VR 2019 1312 DOC.1235/5

Cybersecurity Initiative Flanders

Strategic Research Programme

July 2019

Contributing Authors: Bart De Decker, Bart Jacobs, Bart Preneel, Benedikt Gierlichs, Bert Lagaisse, Bjorn De Sutter, Bruno Crispo, Claudia Diaz, Coen De Roover, Cyprien Delpech de Saint Guilhem, Danny De Cock, Danny Hughes, Dave Singelee, Davy Preuveneers, Dimitri Van Landuyt, Dominique Devriese, Elena Andreeva, Elisa Gonzales Boix , Els Kindt, Emmanuela Orsini, Enrique Argones Rua, Frank Piessens, Frederik Vercauteren, Ingrid Verbauwhede, Jan Tobias Mühlberg, Koen Yskout, Lieven Desmet, Nele Mentens, Nigel Smart, Peggy Valcke, Pieter Maene, Sam Michiels, Stijn Volckaert, Svetla Nikova, Vincent Naessens, Vincent Rijmen, Wolfgang De Meuter, Wouter Castryck, Wouter Joosen

1. Introduction

1.1. Context and Motivation: Cybersecurity, Challenge and Opportunity

While the digital transformation has a strongly increasing and positive impact on our society and our economy, the lack of adequate cybersecurity in our systems, platforms and services can lead to major dangers, risks and problems. More and more information is being collected and analyzed, leading to significant efficiency gains and new applications. In cyber-physical systems, this results in far-reaching automation with, among other things, autonomous robots, cars and drones. The entire infrastructure of society is also being transformed; we get smarter cities (smart cities), smart transport systems (smart transport) and smart electricity grids (smart grids), smarter hospital facilities, etc.

This transformation affects all sectors, both within the government (general policy, education and health, infrastructure, police, defense) and within the private industry: critical infrastructure, transport, manufacturing, financial sector, media, health sector. They are and will all be more competitive and strengthened thanks to digital transformations, but they are also vulnerable. In addition to the far-reaching and limitless possibilities, this digital transformation indeed brings important new risks. Addressing these cybersecurity risks will be essential for economic success in the forthcoming decades.

The cybersecurity risks continuously increase because of the following evolution:

(1) As *cybercrime* becomes more and more attractive and rewarding for malicious organizations, its impact grows with the size of the digital economy. Consequently, the attacks and attackers become more specialized. The cybercrime threat has evolved from break-ins carried out by individuals with simple tools, to sophisticated attacks carried out by organized crime, hacktivists, specialized companies, and nation states. This is not restricted to the passive collection of information (for industrial espionage, among other things) but also actively hacking into systems and creating physical damage (e.g., Wannacry, and attacks on the electricity network in Ukraine in 2015), as well as global hybrid threats. A complex ecosystem has developed in which malicious actors specialize in various aspects of cybercrime.

2) Cybercriminals exploit *large scale attack infrastructures*. By using networks and mechanisms for automatic distribution, it becomes feasible to attack a large number of systems from any location (e.g., the 2016 Mirai botnet). Attribution is very complex. There is also an important problem of proliferation: attacks with sophisticated malware executed by nation states may leave traces. Subsequently, organized crime or other nation states can exploit such malware.

(3) The digital platforms, services and assets that we need to protect also increase in *complexity* – which obviously makes protection more difficult. Modern digital services are very complex and interdependent systems, created through a complex and international supply chain. This means that it is impossible to make such systems perfectly secure: due to the complexity and the dynamism there are always minor errors, which necessitates regular updates that require a complex governance. At the same time, many systems contain deliberate loopholes.

(4) The *societal challenge* goes far beyond the battle between attackers and defenders/guardians. The explosion of the use of computer systems and networks (smartphones, smart cameras, industrial IoT, implantable medical devices, ...) and the sharply decreasing cost of collecting and processing information result in important new privacy risks. For example, the number of data leaks increases exponentially[1]. These risks are not only a threat to human rights (the right to privacy is recognized by Article 8 of the European Convention on Human Rights and Article 7 of the Charter of Fundamental Rights of the European Union), but also the democracy and the rule of law (interference of nation states in elections via social media). In addition, they also pose a security problem: data protection is essential to guarantee safety for citizens as well as governments.

[1] http://informationisbeautiful.net/visualizations/worlds-biggest-data-breaches-hacks/

Internationally leading cybersecurity competence is available in Flanders (see Annex 1). Flanders can therefore meet this challenge, also in the context of a growing international competition. Fortunately, cybersecurity does not only present a challenge, but it also creates an opportunity with economic benefits.

As a starting point, the creation of a *secure infrastructure* (for transport, communication, data storage, etc.) is essential for the economic development of Flanders and becomes an important competitive advantage. The industry is increasingly evolving towards a model that strongly relies on the combination of online services, software and hardware: cybersecurity becomes an essential component of all products and services, but it is also of great importance to prevent the theft of intellectual property.

To meet the complex challenges, there is a need for intensive cooperation according to the triple helix innovation model: academia, government and industry each have an important role to play.

The challenges mentioned above necessitate investments in research, development and implementation, application, innovation and valorization – the entire chain from basic research to effective services in the market – and far beyond common practice. This is illustrated by the realizations and policy decisions in our neighboring countries.

Specifically for this case of cybersecurity, there are a number of important elements:

1) Studies have shown that there are major *market failures*[2] in the area of cybersecurity and privacy: this means that technology users are not prepared to pay more for more secure solutions and the industry investments are insufficient. This implies that the government must intervene, or at least play a supporting and leading role, by playing an active role in stimulating advanced research and innovation.

A Cybersecurity Programme for Flanders must therefore further develop, manage, distribute and share high-level competence in cybersecurity. In an advanced research program, specialized and critical cybersecurity knowledge must be continuously enhanced and maintained, transformed and proven applicable in collaboration with industry, and disseminated with specialized stakeholders.

2) The cybersecurity problem is an international phenomenon. The Flemish economy relies for a large part on hardware, software and ICT services that are produced abroad. We cannot solve the problems locally / regionally – but we can be a strong player and play a central role in the larger European picture.

A Cybersecurity Programme for Flanders must strengthen the core competences, stay comprehensive ("we leave no flank unprotected") while reaching out to industry to ensure applicability of knowledge and technology, and to operate with an up-to-date prioritization of topics in cybersecurity. At the same time, collaboration and synchronization with other leading labs in Europe will ensure that Flanders invests in its strengths in cybersecurity, while creating synergy and collaboration with other leading centers in Europe thus avoiding unproductive duplication of efforts.

3) A number of studies show that there is an important shortage of experts in cybersecurity in all countries. This seems only the beginning of a dangerous evolution. Where this is a (possibly personal) economic opportunity for the limited group of experts in the short term, this will represent a dramatic social and economic risk in the long term. That is why we advocate strengthening knowledge and expertise in an environment where education and training are part of the core business.

A Cybersecurity Programme for Flanders will therefore invest in a rich portfolio of cybersecurity trainings building an offering for different audiences in terms of existing competence levels, as well as focus areas. This training must be backed by state-of-the-art knowledge and competence made available by the top-level academic groups that are present in Flanders.

The aim of this document is to present at a high-level, the Strategic Programme for Cybersecurity Research in Flanders. This Programme has to deliver impactful solutions to real-world challenges, while starting from and building upon academic excellence. The execution of the Programme has to strengthen existing core competences in cybersecurity research, while delivering building blocks and solutions that will benefit

[2] R. Anderson, R. Boehme,R. Clayton,T. Moore, Security Economics and the Internal Market, January 2008, https://www.enisa.europa.eu/publications/archive/economics-sec/

cybersecurity in industry. At the same time, a strong cybersecurity research program will attract international talent and will enable the creation and extension of cybersecurity training programs.

1.2. A Strategic Research Programme

The Strategic Research Programme includes four Tracks (Figure 1-1).

