the development of 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 fake news identification in Online Social Networks as well as developing blockchain based methods for suppression of fake news.

8.1 Overview

During the third year of the project, the research activities of the partners resulted to eighteen publications, one of which received Best Paper award [59]. In the following sections, we present a short summary of the research activities related to Task 1.5.

Privacy

Section 8.3.1 presents YourAdvalue, a privacy-preserving tool for displaying to end-users their advertising value as seen through the Real Time Bidding (RTB) protocol. Section 8.3.2 investigates whether websites use sophisticated tracking techniques in order to track users who have already rejected cookies. The study shows that more than 75% of the tracking activities happened before the users decide whether they will accept or not cookies. The work in Section 8.3.3 presents the results of the first comprehensive analysis on the Indian online news media – with respect to tracking and partisanship – based on a dataset of 103 online news websites. Intense web tracking is verified; this has implications on the overall privacy of users visiting partisan news websites in India. Finally, a novel method is proposed for automatic extraction and categorization of Indian news topical subpages based on their URLs details. In Section 8.3.4, a Privacy-preserving Federated Learning (PPFL) framework for mobile systems is proposed. It is the first practical framework that fully prevents private information leakage at both server and client-side in Federated Learning scenarios. Section 8.3.5 focuses on the federation policies on the Decentralized Web (DW) and the negative impact that may have on users. Section 8.3.6 discusses a novel attack vector that misuses the mobile

advertising ecosystem for delivering sophisticated and stealthy attacks that leverage mobile sensors. Section 8.3.7 discusses the problem of knowledge gap between security experts and data scientists and proposes a solution approach based on the concept of Model-Based Big Data Analytics-as-a-Service (MBDAaaS) that helps security experts in preparing and deploying a data analytics that addresses their requirements. Section 8.3.8 focuses on the privacy-protection problem when publishing knowledge graphs (KGs). For this purpose, the Personalized k-Attribute Degree Principle is proposed in order to protect users' identities in published KGs. Moreover, towards protecting users' sensitive values in anonymized KGs, the (k, l)-Sequence Attribute Degree Principle is introduced along with a novel anonymization algorithm. Section 8.3.9 investigates a risk estimation approach based on apps' static analysis in order to quantify the privacy implications of installing an app. The goal is to determine how much data usage of personal data of a target app diverges from that of apps with the same purpose and this is in turn used to determine the app privacy risk. Finally, in Section 8.3.10 the first performance evaluation of PayString is discussed. PayString is an initiative to make payment identifiers global and human-readable, facilitating the exchange of payment information.

Identity management

Section 8.3.11 presents ITrade, a Blockchain-based secure and scalable IoT data marketplace. ITrade implements the decentralized management of data trading via Smart Contracts (SC) and it brings together individuals, companies, and organizations from private and public sectors, who are interested in IoT-collected data. Section 8.3.12 focuses on the Remote Electronic Voting (REV), which enables vote casting in an uncontrolled environment and vote transmission over communication channels. Section 8.3.13 introduces the DeTi – a Decentralized Ticketing platform which operates through Smart Contracts of Ethereum – for managing the distribution of electronic event tickets and regulating the aftermarket. In Section 8.3.14, the key characteristics of Distributed Storage Systems (DSS) are discussed together with their contributions in decentralized data storage. Section 8.3.15 introduces DIMANDS2, an Identity Association and Discovery system. DIMANDS2 proposes an innovative framework which is based on Decentralized Identifiers (DIDs) and blockchains and it no longer requires a large scale peer-to-peer network infrastructure to operate. Finally, Section 8.3.16 focuses on Hardware Security Module–based wallet methodology and presents an HSM-based working prototype that secures the entire life cycle of Ethereum public and private keys.

Social Networks and Fake News

Section 8.3.17 presents the *meta-graph*, that is, a new approach on detecting disinformation cascades on Twitter. *Meta-graph* is graph data structure that combines underlying users' relational event information, as well as semantic and topi-

cal modeling and it is able to identify disinformation at cascade level, using graph neural network algorithms. A series of studies have been focused on YouTube and Twitter. Specifically, Section 8.3.18 focuses on the role of YouTube's recommendation algorithm on promoting misinformation and conspiracy theories. The study characterizes and detects pseudoscientific misinformation videos related to the Flat Earth theory, as well as the anti-vaccination and anti-mask movements. Section 8.3.19 investigates the Incel community on YouTube and in particular the evolution of this community over the last decade. Incel ideology is associated with misogyny and anti-feminist viewpoints. The study investigates whether the YouTube recommendation algorithm steers users towards Incel-related videos. Section 8.3.20 examines the usefulness of video sharing on social media as a proxy of the amount of time Internet users spend at home. In particular, it focuses on the number of people sharing YouTube videos on Twitter before and during COVID19 lockdown measures were imposed by 109 countries. Section 8.3.21 presents the results of an extended investigation of the fake news websites. The study focuses on the online presence of fake news websites comparing their behavior with the real news websites. In addition, a content-agnostic Machine Learning classifier for automatic detection of fake news websites is proposed. Section 8.3.22 investigates aggression propagation on Twitter. First, a model of propagation of aggression is presented using opinion dynamics. Secondly, various methods have been proposed based on two well-known diffusion models, Independent Cascade (IC) and Linear Threshold (LT) in order to study the aggression evolution in the social networks. Section 8.3.23 presents the findings of a large study on Twitter and YouTube regarding the 2020 US Presidential elections. Specifically, the analysis has been performed on 19.8M tweets along with 28K Youtube links. The study focuses on the connection between the two social networks. The evolution of retweet graph has been studied along with a sentiment analysis on the YouTube links that were embedded in the tweets. Finally, Section 8.3.24 presents a novel system for identifying Twitter bots based on labeled Twitter data. For this purpose, a supervised Machine Learning framework is proposed based on the Extreme Gradient Boosting (XGBoost) algorithm. Section 8.3.25 investigates the gender differences in privacy awareness on online social networks. A comparative analysis has been conducted based on gender, age and number of friends of two Facebook accounts that created for the purpose of this study.

8.2 Link between T1.5 work and CONCORDIA pilots

Strong collaboration has been developed between academia and industry; the research performed within Task 1.5 is linked to the CONCORDIA pilots as follows (see Figure 16):

- Several research activities of TID, FORTH and ICL are related to pilot T2.1. Some of them are the product of collaboration between the three partners

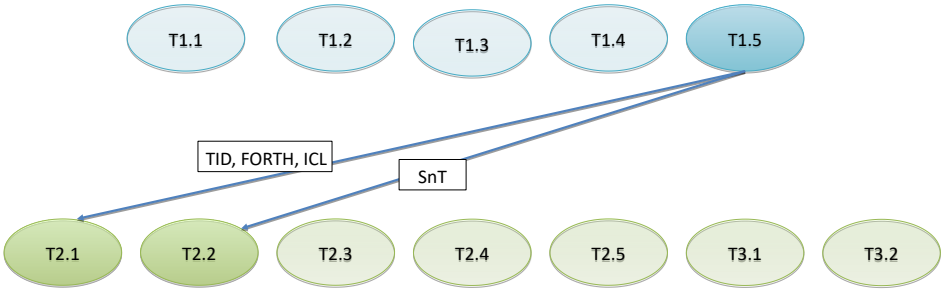

Figure 16: Links between T1.5 and the CONCORDIA pilots

which resulted to numerous publications [60, 61, 62, 59, 63, 64, 65, 66, 67, 68]. Specifically:

  - TID and FORTH investigated: (i) the advertising price dynamics (Section 8.3.1); (ii) the online tracking ecosystem (Sections 8.3.2, 8.3.3); and (iii) traffic analysis of the fake-news sites (Section 8.3.21).
  - TID worked together with ICL on the privacy-preserving Federated Learning framework (Section 8.3.4).
  - TID investigated the media sharing behavior on YouTube, the content moderation in the Decentralised Web and the aggression propagation on social media (Sections 8.3.5, 8.3.20, 8.3.22).

- The research work of SnT is related to pilot T2.2 (Sections 8.3.10, 8.3.16). Specifically, the Concordance[12] data sharing solution has undergone numerous improvements through collaboration with CaixaBank. The project has just seen its very first international test using example customer profiles and documents by initiating the process from Luxembourg with CaixaBank in Spain through a server in the Netherlands. The test saw a request from the bank to the 'customer' to obtain permission to access the data, followed by international data checking and exchange using a blockchain. Work is now underway to integrate an Identity Verification tool developed by CUT.

8.3 Summaries of T1.5 Research Activities

In the next sections we present summaries of T1.5-related research efforts.

[12] https://www.concordanceltd.co.uk/

8.3.1 YourAdvalue: Measuring Advertising Price Dynamics without Bankrupting User Privacy

Contact: Michalis Pachilakis (FORTH), Panagiotis Papadopoulos (TID), Nikolaos Laoutaris (IMDEA), Evangelos P. Markatos (FORTH), Nicolas Kourtellis (TID)

The Real Time Bidding (RTB) protocol is by now more than a decade old. During this time, a handful of measurement papers have looked at bidding strategies, personal information flow, and cost of display advertising through RTB. In [60], we present YourAdvalue, a privacy-preserving tool for displaying to end-users in a simple and intuitive manner their advertising value as seen through RTB. Using YourAdvalue, we measure desktop RTB prices in the wild, and compare them with desktop and mobile RTB prices reported by past work. We present how it estimates ad prices that are encrypted, and how it preserves user privacy while reporting results back to a data-server for analysis. We deployed our system, disseminated its browser extension[13], and collected data from 200 users, including 12000 ad impressions over 11 months.

By analyzing this dataset, we show that desktop RTB prices have grown 4.6× over desktop RTB prices measured in 2013, and 3.8X over mobile RTB prices measured in 2015. We also study how user demographics associate with the intensity of RTB ecosystem tracking, leading to higher ad prices. We find that exchanging data between advertisers and/or data brokers through cookie-synchronization increases the median value of displayed ads by 19%. We also find that female and younger users are more targeted, suffering more tracking (via cookie synchronization) than male or elder users. As a result of this targeting in our dataset, the advertising value (i) of women is 2.4X higher than that of men, (ii) of 25-34 year-olds is 2.5X higher than that of 35-44 year-olds, (iii) is most expensive on weekends and early mornings.

8.3.2 User Tracking in the Post-cookie Era: How Websites Bypass GDPR Consent to Track Users

Contact: Emmanouil Papadogiannakis (FORTH), Panagiotis Papadopoulos (TID), Nicolas Kourtellis (TID), Evangelos P Markatos (FORTH)

During the past few years, mostly as a result of the GDPR and the CCPA, websites have started to present users with cookie consent banners. These banners are web forms where the users can state their preference and declare which cookies they would like to accept, if such option exists. Although requesting consent before storing any identifiable information is a good start towards respecting the user privacy, yet previous research has shown that websites do not always respect user

[13] https://chrome.google.com/webstore/detail/youradvalue/apipekdgpjmbooidaohhecfpaecahaap

choices. Furthermore, considering the ever decreasing reliance of trackers on cookies and actions browser vendors take by blocking or restricting third-party cookies, we anticipate a world where stateless tracking emerges, either because trackers or websites do not use cookies, or because users simply refuse to accept any.

In [61], we explore whether websites use more persistent and sophisticated forms of tracking in order to track users who said they do not want cookies. Such forms of tracking include first-party ID leaking, ID synchronization, and browser fingerprinting. Our results suggest that websites do use such modern forms of tracking even before users had the opportunity to register their choice with respect to cookies. To add insult to injury, when users choose to raise their voice and reject all cookies, user tracking only intensifies. As a result, users' choices play very little role with respect to tracking: we measured that more than 75% of tracking activities happened before users had the opportunity to make a selection in the cookie consent banner, or when users chose to reject all cookies. We make available our code and data for future research on the topic[14].

8.3.3 User Tracking on Indian News Media Websites

Contact: Yash Vekaria (LMNIIT), Vibhor Agarwal (LMNIIT), Pushkal Agarwal (KCL), Sangeeta Mahapatra (GIGA), Shounak Set (KCL), Sakthi Balan Muthiah (LMNIIT), Nishanth Sastry (USurrey), Nicolas Kourtellis (TID)

Online user privacy and tracking have been extensively studied in recent years, especially due to privacy and personal data-related legislations in the EU and the USA, such as the General Data Protection Regulation (GDPR), ePrivacy Regulation, and California Consumer Privacy Act (CCPAA). Meanwhile, in the Asian world, India is experiencing intense political partisanship and sectarian divisions. This paper [65] performs, to the best of our knowledge, the first comprehensive analysis on the Indian online news media with respect to tracking and partisanship. We build a dataset of 103 online, mostly mainstream news websites. With the help of two experts, alongside data from the Media Ownership Monitor of the Reporters without Borders, we label these websites according to their partisanship (Left, Right, or Centre). We study and compare user tracking on these sites with different metrics: numbers of cookies, cookie synchronizations, device fingerprinting, and invisible pixel-based tracking. We find that Left and Centre websites serve more cookies than Right-leaning websites. However, through cookie synchronization, more user IDs are synchronized in Left websites than Right or Centre. Canvas fingerprinting is used similarly by Left and Right, and less by Centre. Invisible pixel-based tracking is 50% more intense in Centre-leaning websites than Right, and 25% more than Left. Desktop versions of news websites deliver more cookies than their mobile counterparts. A handful of third-parties are tracking users in most

[14] https://gitlab.com/papamano/consent-guard

websites in this study. This paper, by demonstrating intense web tracking, has implications for research on overall privacy of users visiting partisan news websites in India. We make available our code and data for future research on the topic[15].

Going beyond this first step, we take a look at the tracking ecosystem beyond the home page of websites. This next work was motivated by the fact that recent research has revealed novel tracking and personal identifiable information leakage methods that first- and third-parties employ on websites around the world, as well as the intensity of tracking performed on such websites. However, for the sake of scaling to cover a large portion of the Web, most past studies focused on homepages of websites, and did not look deeper into the tracking practices on their topical subpages. The majority of studies focused on the Global North markets such as the EU and the USA. Large markets such as India, which covers 20% of the world population and has no explicit privacy laws, have not been studied in this regard.

In this follow-up work [64], we aim to address these gaps and focus on the following research questions: Is tracking on topical subpages of Indian news websites different from their homepage? Do third-party trackers prefer to track specific topics? How does this preference compare to the similarity of content shown on these topical subpages? To answer these questions, we propose a novel method for automatic extraction and categorization of Indian news topical subpages based on the details in their URLs. We study the identified topical subpages and compare them with their homepages with respect to the intensity of cookie injection and third-party embeddedness and type. We find differential user tracking among subpages, and between subpages and homepages. We also find a preferential attachment of third-party trackers to specific topics. Also, embedded third-parties tend to track specific subpages simultaneously, revealing possible user profiling in action. We make available our code and data for future research on the topic[16].