- Track 1 addresses Application and Software Security and aims to support all stakeholders that analyze, develop and deploy new application software, while using an evolving set of technologies in the context of secure software development.
- Track 2 includes Strategic Security Services, such as authentication, authorization and services for data protection. The overall idea is that many security specific building blocks (reusable components or services that are typically offered as security middleware) will not be built from scratch in new applications, and should be evolving with new demands and expectations – typically reaching beyond the state-of-practice in industry offerings.
- Track 3 covers System and Infrastructure Security. Here one expects stable, secured technology that is packaged as a black box in an operating system or in network layers. Software and service developers rely on the robustness of these lower layers – yet we all know that additional research is essential to meet the promise.
- Track 4 covers the Technology Building Blocks for Security: secure hardware, cryptography and secure cryptographic implementations. Needless to state this is of strategic importance.

Many topics and subdomains of cybersecurity can be considered, possibly emerging from a broad survey of academic literature. The specific approach has been to leverage upon available excellence, and therefore on proven Track record, thus maintaining focus and enabling the rapid development of new research results.

Figure 1-1 sketches a helicopter view on the Strategic Research Programme. The Consortium has chosen to model the work as a technology stack, not so much for the elegance of the research representation per se, but to ensure that the programme can be communicated and shared using a comprehensive overview, that can be presented elegantly to industry stakeholders. In fact, many industry stakeholders have a specific interest in one or two of the proposed Research Tracks that are aligned with their business models.

Figure 1-1 A Strategic Programme with four Research Tracks

A further breakdown of these Tracks has been created by identifying and incorporating research themes that are essential for the proposed programme. The proposed structure is not a goal in its own right; themes in each of the Tracks depend on and interact with themes in other Tracks. These dependencies will be documented in the description of each of the Tracks.

Section 1.3 briefly elaborates on the core principles that have been applied while creating the Strategic Programme.

1.3. Principles

The Strategic Research Programme has been defined based on the following principles.

(1) First, the Consortium will address problems and challenges that have been identified in the cybersecurity and ICT security research communities. While describing the research Tracks and the research themes that belong to such a Track (in Chapters 2 to 5), the proposed Programme refers to relevant research from the state-of-the-art. This background information documents the relevance of the proposed work and highlights the starting points for the research Tracks and themes. Notice though that this document does not have the ambition to present an elaborate survey of the scientific and technical literature.

(2) Secondly, the Consortium is built from research groups with a strong and proven track record. The goal is to strengthen existing teams, groups and activities and make progress as fast as possible and feasible. This perspective has some consequences. The selected research themes and topics of the Programme are – in principle – limited by the available expertise in Flanders – even though this does definitely not yield a narrow scope of the research program. This is not a problem but an opportunity. Especially in the European context, existing partnerships and good relationships with scientific peers enable connecting to the right type of additional expertise when and if needed. Annex 1 to this proposal describes the Consortium in detail, by sketching each of the research groups involved, including the faculty members and permanent research staff who work on cybersecurity.

(3) Last but not least, the research themes are driven by a strong interest in industry and by industry demand. The Research Tracks have been identified and confirmed in dialogue with industry and the relevance of each of the proposed research themes has been confirmed by many industry stakeholders. In particular, an elaborate feedback session with a representative group of industry stakeholders has been organized on April 2, 2019. This one-day workshop was of great value in finishing this proposal.

Related to the third principle, it should be noticed that the value and applicability of research results will be pursued by building quality prototypes that will be combined in industry relevant platforms. The Strategic Research Programme will therefore deliver practical assets that enable exploration and discovery of new security solutions. This will be discussed in Chapter 6.

1.4. Outline of the Document

The further description of the Research Programme is structured as follows. The next chapters (Chapters 2 to 5) each contain a high-level overview of one Research Track:

- Chapter 2 presents the Research Track on Application and Software Security.
- Chapter 3 discusses the Research Track on Strategic Security Services.
- Chapter 4 covers the Research Track on System and Infrastructure Security.
- Chapter 5 addresses the fourth Research Track. It covers Technology Building Blocks for Security: Secure Hardware, Cryptography and Secure Implementations.

Each of the Research Tracks includes a collection of strategic Research Themes that are proposed in light of the principles sketched above. Next the scope of each Research Track is summarized. Subsequently each of the Research Themes is introduced by presenting the state-of-the-art, illustrating industry demand and by introducing the Research Activities (RA) that will be performed by members of the Consortium . The related objectives are presented by defining expected outcomes that will be delivered by the end of one and two years of research activity.

Chapter 6 complements the above by discussing in more detail how the strategic research activities that have been defined in Chapters 2-5 contribute to prototypes and platforms that will facilitate technology transfer and adoption in industry.

The final chapter is a conclusion. It summarizes the headlines of the proposal and discusses some of the important actions being taken while preparing for a start on September 1, 2019.

Three annexes complete this work: Annex 1 is an elaborate overview of the Consortium. In fact, the researchers listed in Annex 1 have all contributed to this document. Annex 2 includes the headlines of the budget breakdown. Annex 3 is an overview of all Tracks, Themes and Research Activities in this Research Programme.

Contents

1. Introduction 3
1.1. Context and Motivation: Cybersecurity, Challenge and Opportunity 3
1.2. A Strategic Research Programme 5
1.3. Principles 6
1.4. Outline of the Document 7
2. Research Track 1: Application and Software Security 11
Scope 11
2.1. Secure SDLC – Secure Software Development Life Cycle 11
2.2. Program Verification and Security Testing 14
2.3. Secure Programming Languages and Secure Compilation 17
2.4. Connections with Other Research Tracks 20
2.5. References 21
3. Research Track 2: Strategic Security Services 25
Scope 25
3.1. Identity Management and Authentication 25
3.2. Authorization and Audit 29
3.3. Advanced Encryption Techniques and Data Access Middleware 32
3.4. Policy and Regulation 36
3.5. Connections with other Research Tracks 40
3.6. References 42
4. Research Track 3: System and Infrastructure Security 49
Scope 49
4.1. System Security 49
4.2. Network Security 54
4.3. Security Monitoring and Management 56
4.4. Connections with other Research Tracks 59
4.5. References 60
5. Research Track 4: Technology Building Blocks: Secure Hardware, Cryptography and Secure Implementations 66
Scope 66
5.1. Secure Hardware: Roots of Trust Anchored into Technology Foundations 66
5.2. Cryptographic Algorithms 69
5.3. Cryptographic Protocols 73
5.4. Secure and Efficient Cryptographic Implementations 78
5.5. Connections with other Research Tracks 82
5.6. References 84
6. Prototypes, Validation and Inroads to Industry Implementation 91
6.1. Illustration of Market Relevance:Three Strategic Industry Sectors for Cybersecurity 91
6.2. Essential Types of Technology Platforms 94
6.3. Leveraging on Research Results: Illustration with Three Business Cases 94

6.4. Delivery & Validation of Technology Assets that Emerge from Research 99
7. Conclusion 101
7.1. Summary of the Programme 101
7.2. Refinement and Operational Plan of the Research Program 102
7.3. Evaluation and Kick-off 103
8. Annexes 105
8.1. Consortium 105
8.2. Budget Headlines 105
8.3. Overview Research Activities 106

2. Research Track 1: Application and Software Security

Contributing Authors: ***Bart Jacobs****, Coen De Roover, Dimitri Van Landuyt, Dominique Devriese, Elisa Gonzales Boix, Frank Piessens, Koen Yskout, Lieven Desmet, Vincent Naessens, Wolfgang De Meuter, Wouter Joosen*

Scope

Many security incidents in ICT start from the abuse of software vulnerabilities – be it in mobile devices, web-based desktop applications or services that run in an enterprise, in the cloud or in a hybrid deployment environment. Strategic basic research must therefore invest in capabilities that help improve the security posture of software. In this context, the quality of all software modules that cover the application logic is important – not just the software that is of direct use in a security solution.

The research agenda sketched below contributes to three major challenges: (1) Supporting a *comprehensive set of security methods* to be utilized throughout the life cycle of software and applications. (2) Delivering verification and security testing technologies that can *enhance guarantees* for essential security-critical properties of software. (3) Enhancing core technology at the level of programming language and compilation technology, thus aiming towards a long term goal of *inherently improving* software security.

These three research themes follow in further detail. In the software development life cycle, the programme aims for early stages of the engineering process ("by-Design"). In the context of security testing and verification, the programme aims for fundamental improvements that can in the long run contribute to high-quality certification. In the context of programming language research, the aim is to radically improve the robustness of future code.

2.1. Secure SDLC – Secure Software Development Life Cycle

Creating and deploying (cyber)secure software systems is an inherently hard problem. Vulnerability lists, such as MITRE's CVE list, clearly illustrate the difficulty of such an endeavor. Fortunately, the increased attention to security has naturally led to efforts in improving software security. Multiple improvements are targeted at addressing different aspects of secure software engineering, though the main focus is on the early stages of the life cycle. The techniques developed in this context aim to resolve the inherent tension between practicality (thus relative low cost) on the one hand, and thoroughness on the other hand. This comprises:

1. Improvements in modeling vulnerabilities, threats, attack scenarios and abuses;
2. Improvements in modeling secure solutions (e.g., in the context of Security-by-Design and Privacy-by-Design, for example based on security and privacy patterns);
3. Improvements in techniques that start from existing software artifacts (such as source code) to enable systematic security analysis at the design level.

The application of the corresponding techniques is in itself often part of other efforts focused on the development lifecycle (e.g., Secure Software Development Life-Cycle (Secure SDLC)), or maturity models. A lot of security knowledge has been documented in principles and bodies of knowledge for reuse and reference.

Many enhancements are essential, in combination with a well-orchestrated application of these new techniques. The full lifecycle of software and services demands for many activities, some of these (but obviously not all) focusing on cybersecurity. Many of such best of breed techniques will be combined in practice. Depending on the maturity of the software development organization, and on the perceived or estimated business risk, this will lead to agile security or to intensive processes that include full-fledged security testing and certification.

A) Industry Needs – Use Cases and Technology Outlook

Ultimately, security and privacy must be built into every software system from the start (Security-by-Design and Privacy-by-Design), for example for GDPR compliance. Currently, however, security and privacy are often dealt with in a primarily reactive way (e.g., through patching discovered/reported vulnerabilities, penetration testing, or code inspection tools), or by addressing challenges that are relevant for the application's

infrastructure (e.g., by applying application firewalls, monitoring and intrusion detection, etc.), yet ignoring the application software itself. Gartner's Application Security Hype Cycle (2018) corroborates that a majority of the included technologies follow this philosophy. There are a few noteworthy exceptions, though, which are at the center of the research theme of secure software and applications. Yet many stakeholders in the ICT industry strongly demand for methods and techniques to develop secure software in a systematic way, backed by best of breed techniques. This demand has been shared by many actors, including financial industries developing mobile applications, providers of critical services, on-line businesses such as ticket sales, auctions and e-commerce, to name but a few.

Application security requirements and threat modeling is currently taking up a pivotal role in the Secure SDLC, but exhibits a need for more efficient methods and procedures that demand for *less skills* and *expertise* and thus yield a better cost/benefit ratio. The *privacy by design* approach (which, by necessity, also includes security by design) is emerging, yet remains relatively vague and is not yet backed by clearly articulated, practical techniques – even now, after the EU legislations of GDPR. Security and Privacy by Design are therefore in need of *practical recommendations and guidelines*. Code-based techniques such as (interactive) application security testing are maturing, yet lack support for *an overall and holistic view of the application*.

By fully embracing agile development and the DevOps movement, the software engineering discipline is in need of efficient, automated, repeatable, and integrated solutions for achieving a satisfactory level of security and privacy, connected to the source code of the application and offering end-to-end traceability. This will in the near future be essential for compliance and assessments that yield some form of certification. This leads to a tension between faster development cycles, thereby reducing the time-to-market, and performing an early-stage and holistic security analysis, which requires security expertise and sufficient resources for producing a high-level design of the system.

The techniques developed for the Secure SDLC research theme aim at resolving this tension, by finding a balance between pragmatism and practicality on one hand, and completeness and thoroughness on the other hand. This potential can be achieved, yet only by maximally leveraging upon reusable knowledge and tool support.

B) State-of-the-art: Highlights

Threat modeling [SHOSTACK14] is a fundamental activity to determine the security requirements of a system, to assess the risks to which the system is exposed, and to determine suitable countermeasures [TURPE17]. It is one of the cornerstone practitioner-oriented activities for Security by Design (STRIDE [SHOSTACK14]) and Privacy by Design (LINDDUN [DENG11]). Several reports on empirical evaluations and experiences (e.g., [SCANDARIATO15, DHILLON11]) have concluded that balancing completeness, relevance, and effort spent to obtain an effective threat modeling approach is still challenging but absolutely necessary.

To mitigate the discovered threats, software developers rely on bodies of reusable knowledge and security expertise. In practice, this knowledge is available in practitioner-oriented resources such as OWASP's cheat sheets and Application Security Verification Standard (ASVS) project, or collections of secure design principles [IEEE14], for example. Security and privacy patterns (e.g., [FERNANDEZ13]) are a prime candidate for collecting such design knowledge in a concrete, high-level, reusable manner. Despite their long and active research history, systematic evaluations such as [YSKOUT15] reveal limitations that must be addressed in order to reach the full potential of reusable security design.

Ultimately, the security of software is strongly dependent on *source code*. For many existing code bases, high-level design documentation is outdated or non-existent. This makes analyzing the security of the software much more complex and resource-intensive. Automated reconstruction of the architectural design from code has a long history [DUCASSE09]. For security in specific, research has delivered work towards automatically identifying architectural security tactics from code (e.g., [MIRAKHORLI16]) or towards predicting the occurrence of vulnerabilities based on source code (e.g., [SCANDARIATO14]). Other approaches have been developed to extract the security architecture from applications (e.g., [BERGER13]). Yet these are mainly designed for monolithic applications on a specific platform, and cannot merge information that is spread over different, distributed modules.

C) Main Areas of Work

The execution of this research theme comprises developing both methods and bodies of knowledge for Security-by-Design, focusing on security requirements (RA 1.1.1), secure design solutions (RA 1.1.2), and high-level security analysis starting from existing source code (RA 1.1.3). Overall, this research theme aims for improvements that can be gradually embraced by practitioners, rather than depending on radical change.

(RA 1.1.1) Cybersecurity requirements

The role of security requirements is to capture precisely the security goals of the system-to-build. Unfortunately, devising a sufficient yet realistic set of security requirements proves to be difficult in practice. Pragmatic threat modeling methodologies such as STRIDE already exist and are being used successfully for this purpose. Nevertheless, successfully carrying them out is labor-intensive and requires significant security expertise. Automated techniques, on the other hand, often results in a large set of threats, many of which are irrelevant.

The goal of this activity is to identify the sweet spot for security requirements and threat modeling by making an affordable trade-off between rigorous and systematic approaches on the one hand (to improve accuracy and completeness), and pragmatism on the other hand (to reduce cost and effort). This opens the door for the development of automated, practical, intelligent methods that can be applied with limited effort, while maximizing the relevance of the discovered requirements.

(RA 1.1.2) Cybersecurity-by-Design solutions

Security and privacy must be built in every software system from the start (Security-by-Design and Privacy-by-Design), for example (but not only) for GDPR compliance. Currently, security and privacy are often dealt mainly in a reactive way (e.g., through discovered/reported vulnerabilities, penetration testing, or code inspection tools). Moreover, security expertise is rare among many/most software engineers.