8.3.4 PPFL: Privacy-preserving Federated Learning with Trusted Execution Environments

Contact: Fan Mo (ICL), Hamed Haddadi (ICL), Kleomenis Katevas (TID), Eduard Marin (TID), Diego Perino (TID), Nicolas Kourtellis (TID)

We propose and implement a Privacy-preserving Federated Learning (PPFL) framework for mobile systems to limit privacy leakages in federated learning [59]. Leveraging the widespread presence of Trusted Execution Environments (TEEs) in high-end and mobile devices, we utilize TEEs on clients for local training, and on servers for secure aggregation, so that model/gradient updates are hidden from adversaries. Challenged by the limited memory size of current TEEs, we leverage greedy layer-wise training to train each model's layer inside the trusted area until

[15] http://tiny.cc/india-tracking
[16] http://tiny.cc/india-topic

its convergence. The performance evaluation of our implementation shows that PPFL can significantly improve privacy while incurring small system overheads at the client-side. In particular, PPFL can successfully defend the trained model against data reconstruction, property inference, and membership inference attacks. Furthermore, it can achieve comparable model utility with fewer communication rounds (0.54×) and a similar amount of network traffic (1.002×) compared to the standard federated learning of a complete model. This is achieved while only introducing up to ∼15% CPU time, ∼18% memory usage, and ∼21% energy consumption overhead in PPFL's client-side.

This work has been published in the top conference ACM MobiSys 2021 and received best paper award for its contributions. We make our code available open source for future research on the topic[17].

8.3.5 Exploring Content Moderation in the Decentralised Web: The Pleroma Case

Contact: Anaobi Ishaku Hassan (QMUL), Aravindh Ramanu (TID), Ignacio Castro (QMUL), Haris Bin Zia (QMUL), Emiliano De Cristofaro (UCL) Nishanth Sastry (USurrey), Gareth Tyson (QMUL)

Decentralizing the Web, i.e., creating a Web where it has all services and functionalities as we know them now, but without relying on centralized big operators (that can be monitored, censored or even shut down at a click of a button)[18] is a desirable but challenging goal. One particular challenge is achieving decentralized content moderation in the face of various adversaries (e.g., trolls). To overcome this challenge, many Decentralized Web (DW) implementations rely on federation policies. Administrators use these policies to create rules that ban or modify content that matches specific rules. This, however, can have unintended consequences for many users.

In [68], we present the first study of federation policies on the DW, their in-the-wild usage, and their impact on users. We identify how these policies may negatively impact "innocent" users and outline possible solutions to avoid this problem in the future.

8.3.6 Misusing Mobile Sensors for Stealthy Data Exfiltration

Contact: Michalis Diamantaris (FORTH), Sotiris Ioannidis (FORTH)

In [69] we introduce a novel attack vector that misuses the mobile advertising ecosystem for delivering sophisticated and stealthy attacks that leverage mobile

[17] https://github.com/mofanv/PPFL.

[18] https://amp.theguardian.com/technology/2018/sep/08/decentralisation-next-big-step-for-the-world-wide-web-dweb-data-internet-censorship-brewster-kahle

sensors. These attacks do not depend on any special app permissions or specific user actions, and affect all Android apps that contain in-app advertisements due to the improper access control of sensor data in *WebView*. We outline how motion sensor data can be used to infer users' sensitive touch input (e.g., credit card information) in two distinct attack scenarios, namely intra-app and inter-app data exfiltration. While the former targets the app displaying the ad, the latter affects every other Android app running on the device.

8.3.7 Aiding Security Experts in Big Data Scenarios

Contact: Claudio A. Ardagna

One of the most challenging scenarios is the integration of security within specific environments, due to the mismatch between domain and security experts. This is particularly true in the Big Data world, where fully exploiting data through advanced analytics, machine learning and artificial intelligence becomes crucial for businesses. From micro to large enterprises, it results in a key advantage (or shortcoming) in the global market competition, as well as in a strong market driver for business analytics solutions. This scenario is deeply changing the security landscape, introducing new risks and threats that affect security of systems, on one side, and safety and privacy of users, on the other side, up to the point that malfunctions may endanger users' lives. Many domains that can benefit from novel solutions based on data analytics have stringent security requirements to fulfill. The energy domain's Smart Grid is a major example of systems at the crossroads of security and data-driven intelligence. It plays a crucial role in modern energy infrastructure. However, it must face two major challenges related to security: managing front-end intelligent devices such as power assets and smart meters securely, and protecting the huge amount of data received from these devices. Starting from these considerations, setting up proper analytics is a complex problem because security controls could have the undesired side effect of decreasing the accuracy of the analytics themselves. This is even more critical when the configuration of security controls is let to the security expert, who often has only basic skills in data science. In [70], we proposed a solution based on the concept of Model-Based Big Data Analytics-as-a-Service (MBDAaaS) that bridges the gap between security experts and data scientists. Our solution acts as a middleware allowing a security expert and a data scientist to collaborate to the deployment of an analytics addressing their needs.

8.3.8 Anonymization of Knowledge Graphs

Contact: Barbara Carminati (UI) Elena Ferrari (UI) Anh-Tu Hoang (UI)

During Y3, we have extended our research on protecting privacy when publishing knowledge graphs (KGs). In general, these techniques propose to anonymize the whole data with the same anonymization level (aka the same k). However, in a

context where data are collected from different users, it is crucial to consider also their preferences on the anonymization level to adopt for their data. To cope with this concern, we have presented the Personalized k-Attribute Degree Principle, a personalized version of our previous protection principle [71] for protecting users' identities in published KGs. The new principle allows every user u to specify his/her own k value (i.e., k_u) and ensures that adversaries cannot re-identify user u in an anonymized KG with a confidence higher than $\frac{1}{k_u}$, where k_u is an integer specified by user u. In addition, we proposed a new anonymization algorithm generating anonymized KGs satisfying this principle. The experimental results of this paper have been presented in a paper under submission.

Furthermore, we have investigated protecting users' sensitive values in anonymized KGs while allowing data providers to publish many versions of their KGs sequentially. To this end, we have introduced a new protection principle, namely (k, l)-Sequence Attribute Degree Principle, an extension of the k^w-Time-Varying Attribute Degree Principle [72]. The new principle allows data providers to specify the sensitive attribute (e.g., $disease$) to be protected. Then, adversaries cannot infer the sensitive value of any user with a confidence higher than $\frac{1}{l}$ even if they monitor all published KGs. We also introduced a new anonymization algorithm ensuring that the sequence of its generated anonymized KGs satisfies the proposed principle. We evaluated this work in real-life KGs and these experimental results have been described in a paper under submission.

8.3.9 A Risk Assessment Mechanism for Android Apps

Contact: Barbara Carminati (UI), Elena Ferrari (UI)

Mobile apps have become an integral part of our daily lives in that they can be used for accessing a variety of services everywhere. However, despite the many benefits, there are also risks related to the usage of personal information. Understanding the privacy implications of installing an app could be very difficult, especially for non-skilled users. To cope with this issue, during Y3, we investigated a risk estimation approach based on apps' static analysis [73]. This aims to determine how much data usage of personal data of a target app diverges from that of apps with the same purpose and this is in turn used to determine the app privacy risk. To prove that the proposed risk estimation measure is effective, we run several experiments involving different groups of participants, obtaining promising results.

8.3.10 Privacy-Preserving PayString Service

Contact: Wazen Shbair (SnT),

PayString is an initiative to make payment identifiers global and human-readable, facilitating the exchange of payment information. However, the reference implementation lacks privacy and security features, making it possible for anyone to access the payment information as long as the PayString identifier is known. This

paper presents the first performance evaluation of PayString. Via a large-scale testbed, our experimental results show an overhead which, given the privacy and security advantages offered, is acceptable in practice, thus making the proposed solution feasible [74].

8.3.11 ITrade: A Blockchain-based, Self-Sovereign, and Scalable Marketplace for IoT Data Streams

Contact: Sina Rafati Niya (UZH), Burkhard Stiller (UZH)

In recent years, the interest grew in the Internet-of-Things (IoT) and Blockchain (BC) integration for additional trust and decentralization. This opened potentials in various use cases, such as supply chain tracking, smart cities, and recently IoT data marketplaces. In order to make IoT data more accessible, IoT data marketplaces have been proposed, however as our current analysis shows, these marketplaces are not meeting user demands with respect to privacy, scalability, and data sovereignty aspects. Therefore, in [75] we present the design, implementation, and evaluation of ITrade as a secure and scalable IoT data marketplace based on BCs [75].

ITrade implements the decentralized management of data trading via Smart Contracts (SC). As a data streaming platform, it brings together individuals, companies, and organizations from private and public sectors, which are interested in IoT-collected data. For instance, (a) as within studies performed for health-related use cases, which are most needed in the case of COVID-19 pandemic, and (b) individuals, who are interested in selling data collected by their devices in general. IoT owners can initiate data streaming via ITrade and benefit from its highly scalable architecture. ITrade proposes a highly scalable microservice-based architecture based on clouds. ITrade enables end-to-end data streaming from IoT devices toward data buyers. The Smart Contract (SC)-oriented design of ITrade enables decentralized management of autonomous and distributed IoT data trading. ITrade evaluations attest its scalability as a reliable peer-to-peer data transmission platform.

8.3.12 From Centralized to Decentralized Remote Electronic Voting

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

Elections generally involve the simple tasks of counting votes and publishing the final tally to voters. Depending on the election's scope, these processes require sophisticated methods embedded in the electorate's various technological and societal factors (*e.g.,* the voting culture). An election's integrity is the pinnacle of the trust placed in the voting process and its final results. Previous research on cryptographic voting schemes continuously refined voting protocols to achieve private and verifiable elections. These cryptographic schemes enable novel ways to allow remote voting systems to (1) provide a remote voting alternative to on-site voting

with a ballot box and (2) offer a technical path to build an end-to-end verifiable electronic voting system by applying cryptographic primitives [76].

The work has been submitted as a book chapter [76], and after a review iteration is currently under submission with the editor Prof. Dr. Salil Kanhere. The book chapter [76] will be included in a book about different Blockchain Applications and will be published in 2022. The chapter explicitly addresses Remote Electronic Voting (REV), which enables vote casting in an uncontrolled environment and vote transmission over communication channels (*e.g.,* over the Internet).

8.3.13 DeTi: A Decentralized Ticketing Management Platform

Contact: Sina Rafati Niya (UZH), Burkhard Stiller (UZH)

Event tickets being sold in their electronic instances are subject to counterfeiting, profiteering, and black markets. Therefore, suitable service management mechanisms are required to remedy these problems. This work [77] designs, develops, and evaluates the approach of an open and Decentralized Ticketing platform called *DeTi* for managing the distribution of electronic event tickets and regulating the aftermarket. DeTi operates through Smart Contracts of Ethereum, and users can verify tickets' validity for a given event. By additionally securing technically the set of underlying application processes, DeTi obviates forging, replicating, and scalping of tickets, allowing for a managed resale ecosystem of tickets based on and limited to organizers' initial pricing. Especially, a new mechanism for users to detect fraudulent events is introduced, too. The evaluation performed indicates that DeTi invalidates or validates tickets efficiently via its decentralized and BC-based service management approach.

8.3.14 Enabling Technologies and Distributed Storage

Contact: Sina Rafati Niya (UZH), Burkhard Stiller (UZH)

Blockchains (BC) persist time-stamped blocks of transactions in backward-linked lists, cryptographically chained. BCs are created and maintained within a network of distributed nodes, which are established on the underlying overlay network for applying Peer-to-peer (P2P) communications to store data. Advantages of BCs include immutability and trust without central control. They benefit various application areas, *e.g.,* Decentralized Finance, Internet-of-Things-integrated Smart City monitoring, and Supply Chain Tracking. However, higher data storage costs in BCs hinder the scalability of BC-based applications and, thus, determine an incentive for the development of different Distributed Storage Systems (DSS) motivated by key BC concepts, such as data integrity and tamper-proofness. This work [78] resembles the key characteristics of DSSes that need to be considered by system designers. By putting the users' needs at the center, this work summarizes BCs' contributions to decentralized data storage. Moreover, primary use cases and challenges experienced with the DSS are categorized.

8.3.15 Identity Management through a global Discovery System based on Decentralized Identifiers

Contact: Konstantinos Lampropoulos (UP), Nikos Kyriakoulis (UP), Spyros Denazis (UP)

The different definitions of "identity" and "identity management (IdM)", reveal an underlying issue between these terms in the digital world. The first one defines identity as "the distinguishing character or personality of an individual" while the second one describes the identity management as "a framework of policies and technologies for ensuring that the right users (in an enterprise) have the appropriate access to technology resources. Identities today continue to be a company resource instead of a representation of the person behind it. This results in a continuous proliferation of users' online identities and identity related data, which reside scattered across different domains and are only meaningful within the service context they are created and used. Across different contexts, the identities of the same user appear as unassociated information that cannot be easily managed or shared without significant privacy and security risks. At the same time, the continuous growth of digital services and domains aggravates the "Identity Management problem" (diverse administration models, rigid processes for provisioning and deprovisioning identities, user password fatigue etc.) at a fast pace. This work describes DIMANDS2, an Identity Association and Discovery system which takes advantage of the work on Decentralized Identifiers (DIDs)[19] and Verifiable Credentials (VCs)[20] to a) organize users' identity information which until now reside scattered across multiple isolated contexts (providers, domains, networks etc.) and b) allow service providers exchange/share identity related information in a privacy-enabled manner. DIMANDS2 does not require any authority to overview and maintain its infrastructure and provides to the end user full control over any activity related to its identity data. It is "format-agnostic" in the sense that it can accommodate any type of identifier (existing or new), without requiring from existing providers to change their services and technologies to adopt another new identifier. This paper has been submitted for review.

8.3.16 HSM-based Key Management Solution for Ethereum Blockchain

Contact: Wazen Shbair (SnT)

The security of distributed applications backed by blockchain technology relies mainly on keeping the associated cryptographic keys (i.e. private keys) in well-protected storage. Since they are the unique proof of ownership of the underlying digital assets. If the keys are stolen or lost, there is no way to recover the assets. The cold wallet is a good candidate for basic use cases, but it has a substantial challenge for more complex applications as it does not scale. Warm and hot wallets

[19] https://www.w3.org/TR/did-core/
[20] https://www.w3.org/TR/vc-data-model/

are more convenient options for blockchain-based solutions that aim to transact in a cloud environment. In this work, we focus on Hardware Security Module (HSM) based wallet. The HSM is the de-facto standard device designed to manage high-value cryptographic keys and to protect them against hacks. In this demonstration, we present an HSM-based working prototype that secures the entire life cycle of Ethereum public and private keys [79].