Transitioning towards a more proactive treatment requires the development of a structured, easy-to-access library of practical knowledge (including processes, methods, and solution patterns), and tools that maximally support the exploitation of this knowledge. The goal of this research activity is to gather, validate, and assess relevant knowledge for Security-by-Design and Privacy-by-Design, and making it available in a uniform and automated matter.

(RA 1.1.3) Security analysis for existing applications

Green-field development is rare; often, the code for which a security analysis must be done already exists. Also, many security (improvement) projects aim for enhancing an existing application or system: the starting point then is an existing code base. Reconstructing the high-level design of the code (e.g., on a whiteboard) is a challenging activity, further complicated by the conflicting understanding and assumptions that exist among stakeholders. High-level views are inevitable, though, to efficiently and thoroughly analyze the design-level security of an application, and it is generally understood that design-level security flaws are the costliest to manage and resolve.

The goal of this research activity is to further develop techniques that can extract and link secure design information from existing code, such that the security of the design can be analyzed. Such techniques should take into account the reality that a modern application consist of different parts, for example written in different languages or running on multiple platforms. Where full automation is impossible, developers should be able to provide assistance, for instance through code annotations.

D) Expected Outcomes and Road Map

For (RA 1.1.1), the first expected outcome (Y1) is a catalog of relevant security & privacy threats and design flaws for specific application types (e.g., web, Cloud, IoT or Mobile). The second expected outcome (Y2) incorporates these threats in a tool-supported method in order to identify and prioritize them based on a high-level system description.

For (RA 1.1.2), the first expected outcome (Y1) is a catalog of privacy design patterns, which incorporate strategies to achieve GDPR compliance. The second expected outcome (Y2) is a catalog of concrete and uniform secure design knowledge, focused on resolving the flaws that have been identified in RA 1.1.1.

Finally, for (RA 1.1.3), the first expected outcome (Y1) is advice and selection guidance on existing tools and approaches for code-based security analysis, based upon their strengths and limitations. The second expected outcome (Y2) is a prototype tool that is able to extract secure design information from source code for a specific class of applications.

In the midterm, further extensions of these outcomes (e.g., towards different application types) will lead to a general, consolidated, and widely applicable body of knowledge and to a practical toolset, backed by solid technical research yet aimed directly at application by practitioners. In the long run, the Consortium wants to demonstrate that these tools and prototypes can be combined with best of breed techniques and tools that are delivered by the industry community.

2.2. Program Verification and Security Testing

For most of the security-sensitive components in software, and for critical software in general, scalable program verification is needed. Such verification enables software developers to prove the absence of implementation faults, vulnerabilities and security problems at a reasonable cost – in terms of manpower.

Static program analysis automatically infers information about the behavior of application software by operating on the source code in order to identify weaknesses and undesired behavior. A number of key strategic challenges remain to be addressed: (1) achieving precision and scalability in the context of non-deterministic programs; (2) delivering guarantees in the context of mandatory quality controls (reviews); (3) achieving scalability in a context of changing source code (updates) in order for cost to remain proportional to the change – and not to the size of the overall system.

Dynamic program analysis complements static analysis. Programs are executed to observe behavior, to control behavior and to harness the program in light of threats and weaknesses. The most important challenges in this respect include the combination of transparency (not influencing behavior), completeness and support of dynamic, composite applications that include third party components that may have been integrated at run time (such as for example in the client side of web applications).

A) Industry Needs – Use Cases and Technology Outlook

To this day, critical components of most computer systems are written in the programming languages C and C++, even though it is well known that writing secure programs in these languages is extremely difficult. The reason is that more secure languages generally impose run-time overhead in terms of performance and memory usage, and offer less fine-grained control over and access to low-level aspects of the execution environment. As a result, infrastructural components such as operating systems and device drivers, basic services such as database management systems and web service frameworks, programming platforms, libraries, and runtimes, and even many applications are still written in these languages. Consequently, security vulnerabilities such as buffer overflows continue to be discovered at an alarming rate, enabling adversaries to exfiltrate private data from, corrupt, or even gain full control of systems ranging from cloud infrastructure and applications to mobile and embedded devices. While tools and techniques exist that help C and C++ software developers find bugs, clearly these are insufficient. Much more needs to be done to enable developers to produce software that is secure and delivers optimal performance and low-level control.

At the same time, teams developing software in more secure languages such as Java, too, face formidable correctness challenges. For example, programs running on a smart card such as a debit card or an identity card are often written in Java and as a result are not susceptible to many typical attacks such as buffer overflows; nonetheless, programming errors that cause the programs to crash and the cards to stop functioning could still lead to very costly incidents. Similarly, programming errors in a company's smartphone app, written in Java, could still cause users to be unable to access the company's services, leading to lost revenue and reputation.

It goes without saying that the abovementioned challenges create headaches for software developers in many businesses and vertical segments, including financial services and healthcare applications, but also and probably most prominently in embedded software-oriented domains, such as automotive, avionics and ICS (industrial control systems), for example.

B) State-of-the-art: Highlights

For securing the most safety-critical applications, machine-checked manual proofs provide the highest level of assurance. However, this assurance comes at an enormous cost in terms of highly skilled labor. The L4.verified microkernel, for instance, was proven correct [KLEIN14] at the cost of 20 PhD-years. The CompCert effort (a verified C compiler) [LEROY09] was of a similar scale.

Research in semi-automated formal verification (FV) therefore seeks to provide the same conclusive results at a reduced cost. Manual intervention only happens in the form of source code annotations, rather than by having to interact with a tool for constructing mathematical proofs. Approaches in the semi-automated formal verification category include VCC [COHEN09], Frama-C [CUOQ12], Dafny [LEINO17], and VeriFast [JACOBS11, PHILIPPAERTS14, PENNINCKX15, HAMIN18, JACOBS18], which the research activities on formal verification will build upon.

When the cost of manual interventions is deemed too high, the most advanced industry players deploy fully-automated Static Application Security Testing (SAST) tools such as Coverity, AppScan, Semmle, or the Clang Static Analyzer instead. These can be effective at uncovering common security vulnerabilities. Facebook, for instance, uses its own Infer [CALCAGNO15] tool to check for bugs in its mobile apps. In the interest of performance, however, many industrial SAST tools perform but a superficial scan of the source code and may therefore leave vulnerabilities lingering (i.e. false negatives). Semantics-based SAST tools (e.g., [GUARNIERI09, ARTZ14, JOHNSON15]), and in particular those constructed according to the theory of abstract interpretation [COUSOT77], feature the highest recall of vulnerabilities but often at the cost of precision loss (i.e. warnings need to be inspected manually for false positives). Research in semantics-based SAST tools therefore generally strives to improve performance and precision, in particular for complex application domains such as mobile, web and distributed applications. The research activity on SAST described below will build on process-modular abstract interpretation [STIÉVENART19] for the incremental checking of distributed applications.

In order to secure the whole software development life cycle, it is important to complement both semi-automatic formal verification (FV) and Static Application Security Testing (SAST) with Runtime Application Security Protection (RASP). The latter monitors deployed applications for policy violations at run time. Policy enforcement can be plugged in as part of the target runtime by relying on virtual machine modifications, e.g., ConScript [MEYEROVICH10], JSFlow [HEDIN14]. However, the adoption of such approaches in the web context is very limited due to the many browser and backend implementations for JavaScript. Alternatively, enforcement can be achieved by code instrumentation and rewriting, e.g., CoreScript [KIKUCHI08], If-transpiler [SAYED18], GIFC [SCULLPUPO18B], but this has a negative impact on (real-time) performance. Real-time processing of large-scale event streams is the subject of lots of research [ZAHARIA10, TOSHNIWAL14, NOGHABI17] which is still unable to detect the complex event patterns [HINZE15] needed in incident monitoring. Systems which do offer expressive pattern languages in which infringements can be specified [APACHE18] do not offer strong performance guarantees.

C) Main Areas of Work

The execution of this research line includes activities on semi-automated formal verification (RA 1.2.1), static application security testing (RA 1.2.2), and run-time security protection (RA 1.2.3).