This work is related to this task since digital identities mechanism are built around the concept of public and private keys. Therefore, it's always the question of how to manage the life-cycle of the key. Herby, we propose using the HSM for key management in a secure and trusted environment.

8.3.17 Detecting Disinformation Cascades on Twitter

Contact: Nikos Salamanos (CUT), Michael Sirivianos (CUT)

As recent events have demonstrated, disinformation spread through social networks can have dire political, economic and social consequences. Detecting disinformation must inevitably rely on the structure of the network, on users particularities and on event occurrence patterns. In [80], we present a graph data structure, which we denote as a *meta-graph*, that combines underlying users' relational event information, as well as semantic and topical modeling. We detail the construction of an example meta-graph using Twitter data covering the 2016 US election campaign. Then, we compare the detection of disinformation at cascade level – using well-known graph neural network algorithms – with the detection accuracy when the same algorithms applied on the meta-graph nodes, only. The comparison shows a consistent 3%–4% improvement in accuracy when using the meta-graph, over all considered algorithms, compared to basic cascade classification, and a further 1% increase when topic modeling and sentiment analysis are considered. We carry out the same experiment on two other datasets, HealthRelease and HealthStory, part of the FakeHealth dataset repository, with consistent results. Finally, we discuss further advantages of our approach, such as the ability to augment the graph structure using external data sources, the ease with which multiple meta-graphs can be combined as well as a comparison of our method to other graph-based disinformation detection frameworks.

8.3.18 Assessing the Effect of Watch History on YouTube's Pseudoscientific Video Recommendations

Contact: Nikos Salamanos (CUT), Michael Sirivianos (CUT)

The role played by YouTube's recommendation algorithm in unwittingly promoting misinformation and conspiracy theories is not entirely understood. Yet, this can have dire real-world consequences, especially when pseudoscientific content is promoted to users at critical times, such as the COVID–19 pandemic. In [81], we set out to characterize and detect pseudoscientific misinformation on YouTube.

We collect 6.6K videos related to COVID-19, the Flat Earth theory, as well as the anti-vaccination and anti-mask movements. Using crowdsourcing, we annotate them as pseudoscience, legitimate science, or irrelevant and train a deep learning classifier to detect pseudoscientific videos with an accuracy of 0.79. We quantify user exposure to this content on various parts of the platform and how this exposure changes based on the user's watch history. We find that YouTube suggests more pseudoscientific content regarding traditional pseudoscientific topics (e.g., flat earth, anti–vaccination) than for emerging ones (like COVID–19). At the same time, these recommendations are more common on the search results page than on a user's homepage or in the recommendation section when actively watching videos. Moreover, we shed light on how a user's watch history substantially affects the type of recommended videos. Finally, we make available our code and data for future research on the topic[21].

8.3.19 Understanding the Incel Community on YouTube

Contact: Nikos Salamanos (CUT), Michael Sirivianos (CUT)

YouTube is by far the largest host of user-generated video content worldwide. Alas, the platform has also come under fire for hosting inappropriate, toxic, and hateful content. One community that has often been linked to sharing and publishing hateful and misogynistic content are the Involuntary Celibates (Incels), a loosely defined movement ostensibly focusing on men's issues. In [82], we set out to analyze the Incel community on YouTube by focusing on this community's evolution over the last decade and understanding whether YouTube's recommendation algorithm steers users towards Incel-related videos. We collect videos shared on Incel communities within Reddit and perform a data-driven characterization of the content posted on YouTube. Among other things, we find that the Incel community on YouTube is getting traction and that, during the last decade, the number of Incel–related videos and comments rose substantially. We also find that users have a 6.3% chance of being suggested an Incel–related video by YouTube's recommendation algorithm within five hops when starting from a non Incel–related video. Overall, our findings paint an alarming picture of online radicalization: not only Incel activity is increasing over time, but platforms may also play an active role in steering users towards such extreme content.

8.3.20 YouTubing at Home: Media Sharing Behavior Change as Proxy for Mobility Around COVID-19 Lockdowns

Contact: Yelena Mejova (ISI-Torino) Nicolas Kourtellis (TID)

Compliance with public health measures, such as restrictions on movement and socialization, is paramount in limiting the spread of diseases such as the severe

[21] https://github.com/kostantinos-papadamou/pseudoscience-paper

acute respiratory syndrome coronavirus 2 (also referred to as COVID19). Although large population datasets, such as phone-based mobility data, may provide some glimpse into such compliance, they are often proprietary, and may not be available for all locales.

In [63], we examine the usefulness of video sharing on social media as a proxy of the amount of time Internet users spend at home. In particular, we focus on the number of people sharing YouTube videos on Twitter before and during COVID19 lockdown measures were imposed by 100+ countries around the world. Concretely, we use a dataset of 390 million Twitter posts mentioning YouTube videos spanning the period from July 2019 to September 2020. After geo-locating the users of these posts, we select 109 countries which have had a COVID-related social distancing recommendation or mandate. We find that the media sharing behavior differs widely between countries, in some having immediate response to the lockdown decrees – mostly by increasing the sharing volume dramatically – while in others having a substantial lag. Indeed, the posting trend shifts dramatically around government measures, although the precise timing may differ, depending on the peculiarities of each government's handling of the pandemic.

We confirm that these insights correlate strongly with mobility, as measured using phone data. In fact, these changes in media sharing have a strong negative relationship with the phone-based mobility data, suggesting social media may provide valuable insights, especially where such mobility data are not available due to collection sparsity, or even privacy regulations. Finally, we illustrate that both media sharing and mobility behaviors change more drastically around mandated lockdowns, and less so around more lax recommendations. We make the media sharing volume data available to the research community for continued monitoring of behavior change around public health measures. [22]

8.3.21 The Rise and Fall of Fake News sites: A Traffic Analysis

Contact: Manolis Chalkiadakis (FORTH), Alexandros Kornilakis (FORTH), Panagiotis Papadopoulos (TID), Evangelos P. Markatos (FORTH), Nicolas Kourtellis (TID)

Over the past decade, we have witnessed the rise of misinformation on the Internet, with online users constantly falling victims of fake news. A multitude of past studies have analyzed fake news diffusion mechanics and detection and mitigation techniques. However, there are still open questions about their operational behavior such as: How old are fake news websites? Do they typically stay online for long periods of time? Do such websites synchronize with each other their up and down time? Do they share similar content through time? Which third-parties support their operations? How much user traffic do they attract, in comparison to mainstream or real news websites?

[22] https://github.com/ymejova/yt-tw-covid

In [62], we perform a first of its kind investigation to answer such questions regarding the online presence of fake news websites and characterize their behavior in comparison to real news websites. We collect a dataset of 283 websites tagged from known fact-checking lists as delivering fake news and perform traffic and network analysis of such websites. In particular, and contrary to past work, we study and compare the user engagement of fake and real news sites by analyzing traffic-related metrics.

Additionally, we explore the lifetime of fake news sites and their typical uptime periods over a time range of more than 20 years, and we propose a methodology to detect websites that are synchronizing not only their uptime periods but also the content they serve during these periods. We detect numerous clusters of websites synchronizing their uptime and content for long periods of time within the USA presidential election years 2016-2017.

We also study the third party websites embedded in the different types of news sites (real and fake) and we find that during the aforementioned election years there is a significant increase in the use of analytics in the fake news sites but not an increase in the use of ad related third-parties. Additionally, domains like *doubleblick*, *googleadservices* and *scorecardresearch* tend to have higher presence in real news sites than in fake ones. On the contrary, *facebook* and *quantserve* have higher presence in fake news sites.

Finally, and based on our findings, we build a novel, content-agnostic ML classifier for automatic detection of fake news websites that are not yet included in manually curated blacklists. We tested various supervised and unsupervised ML methods for our classifier which achieved F1 score up to 0.942 and AUC of ROC up to 0.976. We provide our data open sourced for further studies.[23].

8.3.22 Modeling aggression propagation on social media

Cyberaggression has been studied in various contexts and online social platforms, and modeled on different data using state-of-the-art machine and deep learning algorithms to enable automatic detection and blocking of this behavior. Users can be influenced to act aggressively or even bully others because of elevated toxicity and aggression in their own (online) social circle. In effect, this behavior can propagate from one user and neighborhood to another, and therefore, spread in the network. Interestingly, to our knowledge, no work has modeled the network dynamics of aggressive behavior.

In the following pieces of work, this topic was investigated using three types of modeling: 1) Opinion Dynamics [67], 2) Independent Cascades [66] and Linear Threshold Modeling [66].

[23] https://github.com/mxalk/fake_news_resources

Contact: Chrysoula Terizi (UIoannina), Despoina Chatzakou (CERTH-ITI), Evaggelia Pitoura (UIoannina), Panayiotis Tsaparas (UIoannina), Nicolas Kourtellis (TID)

In [67], we took a first step towards this direction by studying propagation of aggression on social media using opinion dynamics. We proposed ways to model how aggression may propagate from one user to another, depending on how each user is connected to other aggressive or regular users. Through extensive simulations on Twitter data, we study how aggressive behavior could propagate in the network. We validate our models with crawled and annotated ground truth data, reaching up to 80% AUC, and discussed the results and implications of our work with respect to how Twitter and its administrators can decide which users (who are deemed important due to their key position in the network) to either ban from the platform from inappropriate behavior, or at least educate for its policies. We make available our code and data for future research on the topic[24].

Contact: Marinos Poiitis (AUTH) Athena Vakali (AUTH) Nicolas Kourtellis (TID)

In a follow-up work, in [66], we address aggression propagation modeling and minimization in Twitter, since it is a popular microblogging platform at which aggression had several onsets. We propose various methods building on two well-known diffusion models, Independent Cascade (IC) and Linear Threshold (LT), to study the aggression evolution in the social network. We experimentally investigate how well each method can model aggression propagation using real Twitter data, while varying parameters, such as seed users selection, graph edge weighting, users' activation timing, etc. It is found that the best performing strategies are the ones to select seed users with a degree-based approach, weigh user edges based on their social circles' overlaps, and activate users according to their aggression levels. We further employ the best performing models to predict which ordinary real users could become aggressive (and vice versa) in the future, and achieve up to AUC=0.89 in this prediction task. Finally, we investigate aggression minimization by launching competitive cascades to "inform" and "heal" aggressors. We show that IC and LT models can be used in aggression minimization, providing less intrusive alternatives to the blocking techniques currently employed by Twitter. We make available our code and data for future research on the topic[25].

8.3.23 Sentiment and Graph analysis on data related to US Elections 2020

Contact: Shevtsov Alexander (FORTH), Despoina Antonakaki (FORTH) Sotiris Ioannidis (FORTH)

[24] https://github.com/cterizi/Modeling-Aggression-Propagation-on-Social-Media

[25] https://anonymous.4open.science/r/95063f31-d3aa-4171-80d7-e0efb890476d/

The presidential elections in the United States on November 3rd, 2020 caused extensive discussions on social media. A part of the content on US elections is organic, coming from users discussing their opinions on the candidates, political positions, or relevant content presented on television. Another significant part originates from organized campaigns, both official, including communication campaigns and dissemination, or unofficial, including astroturfing and targeting manipulation of the electorate. In this study, we obtain approximately 19.8M tweets from 4.5M users, based on prevalent hashtags related to the 2020 US election. From these, we mined 28.343 YouTube links tweeted and obtained the comments of these videos. In this paper, we study the connection between two social networks. We employ an array of techniques, including volume analysis, exploring the retweet graph, sentiment, and graph analysis on the communities formed on YouTube and Twitter. Furthermore, we propose a method to combine the results of community detection on the two social networks and measure the differences between them. Particularly, we study the daily traffic per prevalent hashtags, plot the retweet graph from July to November 2020, highlight the two main entities ('Biden' and 'Trump') and show how the discussion around those entities grow in the period closer to the elections. Additionally, we perform sentiment analysis of both the Twitter corpus and the YouTube comments in tweeted videos. We found that 35,2% of the users contained in our Twitter dataset express positive sentiment towards current president Donald Trump and 28% express positive sentiment towards Joe Biden; while 18% of the users in our YouTube dataset express positive sentiment towards Donald Trump and 12% express positive sentiment towards Joe Biden.

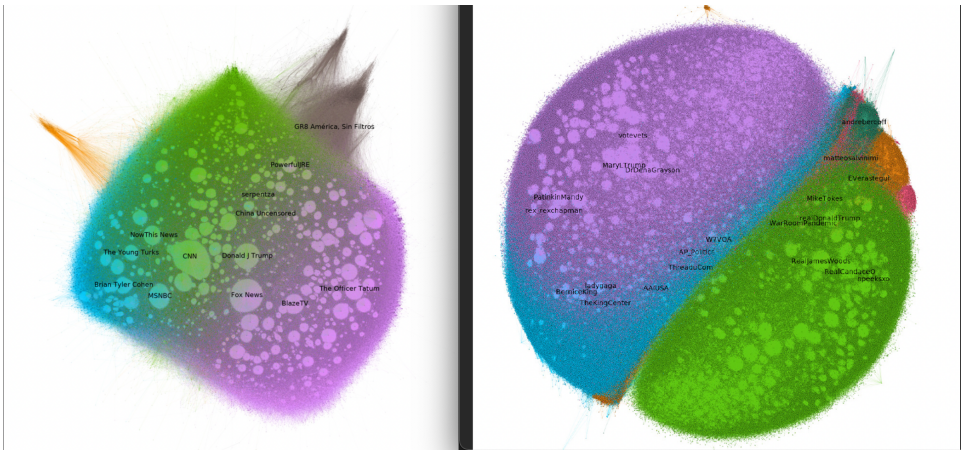

Figure 17: The 3-core of the YouTube comment and Retweet graph, color-coded by community.

Finally, we link the Twitter Retweet graph with the YouTube comment graph using tweeted video links. We measure their similarity and differences and show the

interactions and the correlation between the largest communities on YouTube and Twitter (figure 17).

This work has been submitted to PLOS One journal (currently waiting for third round of reviews). Currently there is only an available version in arXiv [83].

8.3.24 Identification of Twitter Bots Based on an Explainable ML Framework on the US 2020 Elections dataset

Contact: Shevtsov Alexander (FORTH), Despoina Antonakaki (FORTH) Sotiris Ioannidis (FORTH)

Twitter is one of the most popular social networks attracting millions of users, while a considerable proportion of online discourse is captured. It provides a simple usage framework with short messages and an efficient application programming interface (API) enabling the research community to study and analyze several aspects of this social network.

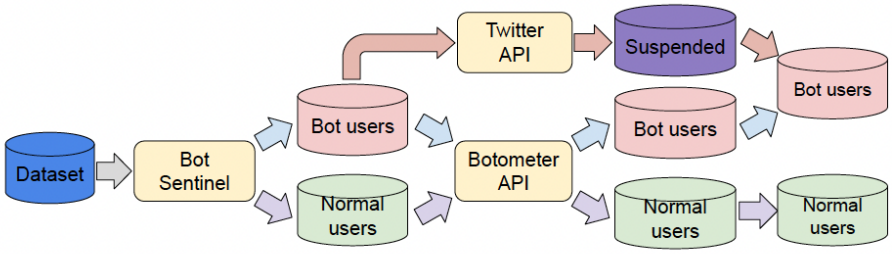

Figure 18: The bot/normal labeling process.

However, the Twitter usage simplicity can lead to malicious handling by various bots. The malicious handling phenomenon expands in online discourse, especially during the electoral periods, where except the legitimate bots used for dissemination and communication purposes, the goal is to manipulate the public opinion and the electorate towards a certain direction, specific ideology, or political party.

This paper focuses on the design of a novel system (figure 18) for identifying Twitter bots based on a labeled Twitter dataset consisting of 15.6M tweets and 3.2M users gathered from 1st of September 2020 to 3rd of November 2020. To this end, a supervised machine learning (ML) framework is adopted using an Extreme Gradient Boosting (XGBoost) algorithm, where the hyper-parameters are tuned via cross-validation. Our study also deploys Shapley Additive Explanations (SHAP) for explaining the ML model predictions by calculating feature importance (figure 19), using the game theoretic-based Shapley values. Experimental evaluation on distinct Twitter datasets demonstrate the superiority of our approach, in terms of

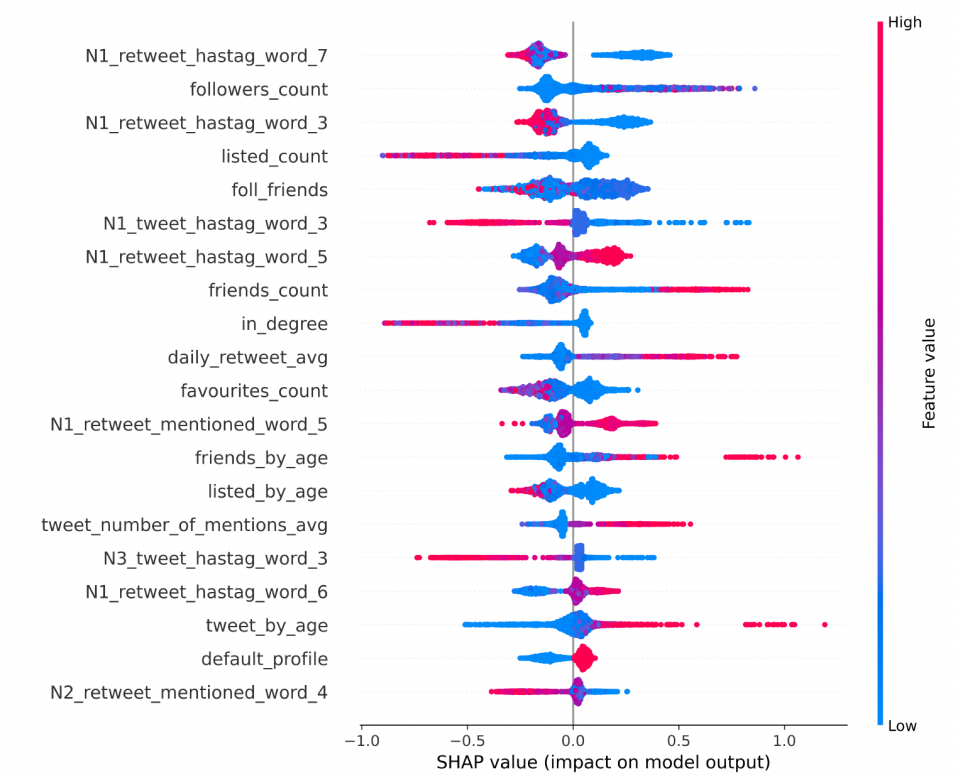

Figure 19: The most important features for our model are the statistical, time and graph based features, which means that combination of them will give us the best possible performance for our model.

bot detection accuracy, when compared against a recent state-of-the-art Twitter bot detection method.

This work has been accepted for publication at the AAAI ICWSM-2022 [84].

8.3.25 Gender differences in privacy awareness on social networking sites

Contact: Simon Kocbek (UM)

We explored gender differences in privacy awareness on social networking sites [85]. Specifically, we created two Facebook profile accounts of an average 18-year-old female and male. We have added 30 female and 30 male profiles on Facebook as friends from created accounts and waited to see if the profiles will get friend request accepted. We have done a comparative analysis based on gender, age and number of friends of profiles that have been added by the created profiles. We have found that male users of Facebook are more inclined towards adding new, unfamiliar friends on Facebook than female users. This might be due to privacy awareness across gender or cultural norms, where female think about their safety before adding someone unknown. Also, users with higher number of friends are willing to add strangers on Facebook more easily.

9 Organization of the Scientific Community and Events

The Description of Action identifies four (SMART) objectives for WP1 (see Section 1). Two of these objectives concern 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.

9.1 Organization of Scientific Events

CONCORDIA members continue to be active in the organization of scientific conferences, whenever possible. These conferences include:

- ACM Conference on Data and Application Security and Privacy
- Workshop on Security Aspects in Autonomous Systems
- Secure Runtime Environments session at HiPEAC CSW Spring
- IFIP/IEEE International Symposium on Integrated Network Management
- ACM SIGCOMM Workshop on Technologies, Applications, and Uses of a Responsible Internet
- IEEE International Conference Web Services
- IEEE International Conference on Cloud Computing
- The Future of Cybersecurity in Slovenia and Europe
- International Conference on Network and Service Management
- IEEE International Conference on Network Protocols
- IEEE International Conference on Trust, Privacy and Security in Intelligent Systems and Applications
- ACM 1st Workshop on Security and Privacy for Mobile AI (MAISP) collocated with ACM MobiSys
- European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases

A more detailed view can be found in Appendix B.

Membership of Technical Program Committees (TPC)

Appendix C contains a table with all TPC related efforts of CONCORDIA researchers during the year 2021. As can be seen, researchers have been active as TPC members in over 45 different conferences and in some cases, quite a few members from CONCORDIA serve on the TPC of specific events.

9.2 Organization of Scientific Community

The organization of the scientific community involves the following three activities:

- Chairing professional organizations (IFIP, CNRS, and ANSSI);
- Editing scientific journals;
- Membership of steering committees.

Chairing Professional Organizations

The scientific community is organized by professional organizations such as IFIP, CNRS, and ANSSI. Various CONCORDIA researchers hold positions within these organizations. The following positions are (or have continued to be through the third year) held by CONCORDIA researchers:

- Chair of IFIP TC6[26] (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.
- Chair of IFIP TC6 Working Group 6.6[27] (Rémi Badonnel, UL)
  The Working Group 6 of IFIP TC6 focuses on the management of network and distributed systems. Management is defined in five functional areas – the so-called "FCAPS" areas. These involve: Fault management, Configuration management, Accounting management, Performance management and Security management. Security management has become the greatest challenge.
- President of the Scientific Council of CNRS[28] working group on Cyber-Security (Jean-Yves Marion)

[26] https://ifip.informatik.uni-hamburg.de/ifip/tc/6
[27] https://ifip.informatik.uni-hamburg.de/ifip/tc/6/wg/6.6/officers
[28] https://gdr-securite.irisa.fr

- Member of the Scientific Council of ANSSI[29](Jean-Yves Marion), French cybersecurity agency.

Editing Scientific Journals

Six researchers from five different partners of the CONCORDIA consortium acted as journal editors in 2021, with varying roles including editorial board members and associate editor. They were active for the following 15 journals:

- ACM Transactions on Data Science
- ACM Transactions on Privacy and Security
- ACM SIGCOMM Computer Communications Review
- IEEE Transactions on Big Data
- Cybersecurity and Privacy of Frontiers in Big Data
- Data Science and Engineering
- IEEE Transactions on Cloud Computing
- IEEE Transactions on Network and Service Management
- Digital Trust: Trust Management in the Cyberspace, IEEE Internet computing
- IEEE Internet Computing
- IEEE Transactions on Service Computing
- International Journal of Cooperative Information Systems
- Journal of Network and Systems Management
- International Journal of Network Management
- Synthesis Lectures on Security, Privacy and Trust

Steering Committee Membership

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

- IFIP TMA – Traffic Measurement and Analysis Conference[30] (Anna Sperotto, UT)
- WTMC – International Workshop on Traffic Measurements for Cybersecurity[31] (Anna Sperotto, UT)

[29]https://www.ssi.gouv.fr
[30]https://tma.ifip.org/
[31]https://wtmc.info/index.html

- ANRW – ACM/IRTF Applied Networking Research Workshop[32] (Rolan van Rijswijk, UT)

Non-Academic Conferences

On occasion, CONCORDIA researchers are also involved in the organization of non-academic events. We list such events here, provided they have a strong focus on cybersecurity and aim for knowledge sharing and technological advancement.

- Talk at FIC[33], The reference event in Europe for digital security and trust, by Jean-Yves Marion (UL) and Sylvain Cecchetto (CYD)

[32] https://irtf.org/anrw/2021/
[33] https://www.forum-fic.com/accueil.htm

10 Contributions to Standards and Open Research Data

One of the objectives of WP1 (see Section 1) is to contribute to standardization, open research data and code. This section outlines the achievements in this area by WP1 researchers.

GitHub, GitLab and other public repositories are used to host around 10 open CONCORDIA projects. More information can be found in each specific Section/research activity.

Standardization efforts take multiple years from the initial research until the establishment of a full standard. These activities within CONCORDIA have started years before, but obtained important outcome during the course of the project and have some relationship with current WP1 research. 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.

Noteworthy, here, is the case of the OpenINTEL measurement project, that was first reported in D1.1 and continues to grow and yield results. The Domain Name System (DNS) plays a key role in almost all Internet services. This is also why, for the past years of the CONCORDIA project, Task 1.2 (see Section 5) has seen a strong focus on research related to DNS security and stability.

The goal of the OpenINTEL project is to collect daily snapshots of the state of the DNS. By analyzing DNS information we are able to detect various types of security issues, such as single points of failure in authoritative nameserver infrastructure, DNS misconfiguration, and the impact of attacks on the normal operation of the DNS. By collecting DNS data over an extended period, OpenINTEL provides insights into the evolution of the DNS. Active measurement, rather than passive collection, can be used to build consistent and reliable time series of the state of the DNS. Since D1.1, in which this collection effort was first reported, a number of zones were added to the project, resulting in 236 million names measured per name. Notably, the project now also measures a sensible and sizable part of the reverse IPv4 address space daily, providing valuable data for new types of research. The details regarding OpenINTEL and the public data repository can be found on https://www.openintel.nl.

A complete overview of all CONCORDIA activities related to standardization and open data is included in deliverables of WP5 (D5.2, D5.3, D5.4) and WP6 (D6.4, D6.5, D6.6).

11 Conclusions and Outlook

Despite the ongoing COVID-19 pandemic, the third year of the CONCORDIA project has been a good one for WP1. In terms of the main objective, stimulating the publication of scientific results in the broad field of cybersecurity, WP1 researchers did exceptionally well. All of the five WP1 tasks saw the successful execution of research directives set in year one, as well as the continuation of ongoing efforts since year two. As a result, more than 75 additional papers were published during the third year, many of them in renowned conferences and journals. This falls relatively short compared to the second year that was the most successful, however, it gives a promising outlook for the next year in terms of writing high-quality papers.

While the quantity and quality of the published papers is very satisfying, the percentage of open access publications (close to 50% in 2021) is below the expectations. In order to improve the situation, we will undertake some measures in 2022 to further educate researchers about the different open access publishing options. We will also encourge researchers to make preprints and postprints of already reported papers openly accessible where this is feasible, considering copyright and embargo rules of the various scientific publishers.

The dissemination of successful WP1 research went beyond academic conferences and journals alone. CONCORDIA researchers contributed to the wider Internet community, targeting a wider audience with research papers on the impact of Twitter and YouTube on a number of subjects.

The collaboration among WP1 tasks and the project pilots, that has been ongoing since the first year of the project, has continued during this third year. A concrete example is the number of joint-work publications (more than 10). Another example is the strong collaboration that has been established between application-centric security research (T1.4) and the Threat Intelligence for the Telco Sector pilot (T2.1), where T1.4's efforts in investigating application security has been instrumental in making sure that requirements set by T2.1 are met.

Concerning the organization of scientific events, CONCORDIA researchers excelled there well. Researchers sat on the technical program committees of over 45 different conferences during the third year. Moreover, they were involved in the organization of 14 academic and non-academic yet cybersecurity-related events. Most of these events had to be held virtually, which created challenges.

The COVID-19 pandemic was and continues to be with us during the third year of the project. From an academic perspective alone, the pandemic has brought about a number of challenges. When it comes to the impact on WP1 activities, the pandemic's effect on research has been limited but certainly not non-existent. Many of WP1 research activities do not involve labs, but on the rare occasion that a lab had to be shut down, research was negatively affected. From a mental

well-being point of view, there are more unknowns. Junior researchers often live abroad and have a smaller social safety net. Many of them have been isolated back home with their families, due to institutions suspending their activities. The pandemic has also made organizing and attending events in-person much harder, which hinders opportunities to meet peers and develop a professional network.

As a concluding note, the consortium is very satisfied with the progress that we made, especially in the face of COVID-19 that affected all aspects of everyday life and work. Research done in WP1 has been very broad, ranging from hardware to social networks, demonstrating the pervasive nature of security concerns and the need to deploy mechanisms and analyze behavior at all levels. Furthermore, research focuses on Internet-enabling technologies, and on building a more secure infrastructure for the future: DNS, blockchain, 5G, software infrastructure, social support infrastructure (social networks, YouTube, etc.). There is also diversity in terms of applications considered and techniques employed: behavior analysis of and defense against malware attacks, cryptographic primitives, machine learning, vehicle security and many more. This demonstrates a wide range of expertise available in the consortium, that will help us move forward with research in the years to come and Europe to build its digital sovereignty when looking towards the future.

For 2022, WP1 has a number of goals. Firstly, we would like to continue strengthening the support for young PhD researchers and help them shape their careers into the future. Whenever possible, we would like to collaborate in this area with the three other pilot projects, as well as other organizations, to ensure that support continues even after the end of this project. Secondly, we would like to extend the collaboration with industry partners and increase the impact of our research, even more than in the previous years. An additional important goal related to this, is to intensify collaboration with T2.3, which has been delayed due to the change in its leadership, as well as the impacts of COVID-19. Furthermore, if the global health crisis permits it, we plan to catch up on physical meetings which have been hindered or even made impossible by the pandemic. Finally, we would like to continue investigating the role and behavior of social media, to increase the security and privacy of our end users.