(RA 1.2.1) Formal Program Verification

Realizing the promise of program verification requires advances both at the level of the underlying program logics that enable concluding properties of a program from properties of its components, and at the level of the algorithms, notations, and tooling that support the application of these logics to programs at an industrial scale.

While great progress has been made in recent years in the power of program logics for proving safety properties of programs, even when they involve fine-grained concurrency, higher-order programming, and behavioral (I/O) properties, the equally important area of liveness properties has so far received much less attention. One goal of this activity is to extend the power of state-of-the-art program logics, so that they can deal with all important security-relevant program patterns and properties, including liveness properties such as fairness and the absence of certain kinds of Denial of Service vulnerabilities, and real-time properties such as the schedulability of a set of real-time tasks.

Applying a program logic to an industrial-scale program requires algorithms, notations, and tooling that support the programming languages, platforms, and frameworks and development processes used, and that minimize the cost of the manual intervention needed to complete the verification, both in terms of the number of person-hours needed, and in terms of the required level of training. The second goal of this activity is therefore to make program verification technology applicable more widely and more cost-effectively; challenges include supporting some of the complex programming language constructs that are not yet supported well by the state of the art, such as C++ templates, and reducing the quantity and the complexity of the source code annotations that developers need to insert into their programs.

(RA 1.2.2) Incremental Static Application Security Testing (SAST) for Distributed Applications

When semi-automated verification is infeasible because of an application's complexity, or because of the frequency at which it is changed and hence needs to be re-verified, fully-automated Static Application Security Testing (SAST) can be effective at identifying common security weaknesses and policy violations before the application is deployed. Semantics-based approaches to SAST (i.e., those performing abstract interpretation or symbolic execution) generally achieve the highest recall but come at a computational cost that precludes their use within contemporary SDLC pipelines —in particular for distributed applications where the possibility of alternative communication orders as well as failures needs to be accounted for.

The goals of this activity are two-fold. First, to move semantics-based SAST approaches within the scope of SDLC pipelines by rendering their analysis engines incremental, and therefore scalable in the size of a committed change. Second, to tailor these analysis engines to the domain of distributed applications by steering their search towards security issues that linger along cross-process execution paths. Key challenges include determining the impact of code changes on previous analysis results, and deriving ordering and multiplicity information about distributed communication as required for the detection of security issues in a sufficiently precise yet scalable manner.

(RA 1.2.3) Efficient Runtime Application Security Protection (RASP) for Distributed Applications

Once deployed, distributed applications still need to be monitored for policy violations, e.g., to detect unauthorized flows of sensitive data, or to detect malicious user behavior in terms of sequences of application-level events, RASP solutions typically monitor the behavior of an application transparently to intercept and report any incident or uncovered policy violations. The goals of this research activity are three-fold: 1) design efficient programming techniques to support the run-time verification of advanced security properties in a transparent manner; 2) improve the capabilities of RASP techniques to the point where security experts can specify fine-grained security policies (e.g., access and information flow control) in one expressive yet familiar specification language so that policies can then be checked continuously upon every commit (and even verified at runtime); and (3) develop an expressive security monitoring language that combines ideas from Complex Event Processing (CEP) and reactive programming to enable expressing behavioral patterns (e.g., complex event sequences) and connecting those to reaction logic.

Runtime monitoring and verification of security policies is challenging for distributed applications in general (cf. above), and full-stack web applications in particular. First, it is important the system itself cannot be successfully attacked in production, which is challenging in the context of a modern language with advanced features like reflection. Moreover, the monitoring infrastructure should have very low overhead such that information gathering for analysis can be always enabled without any noticeable delays in production. Finally, an incident monitoring language should be able to express the envisioned infringement patterns, yet be sufficiently restricted (in a way that is statically known) in order to be guaranteed to run perpetually.

D) Expected Outcomes and Road Map

For (RA 1.2.1), the first expected outcome (Y1) is an approach for adding support for C++ in general, and for C++ templates in particular, to a tool for modular formal verification of C programs such as VeriFast, to support typical use cases such as the C++ Standard Template Library. The second expected outcome (Y2) is a program logic for modular verification of end-to-end fairness properties of applications, such as the property that a server is responsive to each of its clients.

For (RA 1.2.2), the first expected outcome (Y1) is a systematic method for rendering modular static analyzes incremental. We start from modular abstract interpreters for concurrent [STIÉVENART19] and higher-order [NICOLAY19] languages, and investigate how and to what extent they enable scoping the impact of changes to

lexical modules on the analysis results. The second expected outcome (Y2) extends and applies the resulting prototype into an incremental SAST tool capable of detecting at least one representative category of common intra-process security vulnerabilities in a continuous manner. The third expected outcome (Y3) revises the modular and by now incremental abstract interpretation engine so that it computes the inter-process communication information required for detecting inter-process vulnerabilities in a precise yet scalable manner. The remaining outcomes will generalize the incremental SAST tool towards detecting vulnerabilities and policy violations specified by developers, such as in a subset of the language developed in RA 1.2.3.

For (RA 1.2.3), the first expected outcome (Y1) is a first design and prototype implementation of a low-overhead instrumentation platform for monitoring applications in production, and the design of a static analysis method for inferring upper bounds on the resource requirements of a security incident monitoring system. We start from work on client-side run-time verification [ScullPupo18a, ScullPupo18b] of web applications to support the aforementioned developer-specified application-level security policies, and from strongly reactive rule engines to implement the aforementioned security incident monitoring [Kambona18]. The second expected outcome (Y2) extends the resulting monitoring platform with support for distributed interactions, and improves the precision of the analysis when the monitoring system under analysis is elastic in the number of cloud resources it uses.

Throughout the work on this research theme, there will be a constant comparison with related work in order to maintain insight in the best of breed of verification technology, and in the trade-offs between pragmatic approaches to static application testing and full-fledged verification, and between static application security testing and runtime application security protection.

2.3. Secure Programming Languages and Secure Compilation

Many attacks against software happen as a consequence of exploiting low level implementation aspects and/or the infrastructure on top of which the software is running. Well-known classical examples include vulnerabilities resulting from memory management, buffer overflows, code injections (e.g., SQL or scripts), etc. More recent examples include micro-architectural side-channel attacks such as Spectre, Meltdown and Foreshadow. Research is needed to build protection against such attacks into frameworks, compilers, VMs and system software.

Another important challenge corresponds to avoiding logical vulnerabilities; i.e. security-related bugs in the logic of the software which are situated at the application level rather than at the implementation level. Avoiding these is facilitated by programming language support for expressing (and enforcing) application-specific security guarantees. Examples are dedicated language features (e.g., built-in object capabilities to 'isolate' less secure code by design), static annotation systems (e.g., type systems where 'confidentiality' is part of the types), etc.

Application-level vulnerabilities are especially prominent in software that runs on multiple (distributed) computers. Programming language support for security for such systems takes the form of abstractions for deployment, coordination and communication. Moreover, by parameterizing these abstractions further, distributed security aspects can be expressed in a composable way. This is even more relevant now that distributed systems are written by different teams that often use multiple programming paradigms.

A) Industry Needs – Use Cases and Technology Outlook

To significantly reduce the inherent weaknesses/vulnerabilities in newly developed software, it remains a strategic and important challenge to develop new programming languages and frameworks that offer stronger abstractions and guarantees, not just for managing the growing complexity of software, but also to support higher-level reasoning about security.

Many contemporary software applications run on platforms that are connected, open and virtualized. Connectedness results from the fact that the applications are distributed across multiple devices and infrastructures. Openness stems from the fact that the number of and the exact identities of those devices is not known upfront. Virtualized means that the software runs on 'machines' that are software artifacts themselves; usually after being deployed dynamically as a result of mobile code. The Java ecosystem was the first mainstream example of this trend. More recently, the idea has been taken to the next level in platforms

such as Erlang/BEAM (e.g., the backend of WhatsApp) and JavaScript-based browser technology. Applications running on such platforms are prone to a number of vulnerabilities that cannot be relegated to 'the security module'. Instead, they are systemically present in the very computational fabric of the platforms (i.e. in the virtual machines, the compilers and the programming languages). The most extreme manifestation of the problem is typically a consequence of mobile code that is dynamically deployed.