References

[1] Mattijs Jonker, Jean-Yves Marion, Aiko Pras, Jürgen Schönwälder, Nikos Salamanos, Michael Sirivianos, Marinos Tsantekidis, and Ramin Yazdani. Deliverable D1.2: 2nd Year Report on Designing and Developing an European Secure, Resilient and Trusted Ecosystem (ESRTE), 2020.

[2] Gabi Dreo, Corinna Schmitt, and Arthur van der Wees. Deliverable D4.4: Cybersecurity Roadmap for Europe by CONCORDIA, 2020.

[3] George Christou, Giorgos Vasiliadis, Elias Athanasopoulos, and Sotiris Ioannidis. Hard edges: Hardware-based control-flow integrity for embedded devices. In International Conference on Embedded Computer Systems: Architectures, Modeling and Simulation, July 2021.

[4] M. Wolf and J. Schönwälder. Applying Metamorphic Testing to Homomorphic Cryptography. In Proc. 6th IEEE/ACM International Workshop on Metamorphic Testing (MET 2021). IEEE, June 2021.

[5] Charis Dimopoulos, Apostolos P. Fournaris, and Odysseas Koufopavlou. Machine learning attacks and countermeasures on hardware binary edwards curve scalar multipliers. Journal of Sensor and Actuator Networks, 10(3), 2021.

[6] Naila Mukhtar, Apostolos P. Fournaris, Tariq M. Khan, Charis Dimopoulos, and Yinan Kong. Improved hybrid approach for side-channel analysis using efficient convolutional neural network and dimensionality reduction. IEEE Access, 8:184298–184311, 2020.

[7] Evangelos Haleplidis, Thanasis Tsakoulis, Alexander El-Kady, Charis Dimopoulos, Odysseas Koufopavlou, and Apostolos P. Fournaris. Studying opencl-based number theoretic transform for heterogeneous platforms. In 2021 24th Euromicro Conference on Digital System Design (DSD), pages 339–346, 2021.

[8] Alexander ElKady, Apostolos P. Fournaris, Thanasis Tsakoulis, Evangelos Haleplidis, and Vassilis Paliouras. High-level synthesis design approach for number-theoretic transform implementations. In 29th IFIP/IEEE International Conference on Very Large Scale Integration, pages 339–346, 2021.

[9] Claudio A. Ardagna, Rasool Asal, Ernesto Damiani, Nabil El Ioini, Mehdi Elahi, and Claus Pahl. From trustworthy data to trustworthy iot: A data collection methodology based on blockchain. ACM Transactions on Cyber-Physical Systems, December 2021.

[10] Marco Anisetti, Claudio A. Ardagna, Filippo Berto, and Ernesto Damiani. Security certification scheme for content-centric networks. In 2021 IEEE International Conference on Services Computing (SCC), 2021.

[11] Marco Anisetti, Claudio A. Ardagna, Nicola Bena, and Ruslan Bondaruc. Towards an assurance framework for edge and iot systems. In 2021 IEEE International Conference on Edge Computing (EDGE), 2021.

[12] Sina Rafati Niya, Eryk Schiller, and Burkhard Stiller. Architectures for Blockchain-IoT Integration. In Nur Zincir-Heywood, Yixin Diao, and Marco Mellia, editors, Communication Networks and Service Management in the Era of Artificial Intelligence and Machine Learning, IEEE Press Series on Networks and Service Management, pages 100–137, New York, NY, USA, October 2021. Wiley-IEEE Press.

[13] Eryk Schiller, Ramon Huber, and Burkhard Stiller. Python-Based TinyIPFIX in Wireless Sensor Networks. In 2021 IEEE 46th Conference on Local Computer Networks (LCN), pages 431–434, 2021.

[14] Pietro Colombo and Elena Ferrari. Access control enforcement within mqtt-based internet of things ecosystems. In Proceedings of the 23nd ACM on Symposium on Access Control Models and Technologies, SACMAT '18, page 223–234, New York, NY, USA, 2018. Association for Computing Machinery.

[15] Moshe Kravchik, Battista Biggio, and Asaf Shabtai. Poisoning attacks on cyber attack detectors for industrial control systems. In Proceedings of the 36th Annual ACM Symposium on Applied Computing, pages 116–125, 2021.

[16] Lior Sidi, Yisroel Mirsky, Asaf Nadler, Yuval Elovici, and Asaf Shabtai. Helix: Dga domain embeddings for tracking and exploring botnets. In Proceedings of the 29th ACM International Conference on Information & Knowledge Management, pages 2741–2748, 2020.

[17] Lior Sidi, Asaf Nadler, and Asaf Shabtai. Maskdga: An evasion attack against dga classifiers and adversarial defenses. IEEE Access, 8:161580–161592, 2020.

[18] Shimon Harush, Yair Meidan, and Asaf Shabtai. Deepstream: Autoencoder-based stream temporal clustering and anomaly detection. Computers & Security, 106:102276, 2021.

[19] Shimon Harush, Yair Meidan, and Asaf Shabtai. Deepstream: autoencoder-based stream temporal clustering. In Proceedings of the 36th Annual ACM Symposium on Applied Computing, pages 445–448, 2021.

[20] Han Wang, Luis Muñoz-González, David Eklund, and Shahid Raza. Non-IID Data Rebalancing at IoT Edge with Peer-to-Peer Federated Learning for Anomaly Detection. In Proceedings of the 14th ACM Conference on Security and Privacy in Wireless and Mobile Networks, pages 153–163, 2021.

[21] Eder John Scheid, Andreas Knecht, Tim Strasser, Christian Killer, Muriel Franco, Bruno Rodrigues, and Burkhard Stiller. Edge2BC: a Practical Approach for Edge-to-Blockchain IoT Transactions. In IEEE International Conference on Blockchain and Cryptocurrency (ICBC 2021), pages 1–9, Sydney, Australia, May 2021.

[22] Eder John Scheid, Pascal Kiechl, Muriel Franco, Bruno Rodrigues, Christian Killer, and Burkhard Stiller. Security and Standardization of a Notary-based Blockchain Interoperability API. In IEEE International Conference on Blockchain Computing and Applications (BCCA 2021), pages 1–7, Tartu, Estonia, November 2021. Accepted. To be published.

[23] Sina Rafati Niya, Julius Willems, and Burkhard Stiller. On-chain iot data modification in blockchains, 2021. Under Submission.

[24] Sina Rafati Niya, Fabio Maddaloni, Thomas Bocek, and Burkhard Stiller. Toward Scalable Blockchains with Transaction Aggregation. In Symposium on Applied Computing (SAC 2020), page 308–315, Brno, Czech Republic, 2020. ACM.

[25] Eryk Schiller, Elfat Esati, and Burkhard Stiller. IoT-based Access Management Supported by AI and Blockchains. In 2021 17th International Conference on Network and Service Management (CNSM), pages 350–354, 2021.

[26] Lital Badash, Nachiket Tapas, Asaf Nadler, Francesco Longo, and Asaf Shabtai. Blockchain-based bug bounty framework. In Proceedings of the 36th Annual ACM Symposium on Applied Computing, pages 239–248, 2021.

[27] Olivier van der Toorn, Moritz Müller, Sara Dickinson, Cristian Hesselman, Anna Sperotto, and Roland van Rijswijk-Deij. Addressing the Challenges of Modern DNS A Comprehensive Tutorial. Computer Science Review, 2021. (To appear).

[28] Moritz Christian Müller. Making DNSSEC Future Proof. PhD thesis, University of Twente, Netherlands, Sep 2021.

[29] Raffaele Sommese, Gautam Akiwate, Mattijs Jonker, Giovane CM Moura, Marco Davids, Roland van Rijswijk-Deij, Geoffrey M Voelker, Stefan Savage, KC Claffy, and Anna Sperotto. Characterization of Anycast Adoption in the DNS Authoritative Infrastructure. In Network Traffic Measurement and Analysis Conference (TMA'21), 2021.

[30] Olivier van der Toorn, Johannes Krupp, Mattijs Jonker, Roland van Rijswijk-Deij, Christian Rossow, and Anna Sperotto. ANYway: Measuring the Amplification DDoS Potential of Domains. In 17th International Conference on Network and Service Management (CNSM), pages 500–508, 2021.

[31] Eva Papadogiannaki and Sotiris Ioannidis. A survey on encrypted network traffic analysis applications, techniques, and countermeasures. ACM Comput. Surv., 54(6), July 2021.

[32] Eva Papadogiannaki and Sotiris Ioannidis. Acceleration of intrusion detection in encrypted network traffic using heterogeneous hardware. Sensors, 21(4):1140, 2021.

[33] Stanislav Špaček, Daniel Tovarňák, and Pavel Čeleda. Enriching DNS Flows with Host-Based Events to Bypass Future Protocol Encryption. In IFIP International Conference on ICT Systems Security and Privacy Protection, pages 302–316. Springer, 2021.

[34] Eder J. Scheid, Bruno Rodrigues, and Burkhard Stiller. Policy-based Blockchain Selection. IEEE Communications Magazine, 59(10):48–54, 2021.

[35] Edward Chuah, Arshad Jhumka, Samantha Alt, R Todd Evans, and Neeraj Suri. Failure diagnosis for cluster systems using partial correlations. In The 19th IEEE International Symposium on Parallel and Distributed Processing with Applications, pages 1–11, 2021.

[36] Nicolas Coppik, Oliver Schwahn, and Neeraj Suri. Fast kernel error propagation analysis in virtualized environments. In 2021 14th IEEE Conference on Software Testing, Verification and Validation (ICST), pages 159–170. IEEE, 2021.

[37] Binlin Cheng, Jiang Ming, Erika Leal, Haotian Zhang, Jianming Fu, Guojun Peng, and Jean-Yves Marion. Extracting Executable Payloads From Packed Malware: Import Table Reconstruction via Hardware-Assisted API Micro Execution. In USENIX Security '21 Fall, August 2021.

[38] Peter Auer, Nicolo Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed bandit problem. Machine learning, 47(2):235–256, 2002.

[39] Aurélien Garivier and Olivier Cappé. The kl-ucb algorithm for bounded stochastic bandits and beyond. In Proceedings of the 24th annual conference on learning theory, pages 359–376. JMLR Workshop and Conference Proceedings, 2011.

[40] Benjamin Green, Richard Derbyshire, Marina Krotofil, William Knowles, Daniel Prince, and Neeraj Suri. Pcaad: Towards automated determination and exploitation of industrial systems. Computers & Security, 110:102424, 2021.

[41] Tristan Benoit, Jean-Yves Marion, and Sébastien Bardin. Binary level toolchain provenance identification with graph neural networks. In 28th IEEE International Conference on Software Analysis, Evolution and Reengineering, SANER 2021, Honolulu, HI, USA, March 9-12, 2021, pages 131–141. IEEE, 2021.

[42] Jukka Soikkeli, Cora Perner, and Emil C. Lupu. Analyzing the viability of UAV missions facing cyber attacks. In IEEE European Symposium on Security and Privacy Workshops, EuroS&P 2021, Vienna, Austria, September 6-10, 2021, pages 103–112. IEEE, 2021.

[43] Marinos Tsantekidis and Vassilis Prevelakis. Securing Runtime Memory via MMU manipulation. In The Fifteenth International Conference on Emerging Security Information, Systems and Technologies (SECURWARE), 2021.

[44] Marinos Tsantekidis. and Vassilis Prevelakis. Mmu-based access control for libraries. In Proceedings of the 18th International Conference on Security and Cryptography - SECRYPT, pages 686–691. INSTICC, SciTePress, 2021.

[45] Mohammad Hamad, Marinos Tsantekidis, and Vassilis Prevelakis. Intrusion Response System for Vehicles: Challenges and Vision. In Markus Helfert, Cornel Klein, Brian Donnellan, and Oleg Gusikhin, editors, Smart Cities, Green Technologies and Intelligent Transport Systems, pages 321–341, Cham, 2021. Springer International Publishing.

[46] Mohammad Hamad, Emanuel Regnath, Jan Lauinger, Vassilis Prevelakis, and Sebastian Steinhorst. SPPS: Secure Policy-based Publish/Subscribe System for V2C Communication. In 2021 Design, Automation Test in Europe Conference Exhibition (DATE), pages 529–534, 2021.

[47] Boning Feng, Akihiro Kajiwara, Van Thuan Do, Niels Jacot, Bernardo Santos, Bruno Dzogovic, and Thanh van Do. A Secure 5G Eldercare Solution Using Millimeterwave Sensors. In Jamal Bentahar, Irfan Awan, Muhammad Younas, and Tor-Morten Grønli, editors, Mobile Web and Intelligent Information Systems, pages 3–15, Cham, 2021. Springer International Publishing.

[48] Thanh van Do, Boning Feng, Birgitta Langhammer, Van Thuan Do, Niels Jacot, Bruno Dzogovic, Bernardo Santos, and Per Jonny Nesse. Ageing@home: A secure 5G welfare technology solution for elderlies. In 7th EAI International Conference on Smart Objects and Technologies for Social Good (GOODTECHS), 2021.

[49] Bruno Dzogovic, Bernardo Santos, Boning Feng, Van Thuan Do, Niels Jacot, and Thanh Van Do. Optimizing 5g vpn+ transport networks with vector packet processing and fpga cryptographic offloading. In Jamal Bentahar, Irfan Awan, Muhammad Younas, and Tor-Morten Grønli, editors, Mobile Web and Intelligent Information Systems, pages 85–98, Cham, 2021. Springer International Publishing.

[50] Bernardo Santos, Imran Qayyum Khan, Bruno Dzogovic, Boning Feng, Van Thuan Do, and Niels Jacot. Anomaly Detection in Cellular IoT with Machine Learning. In 7th EAI International Conference on Smart Objects and Technologies for Social Good (GOODTECHS), 2021.

[51] Martina Sestak, Marjan Hericko, Tatjana Welzer Druzovec, and Muhamed Turkanovic. Applying k-vertex cardinality constraints on a neo4j graph database. Future Generation Computer Systems, 115:459–474, 2021.

[52] Aida Kamišalić, Renata Kramberger, and Iztok Fister. Synergy of blockchain technology and data mining techniques for anomaly detection. Applied Sciences, 11(17), 2021.

[53] Mohamed Oulaaffart, Remi Badonnel, and Olivier Festor. Towards automating security enhancement for cloud services. In 2021 IFIP/IEEE International Symposium on Integrated Network Management (IM), pages 692–696, 2021.

[54] Nicolas Schnepf, Rémi Badonnel, Abdelkader Lahmadi, and Stephan Merz. Automated Orchestration of Security Chains Driven by Process Learning*, chapter 12, pages 289–319. John Wiley and Sons, Ltd, 2021.

[55] Marco Anisetti, Claudio A. Ardagna, Nicola Bena, and Andrea Foppiani. An assurance-based risk management framework for distributed systems. In 2021 IEEE International Conference on Web Services (ICWS), 2021.