For example, developers of a web application that needs to interact with untrusted native components should be able to rely on convenient infrastructure for invoking these components without worrying about assembly-level security. Those low-level details should be dealt with in a standard way by a security-aware compiler or framework and programmers should be able to reason in terms of the source language. Similar concerns arise as a result of the ongoing virtualization of IoT applications and cyber-physical systems: sensors and actuators need to be protected against unforeseen behavior of newly discovered communication partners and against unforeseen deployments of potentially hostile code.

B) State-of-the-art: Highlights

An important goal of current research is to allow programmers to develop programs and reason about their security in terms of abstractions as offered by high-level source languages. This requires the development of secure compilers [ABADI99, PATRIGNANI19], an active field of research. For many useful source abstractions, the security primitives offered by commodity processors are not well-suited for building a secure compiler enforcing them. A promising alternative are capability machines, which have been recently used as the target of a formally secure compilation phase [SKORSTENGAARD18]. However, some abstractions remain hard to enforce efficiently, for example because their enforcement relies on temporary authority. This includes important properties like the correctness of control flow or ownership and thread-safety in linearly-typed languages like Rust. Recent research has investigated new primitives for enforcing such properties efficiently: local capabilities [SKORSTENGAARD18] and linear capabilities [SKORSTENGAARD19].

To build a secure system, it does not suffice to have a secure compiler which preserves the security properties of the source language. Additionally, we need a source language which enables the programmer to efficiently guarantee relevant security properties in the first place. An important class of properties which is currently hard to guarantee in state-of-the-art programming languages are related to the side effects produced by programs, especially when many different components interact over complex interfaces. For example, it may be important that a component will never access the network, or only certain servers on the network according to a certain protocol. For guaranteeing such properties, recent research has started to explore the use of algebraic effect handlers, to specify and implement the effect interfaces that components have access to [PLOTKIN09]. However, proving the properties that follow from such an approach remains underexplored. A promising approach to do better is the use of a technique called effect polymorphism. This offers a general way to encode and prove properties about programs with side-effects, but has so far only been used for proving properties of programming languages, not individual programs [DEVRIESE16].

Modern distributed applications are conceived as complex interplays of components, some of which sit on clients and others on servers. The components are reusable units of behavior that can be reused, programmed and debugged separately. Recently, the notion of "micro-services" has been proposed as a unifying concept for such a componentization. Micro-services can be distributed between servers and clients and/or on different nodes of the same server. The fact that micro-services are autonomous loosely coupled computational units that communicate via message passing opens up a particular set of security problems. Initially, micro-services were technically realized by means of programming patterns in some mainstream language which makes it very difficult to reason about their security aspects. However, more recently, micro-services are aligned with reactive "actors" (Akka.io, Microsoft Orleans, [SALVANESCHI16], [VANDERVONDER17]) in actor-based languages of which the semantic properties are well-understood by the programming language community. This opens the road to studying vulnerabilities in micro-service architectures in terms of actor semantics and allows to enhance said actor languages with new language features that allow programmers of micro-service based applications to write secure distributed code.

C) Main Areas of Work

The execution of this research line includes 3 research activities:

(RA 1.3.1) Mechanically-verified Security Proofs for Capability Machine Programs

This research activity works towards the development of secure compilers which enforce widespread abstractions. This challenge will be approached by relying on capability machines and security primitives like local or linear capabilities. To realize this potential, it is important to develop a reusable methodology and framework for proving the properties of programs that use these security primitives.

The goal of this first research activity is to develop such a methodology, building on existing general machine-verified reasoning frameworks, as well as existing non-mechanically verified techniques for reasoning about capabilities. The methodology will be designed to be reusable for different applications of capability machine security primitives and will be applied to at least one application.

This application will take the form of a secure compilation technique, i.e. the design of a compiler which enforces a security-relevant source-language property in the presence of target-language interaction with untrusted adversarial code. One promising target is to develop a new technique for enforcing ownership and thread-safety properties in linearly-typed languages like Rust.

(RA 1.3.2) Specifying and Proving Security Properties of Side-Effecting Programs

In this second research activity, we will develop a new way to formulate and prove security properties of programs with side-effects. As a first step, we will do this in languages with support for writing mechanically verified proofs about programs and algorithms. Such languages are known as dependently-typed and have been used for writing the first realistic, provably bug-free C compiler CompCert [LEROY09]. In a second step, we intend to transport these results to the imperative, object-oriented programming languages that are widely used in industrial applications.

Concretely, we will develop a library for writing side-effecting programs and proving their properties, focusing on those that are security-relevant. Initially, we will look at unary properties: for example, that program X will only use the network to connect to servers Y and Z. Subsequently, we intend to extend our approach to relational properties (for example, the data sent by program X to server Y do not leak information about the data received from channel Z). Finally, we intend to construct an imperative programming language that compiles to these new primitives and inherits features for specifying and proving security properties.

(RA 1.3.3) Language-embedded Security Policies for Distributed Micro-Services.

A modern distributed system (e.g., an IoT or Web application) is composed of software components that live on various devices (computers, servers, sensors, actuators, …). Recently, these software components are termed micro-services. Micro-services collaborate by sending messages or by propagating data values to one another. Academically speaking, such micro-services correspond to the actor-model enriched with streaming semantics.

Even if every single micro-service is developed in a secure (sequential) programming language, this does not guarantee that the behavior emerging from a cluster of distributed micro-services is secure. Because of the complex interaction patterns between the loosely-coupled distributed micro-services, excluding all possible insecure executions by means of static techniques becomes impossible. This is aggravated by the fact that 'partial failure' becomes part of the computational fabric that makes up distributed systems: individual micro-services can fail and simply be restarted without affecting the overall semantics of the application.

In this research activity, we envision the design of a language-embedded framework for composable 'first-class' security policies. These policies describe desired and prohibited data-flow and control-flow aspects of a micro-service. Making them first-class allows us to reason about them independently from the functionality of a micro-service, allows us to compose them, and allows us to deploy stable security policies on several micro-services in various applications. Rather than designing a language from scratch (requiring parsing, compiling, etc.), our research builds atop a modern extensible language such as Elixir, Scala or Clojure.

D) Expected Outcomes and Road Map

The research themes and activities within this Research Track encompass the different phases of the software development lifecycle, as shown in the figure 2-1. Furthermore, the activities target different horizons, by ranging from the extension and consolidation of the know-how about current technologies used in practice, over the evolution of these technologies to further improve them, towards the development of new and innovative technologies.

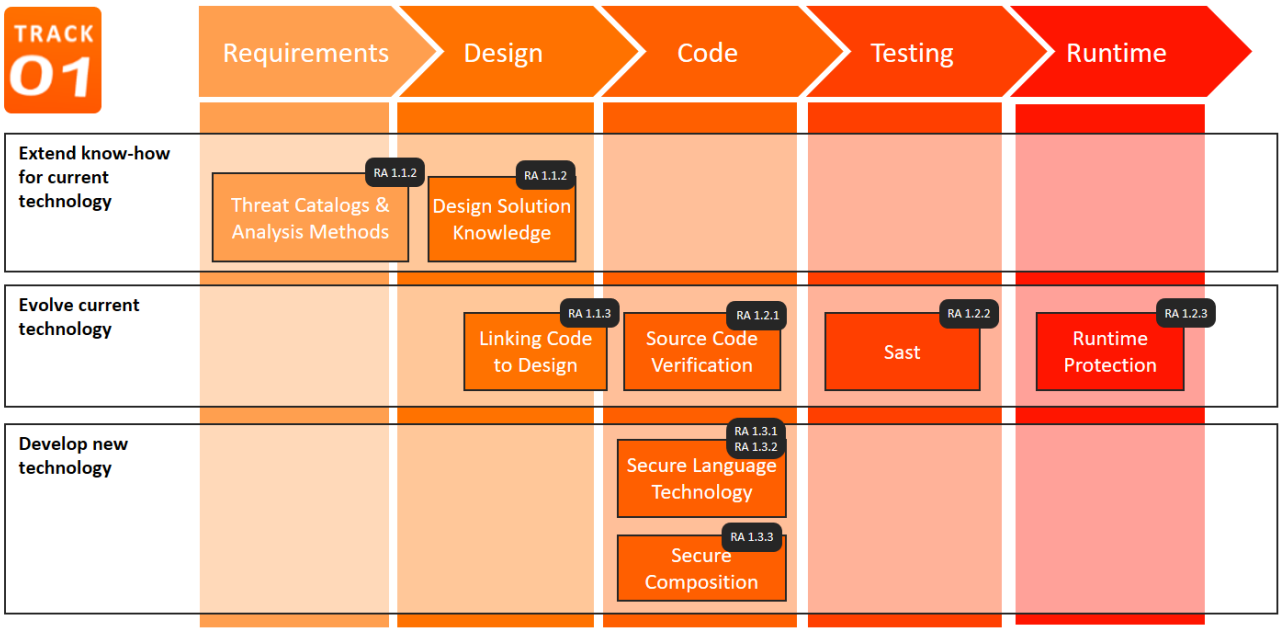

Figure 2-1: Situation of the Research Activities of Track 1 in the Software Development Lifecycle

For (RA 1.3.1), the first expected outcome (Y1) is a design and initial prototype of a novel methodology for proving security properties of programs using security primitives on capability machines. The second expected outcome (Y2) builds on this initial prototype and applies it to at least one challenging application of the primitive.

For (RA 1.3.2), the first expected outcome (Y1) is a library for implementing programs with side-effects in a dependently-typed language using effect polymorphism, with support for at least unary security-relevant properties and to apply it to simple proof-of-concept examples. The second expected outcome (Y2) extends this prototype to support relational properties and uses it to implement a larger application and prove its security properties.

For (RA 1.3.3), the first expected outcome (Y1) is a prototype language that allows us to monitor and guarantee that micro-service compositions do not suffer from unwanted sharing/leakage of data (private to individual micro-services) as a consequence of unanticipated message exchanges. In the second expected outcome (Y2), we will study how vulnerabilities that are the result of partial failures in micro-service compositions (in the 'surviving' part of the application) can be prevented and/or dealt with at the language level.

2.4. Connections with Other Research Tracks

In relation to the other Research Tracks, there is a strong connection between (RA 1.3.1) in this Research Track and (RA 3.1.2) of Research Track 3. The latter research activity will study and evaluate the implementation of capability-based security primitives in processors, while the former will study how we can prove security properties of programs that rely on those primitives, and extensions of them. It is expected that both activities will yield mutually beneficial impact, and strengthen each of the results. This will definitely create additional opportunities for collaboration.

There is also a connection between (RA 1.2.3) of this Research Track, and (RA 2.2.1) in Research Track 2. The latter will study security services offering enforcement of access control policies in an externalized and modular way. The former focuses on runtime verification of application-level security policies in general, of which access and information flow control policies are the two best-known types. While the work in Research Track 2 aims to analyze network packages produced by distributed applications, the work in this Research Track aims to analyze operations performed at language runtime/VM level during the execution of a distributed application. Since policy enforcement is crosscutting through the whole software stack, we believe both efforts are necessary and complementary and we foresee synergies and opportunities for collaboration to align language and network-level access control enforcement.

2.5. References

[ABADI99] Abadi, Martin. 1999. Protection in programming-language translations. In Secure Internet programming.

[APACHE18] Apache Software Foundation, 2018. FlinkCEP - Complex event processing for Flink. https://ci.apache.org/projects/flink/flink-docs-stable/dev/libs/cep.html.

[ARTZ14] Steven Arzt, Siegfried Rasthofer, Christian Fritz, Eric Bodden, Alexandre Bartel, Jacques Klein, Yves Le Traon, Damien Octeau, and Patrick D. McDaniel. 2014. Flowdroid: Precise context, flow, field, object-sensitive and lifecycle-aware taint analysis for Android apps. In Proceedings of the 35th ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI14)

[BERGER13] Berger, B.J., Sohr, K. and Koschke, R., 2013. Extracting and analyzing the implemented security architecture of business applications. Proceedings of the European Conference on Software Maintenance and Reengineering, CSMR, pp.285–294.

[CALCAGNO15] Calcagno, C., Distefano, D., Dubreil, J., Gabi, D., Hooimeijer, P., Luca, M., O'Hearn, P., Papakonstantinou, I., Purbrick, J., Churchill, D., 2015. Moving fast with software verification. Proceedings of the NASA Formal Methods Symposium, pp.3--11.

[COHEN09] Cohen, E., Dahlweid, M., Hillebrand, M. A., Leinenbach, D., Moskal, M., 2009. VCC: A practical system for verifying concurrent C. Proceedings of the International Conference on Theorem Proving in Higher Order Logics, pp.23-42.

[COUSOT77] Patrick Cousot and Radia Cousot. 1977. Abstract interpretation: A unified lattice model for static analysis of programs by construction or approximation of fixpoints. In Symposium on Principles of Programming Languages (POPL77).

[CUOQ12] Cuoq, P., Kirchner, F., Kosmatov, N., Prevosto, V., Signoles, J., Yakobowski, B., 2012. Frama-C, a software analysis perspective. Proceedings of the International Conference on Software Engineering and Formal Methods, pp.233-247.

[DENG11] Deng, M., Wuyts, K., Scandariato, R., Preneel, B. and Joosen, W., 2011. A privacy threat analysis framework: supporting the elicitation and fulfillment of privacy requirements. Requirements Engineering, 16(1), pp.3–32.

[DEVRIESE16] Devriese, D., Birkedal, L., and Piessens, F. 2016. Reasoning about Object Capabilities Using Logical Relations and Effect Parametricity. In European Symposium on Security and Privacy.

[DHILLON11] Dhillon, D., 2011. Developer-Driven Threat Modeling - Lessons Learned in the Trenches. IEEE Security & Privacy, 9(4), pp.41–47.

[DUCASSE09] Ducasse, S. and Pollet, D., 2009. Software architecture reconstruction: A process-oriented taxonomy. IEEE Transactions on Software Engineering, 35(4), pp.573–591.

[FERNANDEZ13] Fernandez-Buglioni, E., 2013. Security patterns in practice: designing secure architectures using software patterns. John Wiley & Sons.

[GUARNIERI09] Salvatori Guarnieri and Benjamin Livshits. 2009. Gatekeeper: mostly static enforcement of security and reliability policies for JavaScript code. In Proceedings of the 18th USENIX Security Symposium (SSYM09).

[HAMIN18] Hamin, J., Jacobs, B., 2018. Deadlock-free monitors. Proceedings of the European Symposium on Programming, p.415-441.

[HEDIN14] Daniel Hedin, Arnar Birgisson, Luciano Bello, and Andrei Sabelfeld. 2014. JSFlow: Tracking Information Flow in JavaScript and Its APIs. In Proceedings of the 29th Annual ACM Symposium on Applied Computing. ACM, New York, NY, USA, 1663–1671.

[HINZE15] Hinze, A. and Voisard, A. EVA: an Event Algebra Supporting Complex Event Specification. Information Systems, 48:1–25, 2015.

[IEEE14] Arce, I. et al., 2014. Avoiding the top 10 software security design flaws. IEEE Center for secure design. Available from https://cybersecurity.ieee.org/center-for-secure-design/

[JACOBS11] Jacobs, B., Piessens, F., 2011. Expressive modular fine-grained concurrency specification. Proceedings of the International Symposium on Principles of Programming Languages, pp.271-282.