[56] Marco Anisetti, Claudio A. Ardagna, Nicola Bena, and Ernesto Damiani. An assurance framework and process for hybrid systems. In E-Business and Telecommunications, pages 79–101, 2021.

[57] Ammar Battah, Youssef Iraqi, and Ernesto Damiani. Blockchain-based reputation systems: Implementation challenges and mitigation. Electronics, 10(3), 2021.

[58] Federico Daidone, Barbara Carminati, and Elena Ferrari. Blockchain-based privacy enforcement in the iot domain. IEEE Transactions on Dependable and Secure Computing, pages 1–1, 2021.

[59] Fan Mo, Hamed Haddadi, Kleomenis Katevas, Eduard Marin, Diego Perino, and Nicolas Kourtellis. Ppfl: Privacy-preserving federated learning with trusted execution environments. In Proceedings of the 19th Annual International Conference on Mobile Systems, Applications, and Services, MobiSys '21, page 94–108, 2021.

[60] Michalis Pachilakis, Panagiotis Papadopoulos, Nikolaos Laoutaris, Evangelos P. Markatos, and Nicolas Kourtellis. Youradvalue: Measuring advertising price dynamics without bankrupting user privacy. ACM Measurement and Analysis of Computing Systems (POMACS), 5(3), 12 2021.

[61] Emmanouil Papadogiannakis, Panagiotis Papadopoulos, Nicolas Kourtellis, and Evangelos P. Markatos. User tracking in the post-cookie era: How websites bypass gdpr consent to track users. In Proceedings of the Web Conference 2021, WWW '21, page 2130–2141, 2021.

[62] Manolis Chalkiadakis, Alexandros Kornilakis, Panagiotis Papadopoulos, Evangelos Markatos, and Nicolas Kourtellis. The rise and fall of fake news sites: A traffic analysis. In 13th ACM Web Science Conference 2021, WebSci '21, page 168–177, 2021.

[63] Yelena Mejova and Nicolas Kourtellis. Youtubing at home: Media sharing behavior change as proxy for mobility around covid-19 lockdowns. In 13th ACM Web Science Conference 2021, WebSci '21, page 272–281, 2021.

[64] Yash Vekaria, Vibhor Agarwal, Pushkal Agarwal, Sangeeta Mahapatra, Sakthi Balan Muthiah, Nishanth Sastry, and Nicolas Kourtellis. Differential tracking across topical webpages of indian news media. In 13th ACM Web Science Conference 2021, WebSci '21, page 299–308, 2021.

[65] Vibhor Agarwal, Yash Vekaria, Pushkal Agarwal, Sangeeta Mahapatra, Shounak Set, Sakthi Balan Muthiah, Nishanth Sastry, and Nicolas Kourtellis. Under the spotlight: Web tracking in indian partisan news websites. Proceedings of the International AAAI Conference on Web and Social Media, 15(1):26–37, May 2021.

[66] Marinos Poiitis, Athena Vakali, and Nicolas Kourtellis. On the aggression diffusion modeling and minimization in twitter. Transactions on the Web, 2021.

[67] Chrysoula Terizi, Despoina Chatzakou, Evaggelia Pitoura, Panayiotis Tsaparas, and Nicolas Kourtellis. Modeling aggression propagation on social media. Online Social Networks and Media, 24:100137, 2021.

[68] Anaobi Ishaku Hassan, Aravindh Raman, Ignacio Castro, Haris Bin Zia, Emiliano De Cristofaro, Nishanth Sastry, and Gareth Tyson. Exploring Content Moderation in the Decentralised Web: The Pleroma Case. In Proceedings of the 17th International Conference on Emerging Networking Experiments and Technologies (ACM CoNext 2021), December 2021.

[69] Michalis Diamantaris, Serafeim Moustakas, Lichao Sun, Sotiris Ioannidis, and Jason Polakis. This Sneaky Piggy Went to the Android Ad Market: Misusing Mobile Sensors for Stealthy Data Exfiltration. In Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security (CCS '21), pages 390–398, November 2021.

[70] Claudio A. Ardagna, Valerio Bellandi, Ernesto Damiani, Michele Bezzi, and Cedric Hebert. Big data analytics-as-a-service: Bridging the gap between security experts and data scientists. Computers & Electrical Engineering, 93:107215, 2021.

[71] Anh-Tu Hoang, Barbara Carminati, and Elena Ferrari. Cluster-based anonymization of knowledge graphs. In 18th International Conference on Applied Cryptography and Network Security, 2020.

[72] Anh-Tu Hoang, Barbara Carminati, and Elena Ferrari. Privacy-preserving sequential publishing of knowledge graphs. In 37th IEEE International Conference on Data Engineering, 2021.

[73] Ha Xuan Son, Barbara Carminati, and Elena Ferrari. A risk assessment mechanism for android apps. In IEEE International Conference on Smart Internet of Things, pages 237–244. IEEE, August 2021.

[74] Flaviene Scheidt de Cristo, Wazen M Shbair, Lucian Trestioreanu, Aanchal Malhotra, and Radu State. Privacy-preserving paystring service. In 2021 IEEE International Conference on Blockchain and Cryptocurrency (ICBC), pages 1–3. IEEE, 2021.

[75] Sina Rafati Niya, Danijel Dordevic, and Burkhard Stiller. ITrade: A Blockchain-based, Self-Sovereign, and Scalable Marketplace for IoT Data Streams. In IFIP/IEEE International Symposium on Integrated Network Management (IM 2021), pages 530–536, Bordeaux, France, May 2021. IFIP/IEEE.

[76] Christian Killer, Bruno Rodrigues, Eder J. Scheid, Muriel Franco, and Burkhard Stiller. From Centralized to Decentralized Remote Electronic Voting. In To appear, To appear, pages 1–38. To appear, Cham, Switzerland, To appear 2021.

[77] Sina Rafati Niya, Raphael Beckmann, Claudio Brasser, Michael Bucher, Nicolas Spielmann, and Burkhard Stiller. Deti: A decentralized ticketing management platform. Under Submision.

[78] Sina Rafati Niya and Burkhard Stiller. Enabling Technologies and Distributed Storage, 2021. Under Submision.

[79] Wazen Shbair, Eugene Gavrilov, et al. Hsm-based key management solution for ethereum blockchain. In IEEE International Conference on Blockchain and Cryptocurrency, 3-6 May 2021. IEEE, 2021.

[80] Marius Paraschiv, Nikos Salamanos, Costas Iordanou, Nikolaos Laoutaris, and Michael Sirivianos. A unified graph-based approach to disinformation detection using contextual and semantic relations. In 16th International Conference on Web and Social Media (ICWSM 2022), June 2022.

[81] Kostantinos Papadamou, Savvas Zannettou, Jeremy Blackburn, Emiliano De Cristofaro, Gianluca Stringhini, and Michael Sirivianos. "it is just a flu": Assessing the effect of watch history on youtube's pseudoscientific video recommendations. 16th International Conference on Web and Social Media (ICWSM 2022), 2022.

[82] Kostantinos Papadamou, Savvas Zannettou, Jeremy Blackburn, Emiliano De Cristofaro, Gianluca Stringhini, and Michael Sirivianos. "how over is it?" understanding the incel community on youtube. Proc. ACM Hum.-Comput. Interact., 5(CSCW2), October 2021.

[83] Alexander Shevtsov, Maria Oikonomidou, Despoina Antonakaki, Polyvios Pratikakis, and Sotiris Ioannidis. Sentiment and graph analysis on data related to us elections 2020. arXiv preprint: available at https://arxiv.org/abs/2010.08183, 2020.

[84] Alexander Shevtsov, Maria Oikonomidou, Despoina Antonakaki, Polyvios Pratikakis, and Sotiris Ioannidis. Identification of twitter bots based on an explainable ml framework on the us 2020 elections dataset. In 16th International Conference on Web and Social Media (AAAI ICWSM-2022), June 2022.

[85] Lili Nemec Zlatolas and Tatjana Welzer. Gender differences in privacy awareness on social networking sites. In 8th ACM Celebration of Women in Computing: womENcourage™ 2021, Virtual, September 22-24, 2021. ACM, 2021.

[86] Pietro Colombo, Elena Ferrari, and Engin Deniz Tümer. Regulating data sharing across mqtt environments. Journal of Network and Computer Applications, 174:102907, 2021.

[87] Marcin Nawrocki, Mattijs Jonker, Thomas C Schmidt, and Matthias Wählisch. The Far Side of DNS Amplification: Tracing the DDoS Attack Ecosystem from the Internet Core. In Proceedings of the 21st ACM Internet Measurement Conference, pages 419–434, 2021.

[88] Leandro M Bertholdo, Joao M Ceron, Wouter B de Vries, Ricardo de Oliveira Schmidt, Lisandro Zambenedetti Granville, Roland van Rijswijk-Deij, and Aiko Pras. Tangled: A Cooperative Anycast Testbed. In 2021 IFIP/IEEE International Symposium on Integrated Network Management (IM), pages 766–771. IEEE, 2021.

[89] Leandro M Bertholdo, João M Ceron, Lisandro Z Granville, and Roland van Rijswijk-Deij. Forecasting the Impact of IXP Outages Using Anycast. In Network Traffic Measurement and Analysis Conference (TMA'21), 2021.

[90] Muriel Franco, Jan Von der Assen, Luc Boillat, Christian Killer, Bruno Rodrigues, Eder J Scheid, Lisandro Granville, and Burkhard Stiller. SecGrid: a Visual System for the Analysis and ML-based Classification of Cyberattack Traffic. In 2021 IEEE 46th Conference on Local Computer Networks (LCN), pages 140–147. IEEE, 2021.

[91] Eder J Scheid, Bruno B Rodrigues, Christian Killer, Muriel F Franco, Sina Rafati, and Burkhard Stiller. Blockchains and Distributed Ledgers Uncovered: Clarifications, Achievements, and Open Issues. In Advancing Research in Information and Communication Technology, pages 289–317. Springer, 2021.

[92] Jean-Philippe Eisenbarth, Thibault Cholez, and Olivier Perrin. An open measurement dataset on the Bitcoin P2P Network. In IM 2021 - The 17th IFIP/IEEE International Symposium on Integrated Network Management, page 5, Bordeaux / Virtual, France, May 2021.

[93] Jean-Philippe Eisenbarth, Thibault Cholez, and Olivier Perrin. A Comprehensive Study of the Bitcoin P2P Network. In 3rd Conference on Blockchain Research and Applications for Innovative Networks and Services (BRAINS 2021), page 8, Paris/ Virtuel, France, September 2021. IEEE.

[94] Edward Chuah, Neeraj Suri, Arshad Jhumka, and Samantha Alt. Challenges in identifying network attacks using netflow data. In The 20th IEEE International Symposium on Network Computing and Applications (NCA), 2021.

[95] Bruno Dzogovic, Tariq Mahmood, Bernardo Santos, Boning Feng, Van Thuan Do, Niels Jacot, and Thanh van Do. Advanced 5G Network Slicing Isolation Using Enhanced VPN+ for Healthcare Verticals. In 7th EAI International Conference on Smart Objects and Technologies for Social Good (GOODTECHS), 2021.

[96] Martina Šestak, Marjan Heričko, Tatjana Welzer Družovec, and Muhamed Turkanović. Applying k-vertex cardinality constraints on a neo4j graph database. Future Generation Computer Systems, 115:459–474, 2021.

[97] Christian Rondanini, Federico Daidone, Barbara Carminati, and Elena Ferrari. Blockchain-based controlled information sharing in inter-organizational workflows. In Proceeding of International Conference on Services Computing (SCC 2020), 2020.

[98] Anh-Tu Hoang, Barbara Carminati, and Elena Ferrari. Cluster-based anonymization of directed graphs. In 2019 IEEE 5th International Conference on Collaboration and Internet Computing (CIC), pages 91–100, 2019.

Appendices

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, whether the paper is available as open access, and the ranking of the publication (the SJR 2021 for journals and the CORE 2021 ranking for conferences). Note that some conferences are not ranked and that some publications are book chapters that are not ranked.

It should be noted that CONCORDIA has entered into the EU-ECAS system even more papers than what is shown in the list below. There are several reasons for that, including that white papers that were not peer-reviewed have been excluded from this Appendix, although they were found by the OpenAire system and therefore included into ECAS.

No	Task	Partner	Year	Title	OA	Rank
1	T1.1	JUB	2021	Applying Metamorphic Testing to Homomorphic Cryptography [4]	No	
2	T1.1	FORTH	2021	Hard edges: Hardware-based Control-Flow Integrity for Embedded Devices [3]		
3	T1.1	UZH	2021	Architectures for Blockchain-IoT Integration [12]	No	
4	T1.1	UZH	2021	Python-Based TinyIPFIX in Wireless Sensor Networks [13]	Yes	B
5	T1.1	UZH	2021	IoT-based Access Management Supported by AI and Blockchains [25]	Yes	B
6	T1.1	UZH	2021	Security and Standardization of a Notary-based Blockchain Interoperability API [22]	No	
7	T1.1	UZH	2021	Edge2BC: a Practical Approach for Edge-to-Blockchain IoT Transactions [21]	No	
8	T1.1	UI	2021	Regulating data sharing across MQTT environments [86]	Yes	2.193
9	T1.1	UP/ISI	2021	Machine Learning Attacks and Countermeasures on Hardware Binary Edwards Curve Scalar Multipliers [5]	Yes	0.965
10	T1.1	UP/ISI	2021	Studying OpenCL-based Number Theoretic Transform for heterogeneous platforms [7]	No	
11	T1.1	ISI	2021	High-Level Synthesis design approach for Number-Theoretic Transform Implementations [8]	No	
12	T1.1	UMIL	2021	From Trustworthy Data to Trustworthy IoT: A Data Collection Methodology Based on Blockchain [9]	No	0.979
13	T1.1	UMIL	2021	Security Certification Scheme for Content-centric Networks [10]	No	B
14	T1.1	UMIL	2021	Towards an Assurance Framework for Edge and IoT Systems [11]	No	
15	T1.1	BGU	2021	Poisoning attacks on cyberattack detectors for industrial control systems [15]	No	B

Continued on next page . . .

No	Task	Partner	Year	Title	OA	Rank
16	T1.1	BGU	2020	MaskDGA: An evasion attack against dga classifiers and adversarial defenses [17]	Yes	0.927
17	T1.1	BGU	2021	Deepstream: Autoencoder-based stream temporal clustering and anomaly detection [18]	No	1.726
18	T1.1	BGU	2021	Deepstream: autoencoder-based stream temporal clustering [19]	No	B
19	T1.1	BGU	2021	Blockchain-based bug bounty framework [26]	No	B
20	T1.2	FORTH	2021	A Survey on Encrypted Network Traffic Analysis Applications, Techniques, and Countermeasures [31]		5.090
21	T1.2	FORTH	2021	Acceleration of Intrusion Detection in Encrypted Network Traffic Using Heterogeneous Hardware [32]		0.803
22	T1.2	UT	2021	The Far Side of DNS Amplification: Tracing the DDoS Attack Ecosystem from the Internet Core [87]		A
23	T1.2	UT	2021	ANYway: Measuring the Amplification DDoS Potential of Domains [30]		B
24	T1.2	UT	2021	Characterization of Anycast Adoption in the DNS Authoritative Infrastructure [29]		
25	T1.2	UT	2021	Tangled: A Cooperative Anycast Testbed [88]		B
26	T1.2	UT	2021	Forecasting the Impact of IXP Outages Using Anycast [89]		
27	T1.2	UT	2021	Addressing the Challenges of Modern DNS A Comprehensive Tutorial [27]		
28	T1.2	UZH	2021	Policy-based Blockchain Selection [34]	No	5.147
29	T1.2	UZH	2021	SecGrid: a Visual System for the Analysis and ML-based Classification of Cyberattack Traffic [90]		B
30	T1.2	UZH	2021	Blockchains and Distributed Ledgers Uncovered: Clarifications, Achievements, and Open Issues [91]	No	

Continued on next page . . .

No	Task	Partner	Year	Title	OA	Rank
31	T1.2	UL	2021	An open measurement dataset on the Bitcoin P2P Network [92]	Yes	B
32	T1.2	UL	2021	A Comprehensive Study of the Bitcoin P2P Network [93]	Yes	
33	T1.2	RISE, ICL	2021	Non-IID Data Re-balancing at IoT Edge with Peer-to-peer Federated Learning for Anomaly Detection [20]	Yes	
34	T1.2	MUNI	2021	Enriching DNS Flows with Host-Based Events to Bypass Future Protocol Encryption [33]	No	
35	T1.3	ULANC	2021	Fast Kernel Error Propagation Analysis in Virtualized Environments [36]	Yes	A
36	T1.3	ULANC	2021	Failure Diagnosis for Cluster Systems using Partial Correlations [35]	Yes	C
37	T1.3	ULANC	2021	PCaaD: Towards automated determination and exploitation of industrial systems [40]	Yes	1.726
38	T1.3	ULANC	2021	Challenges in Identifying Network Attacks Using Netflow Data [94]	Yes	B
39	T1.3	UL	2021	Extracting Executable Payloads From Packed Malware [37]	Yes	A*
40	T1.4	FORTH	2021	Securing Runtime Memory via MMU manipulation [43]		
41	T1.4	FORTH	2021	MMU-based Access Control for Libraries [44]		B
42	T1.4	TUBS	2021	Intrusion Response System for Vehicles: Challenges and Vision [45]		
43	T1.4	TUBS	2021	SPPS: Secure Policy-based Publish/Subscribe System for V2C Communication [46]		B
44	T1.4	OsloMet/Telenor	2021	A Secure 5G Eldercare Solution Using Millimeterwave Sensors [47]		
45	T1.4	OsloMet/Telenor	2021	Optimizing 5G VPN+ Transport Networks with Vector Packet Processing and FPGA Cryptographic Offloading [49]		
46	T1.4	OsloMet/Telenor	2021	Anomaly Detection in Cellular IoT with Machine Learning [50]		
47	T1.4	OsloMet/Telenor	2021	Ageing@home: A secure 5G welfare technology solution for elderlies [48]		
48	T1.4	OsloMet/Telenor	2021	Advanced 5G Network Slicing Isolation Using Enhanced VPN+ for Healthcare Verticals [95]		

Continued on next page . . .

No	Task	Partner	Year	Title	OA	Rank
49	T1.4	UM	2021	Applying k-vertex cardinality constraints on a Neo4j graph database [96]		2.233
50	T1.4	UM	2021	Synergy of Blockchain Technology and Data Mining Techniques for Anomaly Detection [52]		0.507
51	T1.4	UL	2021	Towards Automating Security Enhancement for Cloud Services [53]	Yes	B
52	T1.4	UL	2021	Automated Orchestration of Security Chains Driven by Process Learning [54]	No	
53	T1.4	UMIL	2021	An Assurance-Based Risk Management Framework for Distributed Systems [55]		A
54	T1.4	UMIL	2021	An Assurance Framework and Process for Hybrid Systems [56]		
55	T1.4	UMIL	2021	Blockchain-Based Reputation Systems: Implementation Challenges and Mitigation [57]	Yes	0.59
56	T1.4	UI	2020	Blockchain-based controlled information sharing in inter-organizational workflows [97]	Yes	B
57	T1.4	UI	2021	Blockchain-based Privacy Enforcement in the IoT domain [58]	Yes	2.265
58	T1.5	UMIL	2021	Big Data Analytics-as-a-Service: Bridging the gap between security experts and data scientists [70]	Yes	1.112
59	T1.5	TID+FORTH	2021	YourAdvalue: Measuring Advertising Price Dynamics without Bankrupting User Privacy [60]	Yes	
60	T1.5	TID+FORTH	2021	User Tracking in the Post-Cookie Era: How Websites Bypass GDPR Consent to Track Users [61]	Yes	A*
61	T1.5	TID+FORTH	2021	The Rise and Fall of Fake News Sites: A Traffic Analysis [62]	Yes	
62	T1.5	TID+ICL	2021	PPFL: Privacy-Preserving Federated Learning with Trusted Execution Environments [59]	Yes	B
63	T1.5	TID	2021	YouTubing at Home: Media Sharing Behavior Change as Proxy for Mobility Around COVID-19 Lockdowns [63]	Yes	

Continued on next page . . .

No	Task	Partner	Year	Title	OA	Rank
64	T1.5	TID	2021	Differential Tracking Across Topical Webpages of Indian News Media [64]	Yes	
65	T1.5	TID	2021	Under the Spotlight: Web Tracking in Indian Partisan News Websites [65]	Yes	
66	T1.5	TID	2021	On the Aggression Diffusion Modeling and Minimization in Twitter [66]	Yes	0.839
67	T1.5	TID	2021	Modeling aggression propagation on social media [67]	Yes	0.929
68	T1.5	TID	2021	Exploring Content Moderation in the Decentralised Web: The Pleroma Case [68]	Yes	A
69	T1.5	UI	2019	Cluster-based anonymization of directed graphs [98]	Yes	
70	T1.5	UI	2020	Cluster-based anonymization of knowledge graphs [71]	Yes	B
71	T1.5	UI	2021	Privacy-Preserving Sequential Publishing of Knowledge Graphs [72]	Yes	A*
72	T1.5	UI	2021	A Risk Assessment Mechanism for Android Apps [73]	Yes	
73	T1.5	CUT	2022	A Unified Graph-Based Approach to Disinformation Detection using Contextual and Semantic Relations [80]	Yes	
74	T1.5	CUT	2021	”How over is It?” Understanding the Incel Community on YouTube [82]	Yes	
75	T1.5	CUT	2022	”It is just a flu”: Assessing the Effect of Watch History on YouTube’s Pseudoscientific Video Recommendations [81]	Yes	
76	T1.5	FORTH	2021	Misusing Mobile Sensors for Stealthy Data Exfiltration [69]	Yes	A*
77	T1.5	UZH	2021	”ITrade: A Blockchain-based, Self-Sovereign, and Scalable Marketplace for IoT Data Streams [75]	No	B

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 order by date.

No	**Scientific Event** Name	**Abbr.** Partner	**Where**	**When** Role	**URL**
1	**2021 ACM Conference on Data and Application Security and Privacy**	**CODASPY 2021**	**Baltimore, USA**	**22-24 March 2021**	http://www.codaspy.org/2021/
	Name: Elena Ferrari	Partner:	UI	Role:	Workshop Chair
	Name: Barbara Carminati	Partner:	UI	Role:	Program Co-chair
2	**Workshop on Security Aspects in Autonomous Systems**	**Lancaster 2021**	**Virtual Workshop**	**29 March 2021**	https://tas-security.lancs.ac.uk/esg1/
	Name: Neeraj Suri	Partner:	ULANC	Role:	Organizing Committee
3	**Secure Runtime Environments session at HiPEAC CSW Spring 2021**	**HiPEAC CSW Spring 2021**	**Virtual Session**	**8 April 2021**	https://www.hipeac.net/csw/2021/spring-webinars/#/program/sessions/7889/
	Name: Marinos Tsantekidis	Partner:	FORTH	Role:	Organizer
4	**IFIP/IEEE International Symposium on Integrated Network Management 2021**	**IFIP/IEEE IM 2021**	**Virtual Conference**	**17-21 May 2021**	https://im2021.ieee-im.org/
	Name: Olivier Festor	Partner:	UL	Role:	General Co-chair
	Name: Rémi Badonnel	Partner:	UL	Role:	Experience Co-chair

Continued on next page...

No	**Scientific Event** Name	**Abbr.** Partner	**Where**	**When** Role	**URL**
5	**1st Workshop on Security and Privacy for Mobile AI 2021**	**ACM MAISP 2021**	**Virtual Workshop**	**24 June 2021**	https://maisp.gitlab.io
	Name: Nicolas Kourtellis	Partner:	TID	Role:	Steering Committee
6	**ACM SIGCOMM 2021 Workshop on Technologies, Applications, and Uses of a Responsible Internet**	**TAURIN 2021**	**Virtual Conference**	**23 August 2021**	https://conferences.sigcomm.org/sigcomm/2021/workshop-taurin.html
	Name: Ralph Holz Name: Cristian Hesselman	Partner: UT Partner: SIDN		Role: Organizing Committee Role: Organizing Committee	
7	**13th IEEE International Conference Web Services**	**IEEE ICWS 2021**	**Chicago, IL, USA**	**5-10 September 2021**	https://conferences.computer.org/icws/2021/
	Name: Elena Ferrari	Partner:	UMIL	Role:	General Co-chair
8	**IEEE International Conference on Cloud Computing**	**IEEE CLOUD 2021**	**Virtual Conference**	**5-11 September 2021**	https://conferences.computer.org/cloud/2021/
	Name: Claudio Ardagna	Partner:	UMIL	Role:	Program Co-chair

Continued on next page...

No	Scientific Event Name	Abbr. Partner	Where	When Role	URL
9	**European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases**	**ECML PKDD 2021**	**Virtual Conference**	**13-17 September 2021**	https://2021.ecmlpkdd.org
	Name: Nicolas Kourtellis	Partner:	TID	Role:	Applied Data Science Track Program Chairs
10	**The Future of Cybersecurity in Slovenia and Europe**		**Ljubljana, Slovenia (hybrid)**	**4 October 2021**	https://eu2021.dihslovenia.si/en/events/the-future-of-cybersecurity-in-slovenia-and-europe/
	Name: Tatjana Welzer Name: Lili Nemec Zlatolas	Partner: UM Partner: UM		Role: Organizer Role: Organizer	
11	**17th International Conference on Network and Service Management**	**CNSM 2021**	**Izmir, Turkey**	**25-29 October 2021**	http://www.cnsm-conf.org/2021/
	Name: Rémi Badonnel	Partner:	UL	Role:	Publicity Co-chair

Continued on next page...

No	**Scientific Event** Name	**Abbr.** Partner	**Where**	**When** Role	**URL**
12	**The 29th IEEE International Conference on Network Protocols**	**IEEE ICNP 2021**	**Dallas, Texas, USA 2021 (Virtual Event)**	**1-5 November 2021**	https://icnp21.cs.ucr.edu/
	Name: Michael Sirivianos	Partner:	CUT	Role:	Area Chair
13	**3rd IEEE International Conference on Trust, Privacy and Security in Intelligent Systems and Applications**	**IEEE TPS-ISA 2021**	**Virtual Conference**	**13-15 December 2021**	http://www.sis.pitt.edu/lersais/conference/tps/2021/
	Name: Elena Ferrari	Partner:	UI	Role:	Program Co-chair

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 order by date.

No	**Scientific Event** Name	**Abbr.**	**Where**	**When** Partner	**URL**
1	**16th Dependable and Adaptive Distributed Systems, Track of the 36th ACM Symposium on Applied Computing**	**DADS 2021**	**Gwangju, Korea**	**22-26 March 2021**	https://www.dedisys.org/sac21/
	Name:	Claudio Ardagna		Partner:	UMIL
2	**11th IFIP International Conference on New Technologies, Mobility and Security - Security Track**	**NTMS 2020**	**Paris, France**	**19-21 April 2021**	https://www.ntms-conf.org/ntms2021/
	Name:	Claudio Ardagna		Partner:	UMIL
3	**The Web Conference**	**WWW 2021**	**Virtual Conference**	**19-23 April 2021**	https://www2021.thewebconf.org/
	Name:	Neeraj Suri		Partner:	ULANC
4	**The European Conference on Computer Systems**	**EuroSys 2021**	**Virtual Conference**	**26-28 April 2021**	https://2021.eurosys.org/
	Name:	Neeraj Suri		Partner:	ULANC
5	**3rd IEEE Conference on Blockchain and Cryptocurrency**	**ICBC 2021**	**Virtual Conference**	**3-6 May 2021**	https://icbc2021.ieee-icbc.org/

Continued on next page...

No	**Scientific Event** Name	**Abbr.**	**Where**	**When** Partner	**URL**
	Name:	Thibault Cholez		Partner:	UL
6	**IFIP/IEEE International Symposium on Integrated Network Management 2021**	**IFIP/IEEE IM 2021**	**Virtual Conference**	**17-21 May 2021**	https://im2021.ieee-im.org/
	Name:	Anna Sperotto		Partner:	UT
	Name:	Jürgen Schönwälder		Partner:	JUB
	Name:	Thibault Cholez		Partner:	UL
	Name:	Emil Lupu		Partner:	ICL
7	**International Conference on Computing, Networks and Internet of Things**	**CNIOT 2021**	**Beijing, China**	**20-22 May 2021**	http://www.cniot2021.net
	Name:	Nicola Bena		Partner:	UMIL
8	**6th International Workshop on Traffic Measurements for Cybersecurity**	**WTMC 2021**	**Virtual Workshop**	**25 May 2021**	https://wtmc.info/index.html
	Name:	Roland van Rijswijk		Partner:	UT
	Name:	Jair Santanna		Partner:	UT
	Name:	Moritz Müller		Partner:	SIDN
	Name:	Giovane Moura		Partner:	SIDN
9	**15th International Conference on Web and Social Media**	**AAAI ICWSM-2021**	**Virtual Conference**	**7-10 June 2021**	https://icwsm.org/2021/

Continued on next page...

No	**Scientific Event** Name	**Abbr.**	**Where**	**When** Partner	**URL**
	Name:	Michael Sirivianos		Partner:	CUT
10	**2021 IEEE International Conference on Communications - Communication and Information Systems Security Symposium (CISS Symposium)**	**ICC 2021**	**Montreal, Canada**	**14-23 June 2021**	https://icc2021.ieee-icc.org
	Name:	Claudio Ardagna		Partner:	UMIL
11	**ACM International Conference on Management of Data**	**SIGMOD 2021**	**Xian, China**	**20-25 June 2021**	https://2021.sigmod.org/
	Name:	Elena Ferrari		Partner:	UI
12	**19th International Conference on Applied Cryptography and Network Security**	**ACNS 2021**	**Virtual Conference**	**21-24 June 2021**	https://sulab-sever.u-aizu.ac.jp/ACNS2021/
	Name:	Emil Lupu		Partner:	ICL
13	**IFIP Networking**	**IFIP NETWORKING 2021**	**Virtual Presentations**	**21-24 June 2021**	https://networking.ifip.org/2021/
	Name:	Jürgen Schönwälder		Partner:	JUB

Continued on next page...

No	**Scientific Event** Name	**Abbr.**	**Where**	**When** Partner	**URL**
	Name:	Ralph Holz		Partner:	UT
14	**IFIP Security**	**IFIP SEC 2021**	**Virtual Presentations**	**22-24 June 2021**	https://ifipsec.org/2021/
	Name: Name:	Jürgen Schönwälder Claudio Ardagna		Partner: Partner:	JUB UMIL
15	**7th IEEE International Conference on Network Softwarization**	**Netsoft 2021**	**Virtual Conference**	**28 June - 2 July 2021**	https://netsoft2021.ieee-netsoft.org/
	Name:	Olivier Festor		Partner:	UL
16	**IEEE International Conference on Distributed Computing Systems**	**ICDCS 2021**	**Virtual Conference**	**7-10 July 2021**	https://icdcs2021.us/cfp-html/
	Name:	Neeraj Suri		Partner:	ULANC
17	**2021 IEEE International Conference on Cyber-Security and Resilience**	**IEEE CSR 2020**	**Rhodes, Greece**	**26-28 July 2021**	https://www.ieee-csr.org/
	Name:	Claudio Ardagna		Partner:	UMIL
18	**ACM Special Interest Group on Knowledge Discovery and Data Mining**	**ACM SIGKDD 2021**	**Virtual**	**14-18 August 2021**	https://kdd.org/kdd2021/

Continued on next page...

No	**Scientific Event** Name	**Abbr.**	**Where**	**When** Partner	**URL**
	Name:	Nicolas Kourtellis		Partner:	TID
19	**International Conference on Very Large Databases**	**PVLDB 2021**	**Copenhagen, Denmark (hybrid)**	**16-20 August 2021**	https://vldb.org/pvldb/reproducibility/
	Name:	Elena Ferrari		Partner:	UI
20	**International Workshop on Big Data And IoT Security in Smart Computing During SMARTCOMP**	**IEEE BITS 2021**	**Virtual Conference**	**23 August 2021**	https://www.yama.info.waseda.ac.jp/bits2021/
	Name:	Nicola Bena		Partner:	UMIL
21	**ACM SIGCOMM 2021 Workshop on Technologies, Applications, and Uses of a Responsible Internet**	**TAURIN 2021**	**Virtual Workshop**	**23 August 2021**	https://conferences.sigcomm.org/sigcomm/2021/workshop-taurin.html
	Name:	Mattijs Jonker		Partner:	UT
22	**European Conference on Advances in Databases and Information Systems**	**ADBIS 2021**	**Tartu, Estonia**	**24-26 August 2021**	https://adbis2021.cs.ut.ee/adbis-committees/
	Name:	Tatjana Welzer		Partner:	UM

Continued on next page...

No	**Scientific Event** Name	**Abbr.**	**Where**	**When** Partner	**URL**
23	**30th Annual Conference of the European Association for Education in Electrical and Information Engineering**	**EAEEIE 2021**	**Prague, Czech Republic**	**1–4 September 2021**	https://eaeeie.cvut.cz/committees/
	Name:	Tatjana Welzer		Partner:	UM
24	**24th International Conference on Database Technology**	**EDBT 2021**	**Nicosia, Cyprus**	**1–4 September 2021**	https://edbticdt2021.cs.ucy.ac.cy/
	Name:	Elena Ferrari		Partner:	UI
25	**IEEE International Conference on Cloud Computing**	**IEEE CLOUD 2021**	**Virtual Conference**	**5-11 September 2021**	https://conferences.computer.org/cloud/2021/
	Name:	Nicola Bena		Partner:	UMIL
26	**Network Traffic Measurement and Analysis Conference**	**TMA 2021**	**Virtual Conference**	**14-15 September 2021**	https://tma.ifip.org/2021
	Name: Name:	Roland van Rijswijk Ralph Holz		Partner: Partner:	UT UT
27	**3rd conference on Blockchain Research and Applications for Innovative Networks and Services**	**BRAINS 2021**	**Virtual Conference**	**27-30 September 2021**	https://brains.dnac.org/

Continued on next page...

No	**Scientific Event** Name	**Abbr.**	**Where**	**When** Partner	**URL**
	Name:	Thibault Cholez		Partner:	UL
28	**4th International Workshop on Attacks and Defenses for Internet-of-Things**	**ADIoT 2021**	**Darmstadt, Germany**	**4-8 October 2021**	https://adiot.compute.dtu.dk/2021/
	Name:	Claudio Ardagna		Partner:	UMIL
29	**25th IEEE International Conference on Enterprise Distributed Object Computing**	**EDOC 2021**	**Virtual Conference**	**25-29 October 2021**	http://ieee-edoc.org/2021/
	Name:	Emil Lupu		Partner:	ICL
30	**17th International Conference on Network and Service Management**	**CNSM 2021**	**Virtual Conference**	**25-29 October 2021**	http://www.cnsm-conf.org/2021/
	Name:	Mattijs Jonker		Partner:	UT
	Name:	Jürgen Schönwälder		Partner:	JUB
	Name:	Jair Santanna		Partner:	UT
	Name:	Roland van Rijswijk-Deij		Partner:	UT
	Name:	Olivier Festor		Partner:	UL
	Name:	Emil Lupu		Partner:	ICL

Continued on next page...

No	**Scientific Event** Name	**Abbr.**	**Where**	**When** Partner	**URL**
31	**3rd International Workshop on High-Precision, Predictable, and Low-Latency Networking**	**HiP-Net 2021**	**Virtual Conference**	**29 October 2021**	http://www.cnsm-conf.org/2021/workshop_HiPNet.html
	Name:	Thibault Cholez		Partner:	UL
32	**The 29th IEEE International Conference on Network Protocols**	**IEEE ICNP 2021**	**Dallas, Texas, USA (Virtual Event)**	**1-5 November 2021**	https://icnp21.cs.ucr.edu/
	Name:	Michael Sirivianos		Partner:	CUT
33	**2021 Internet Measurement Conference**	**IMC 2021**	**Virtual Conference**	**2-4 November 2021**	https://conferences.sigcomm.org/imc/2021/
	Name: Name:	Mattijs Jonker Ralph Holz		Partner: Partner:	UT UT
34	**IFIP Internet of Things**	**IFIP IOT 2021**	**Virtual Presentations**	**4-5 November 2021**	http://ifip-iotconference.org/
	Name:	Jürgen Schönwälder		Partner:	JUB
35	**IEEE/ACM Conference on Advances in Social Networks Analysis and Mining**	**ASONAM 2021**	**Virtual Conference**	**8-11 November 2021**	https://asonam.cpsc.ucalgary.ca/2021/

Continued on next page...

No	**Scientific Event** Name	**Abbr.**	**Where**	**When** Partner	**URL**
	Name:	Elena Ferrari		Partner:	UI
36	**ACM/IRTF Applied Networking Research Prize**	**ANRP 2021**	**Virtual Presentations**	**8-12 November 2021**	https://www.ietf.org/how/meetings/112/
	Name:	Anna Sperotto		Partner:	UT
37	**2021 Conference on Computer and Communications Security**	**CCS 2021**	**Virtual Conference**	**15-19 November 2021**	https://www.sigsac.org/ccs/CCS2021/
	Name:	Mattijs Jonker		Partner:	UT
38	**2nd Joint Workshop on CPS and IoT Security and Privacy**	**CPSIoTSec 2021**	**Virtual Conference**	**15 November 2021**	https://cpsiotsec.github.io
	Name:	Emil Lupu		Partner:	ICL
39	**DNS and Internet Naming Research Directions 2021**	**DINR 2021**	**Virtual Workshop**	**16 November 2021**	https://ant.isi.edu/events/dinr2021/
	Name:	Moritz Müller		Partner:	SIDN

Continued on next page...

No	**Scientific Event** Name	**Abbr.**	**Where**	**When** Partner	**URL**
40	**IEEE Uruguay Conference**	**IEEE URUCON 2020**	**Montevideo, Uruguay**	**24-26 November 2021**	http://urucon2021.org/index.html
	Name:	Abhilash Hota		Partner:	JUB
41	**17th International Conference on Information Systems Security**	**SECITC 2021**	**Virtual Conference**	**25-26 November 2021**	http://secitc.eu/
	Name:	Claudio Ardagna		Partner:	UMIL
42	**14th IEEE/ACM International Conference on Utility and Cloud Computing**	**UCC 2021**	**Leicester, UK**	**6-9 December 2021**	https://www.cs.le.ac.uk/events/UCC2021/
	Name:	Claudio Ardagna		Partner:	UMIL
43	**The 17th International Conference on emerging Networking EXperiments and Technologies – Artifact Evaluation**	**CoNEXT 2021 (AeC)**	**Virtual Conference**	**7-10 December 2021**	https://conferences2.sigcomm.org/co-next/2021/
	Name: Name: Name:	Mattijs Jonker Raffaele Sommese Moritz Müller		Partner: Partner: Partner:	UT UT SIDN

Continued on next page...

No	**Scientific Event** Name	**Abbr.**	**Where**	**When** Partner	**URL**
44	**2021 IEEE Global Communications Conference: Communication and Information System Security**	**GLOBECOM CISS 2021**	**Madrid, Spain**	**7-11 December 2021**	https://globecom2021.ieee-globecom.org/
	Name:	Claudio Ardagna		Partner:	UMIL
45	**The 14th International Symposium on Foundations and Practice of Security**	**FPS 2021**	**Hybrid Conference**	**8-10 December 2021**	https://www.fps-2021.com
	Name:	Jean-Yves Marion		Partner:	UL
46	**IEEE/WIC/ACM International Conference on Web Intelligence and Intelligent Agent Technology**	**WI-IAT 2021**	**Melbourne, Australia (hybrid)**	**14-17 December 2021**	https://www.wi-iat.com/wi-iat2021/index.html
	Name:	Elena Ferrari		Partner:	UI
47	**17th International Conference on Information Systems Security**	**ICISS 2021**	**Patna, India**	**16-20 December 2021**	https://www.iitp.ac.in/~iciss2021/
	Name:	Claudio Ardagna		Partner:	UMIL

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. The table is organized in order by publisher.

No	**Description** Name	**Publisher** Partner	**URL** Role
1	**ACM Transactions on Data Science**	**ACM**	https://tds.acm.org/editorial.cfm
	Name: Elena Ferrari	Partner: UI	Role: Associate Editor
2	**ACM Transactions on Privacy and Security**	**ACM**	https://dl.acm.org/journal/tops
	Name: Elena Ferrari	Partner: UI	Role: Associate Editor
3	**ACM SIGCOMM Computer Communications Review**	**ACM**	https://ccronline.sigcomm.org/
	Name: Anna Sperotto Name: Ralph Holz	Partner: UT Partner: UT	Role: Editorial board member Role: Editorial board member
4	**IEEE Transactions on Big Data**	**IEEE**	https://ieeexplore.ieee.org/xpl/RecentIssue.jsp?punumber=6687317
	Name: Neeraj Suri	Partner: ULANC	Role: Associate Editor
5	**Cybersecurity and Privacy of Frontiers in Big Data**	**Frontiers Media SA**	https://www.frontiersin.org/journals/big-data/sections/cybersecurity-and-privacy
	Name: Elena Ferrari	Partner: UI	Role: Specialty Chief Editor
6	**Data Science and Engineering**	**Springer**	https://www.springer.com/journal/41019

Continued on next page...

No	**Description** Name	**Publisher** Partner	**URL** Role
	Name: Elena Ferrari	Partner: UI	Role: Associate Editor
7	**IEEE Transactions on Cloud Computing**	**IEEE**	https://ieeexplore.ieee.org/xpl/RecentIssue.jsp?punumber=6245519
	Name: Neeraj Suri	Partner: ULANC	Role: Associate Editor
8	**IEEE Transactions on Network and Service Management**	**IEEE**	https://www.comsoc.org/publications/journals/ieee-tnsm
	Name: Jürgen Schönwälder	Partner: JUB	Role: Editorial board member
9	**Digital Trust: Trust Management in the Cyberspace, IEEE Internet computing**	**IEEE**	https://www.computer.org/digital-library/magazines/ic/call-for-papers-special-issue-on-digital-trust-trust-management-in-the-cyberspace
	Name: Elena Ferrari	Partner: UI	Role: Special Issue Editor
10	**IEEE Internet Computing**	**IEEE**	https://www.computer.org/csdl/magazine/ic/about/15624?title=Editorial%20Board&periodical=IEEE%20Internet%20Computing
	Name: Elena Ferrari	Partner: UI	Role: Associate Editor in Chief

Continued on next page...

No	**Description** Name	**Publisher** Partner	**URL** Role
11	**IEEE Transactions on Service Computing**	**IEEE**	https://www.computer.org/csdl/journal/sc/misc/14407?title=About&periodical=IEEE%20Transactions%20on%20Services%20Computing
	Name: Elena Ferrari	Partner: UI	Role: Associate Editor
12	**International Journal of Cooperative Information Systems**	**World Scientific**	https://www.worldscientific.com/page/ijcis/editorial-board
	Name: Elena Ferrari	Partner: UI	Role: Associate Editor
13	**Journal of Network and Systems Management**	**Springer**	https://www.springer.com/journal/10922
	Name: Anna Sperotto Name: Jürgen Schönwälder Name: Ralph Holz Name: Emil Lupu	Partner: UT Partner: JUB Partner: UT Partner: ICL	Role: Editorial board member Role: Editorial board member Role: Editorial board member Role: Editorial board member
14	**International Journal of Network Management**	**Wiley**	https://onlinelibrary.wiley.com/journal/10991190
	Name: Anna Sperotto Name: Jürgen Schönwälder Name: Emil Lupu	Partner: UT Partner: JUB Partner: ICL	Role: Editorial board member Role: Editorial board member Role: Editorial board member

Continued on next page...

No	**Description** Name	**Publisher** Partner	**URL** Role
15	**Synthesis Lectures on Security, Privacy and Trust**	**Morgan and Claypool**	https://www.morganclaypool.com/toc/spt/1/1
	Name: Elena Ferrari	Partner: UI	Role: Editor