[JACOBS18] Jacobs, B., Bsonacki, D., Kuiper, R., 2018. Modular termination verification of single-threaded and multithreaded programs. ACM Transactions on Programming Languages and Systems 40(3), pp.12:1-12:59.

[JOHNSON15] Andrew Johnson, Lucas Waye, Scott Moore, and Stephen Chong. 2015. Exploring and enforcing security guarantees via program dependence graphs. In Proceedings of the 36th ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI15).

[KAMBONA18] Kambona K., Renaux T., De Meuter W. (2018) Harnessing Community Knowledge in Heterogeneous Rule Engines. In: Web Information Systems and Technologies. WEBIST 2017. Lecture Notes in Business Information Processing, vol 322. Springer.

[KIKUCHI08] Haruka Kikuchi, Dachuan Yu, Ajay Chander, Hiroshi Inamura, and Igor Serikov. 2008. JavaScript Instrumentation in Practice. In Programming Languages and Systems. Springer Berlin Heidelberg, Berlin, Heidelberg, 326–341.

[KLEIN14] Klein, G., Andronick, J., Elphinstone, K., Murray, T., Sewell, T., Kolanski, R., Heiser, G., 2014. Comprehensive formal verification of an OS microkernel. ACM Transactions on Computer Systems, 32(1), pp.2:1-2:70.

[LEINO17] Leino, K. R. M., 2017. Accessible software verification with Dafny. IEEE Software 34(6), pp.94-97.

[LEROY09] Leroy, X., 2009. Formal verification of a realistic compiler. Communications of the ACM, 52(7), pp. 107—115.

[MEYEROVICH10] Leo A Meyerovich and Benjamin Livshits. 2010. ConScript: Specifying and Enforcing Fine-Grained Security Policies for JavaScript in the Browser. IEEE Symposium on Security and Privacy. IEEE, 481–496.

[MIRAKHORLI16] Mirakhorli, M. and Cleland-Huang, J., 2016. Detecting, Tracing, and Monitoring Architectural Tactics in Code. IEEE Transactions on Software Engineering, 42(3), pp.206–221.

[NICOLAY19] Nicolay J., Stiévenart Q., De Meuter W., De Roover C., 2019. Effect-Driven Flow Analysis. Verification Model Checking Abstract Interpretation (VMCAI): 247-274.

[NOGHABI17] Noghabi, S. A., Paramasivam, K., Pan, Y., Ramesh, N., Bringhurst, J., Gupta, I., and Campbell, R. H., 2017. Samza: Stateful Scalable Stream Processing at LinkedIn. Proceedings of the International Conference on Very Large Data Bases Endowment.

[PATRIGNANI19] Patrignani, M., Ahmed, A., and Clarke, D. 2019. Formal Approaches to Secure Compilation: A Survey of Fully Abstract Compilation and Related Work. ACM Comput. Surv. 51, 6 (February 2019), 125:1–125:36.

[PENNINCKX15] Penninckx, W., Jacobs, B., Piessens, F., 2015. Sound, modular and compositional verification of the input/output behavior of programs. Proceedings of the European Symposium on Programming, pp.158-182.

[PHILIPPAERTS14] Philippaerts, P., Mühlberg, J. T., Penninckx, W., Smans, J., Jacobs, B., Piessens, F., 2014. Software verification with VeriFast: industrial case studies. Science of Computer Programming 82, pp.77-97.

[PLOTKIN09] Plotkin, G. and Pretnar, M., 2009. Handlers of Algebraic Effects. In Programming Languages and Systems.

[SALVANESCHI16] Salvaneschi, G. and Mezini, M., 2016. Debugging reactive programming with reactive inspector. In Proceedings of the International Conference on Software Engineering (ICSE '16).

[SAYED18] Bassam Sayed, Issa Traoré, and Amany Abdelhalim. 2018. If-transpiler: Inlining of hybrid flow-sensitive security monitor for JavaScript. Computers & Security 75 (June 2018), 92–117.

[SCANDARIATO14] Scandariato, R., Walden, J., Hovsepyan, A. and Joosen, W., 2014. Predicting vulnerable software components via text mining. IEEE Transactions on Software Engineering, 40(10), pp.993–1006.

[SCANDARIATO15] Scandariato, R., Wuyts, K. and Joosen, W., 2015. A descriptive study of Microsoft's threat modeling technique. Requirements Engineering, 20(2), pp.163–180.

[SCULLPUPO18A] Scull Pupo A., Nicolay J., Gonzalez Boix E., 2018. GUARDIA: Specification and Enforcement of JavaScript Security Policies without VM modifications. ManLang 17:1-17:15.

[SCULLPUPO18B] Scull Pupo A., Nicolay J., De Roover C., Gonzalez Boix E., 2018. Practical Information Flow Control for Web Applications. Runtime Verification (RV) 2018: 372-388.

[SHOSTACK14] Shostack, A., 2014. Threat Modeling: Designing for Security, Wiley.

[SKORSTENGAARD18] Skorstengaard, L., Devriese, D., and Birkedal, L.. 2018. Reasoning About a Machine with Local Capabilities. In Programming Languages and Systems, pp. 475–501.

[SKORSTENGAARD19] Skorstengaard, L., Devriese, D., and Birkedal, L.. 2019. StkTokens: Enforcing Well-bracketed Control Flow and Stack Encapsulation Using Linear Capabilities. Proc. ACM Program. Lang. 3, POPL (January 2019), 19:1–19:28.

[STIÉVENART19] Stiévenart Q., Nicolay J., De Meuter W., De Roover C., 2019. A general method for rendering static analyzes for diverse concurrency models modular. Journal of Systems and Software, 147: 17-45.

[TOSHNIWAL14] Toshniwal, A., Taneja, S., Shukla, A., Ramasamy, K., Patel, J.M., Kulkarni, S., Jackson, J., Gade, K., Fu, M., Donham, J., Bhagat, N., Mittal, S., and Ryaboy, D., 2014. Storm@Twitter. In Proceedings of the International Conference on Management of Data (SIGMOD).

[TURPE17] Turpe, S., 2017. The Trouble with Security Requirements. In 2017 IEEE 25th International Requirements Engineering Conference (RE). IEEE, pp. 122–133.

[VANDENVONDER17] Van den Vonder, S., De Koster, J., Myter, F., De Meuter, W., 2017. Tackling the awkward squad for reactive programming: the actor-reactor model. REBLS@SPLASH 2017, pp. 27-33.

[YSKOUT15] Yskout, K., Scandariato, R. and Joosen, W., 2015. Do security patterns really help designers? Proceedings - International Conference on Software Engineering, 1, pp.292–302.

[ZAHARIA10] Zaharia, M., Chowdhury, M., Franklin, M. J., Shenker, S., and Stoica, I., 2010. Spark: Cluster Computing with Working Sets. Proceedings of the USENIX Conference on Hot Topics in Cloud Computing.

3. Research Track 2: Strategic Security Services

Contributing Authors: Bart De Decker, Bart Preneel, Bert Lagaisse, Bruno Crispo, Dave Singelée, Davy Preuveneers, Els Kindt, Enrique Argones Rua, ***Frederik Vercauteren****, Nigel Smart, Peggy Valcke, Wouter Joosen*

Scope

Many security solutions are based on controls that allow or deny users to perform specific actions. Such security solutions are typically based on identity management, on the distribution and validation (and revocation) of credentials and on the management of permissions. Correct evaluation of credentials will determine whether the user (and by extension component, service etc.) is allowed to perform certain functionalities, whether he/she gets or is granted access to specific information, etc. The rules for granting access and auditing access are often driven by domain-specific policies and legal regulations on both security as well as privacy.

The first type of security functionality is typically delivered by a composition of strategic and critical security services such as identity services, authentication services, and authorization services, and the creation, utilization and management of audit trails. In practice, practical solutions need to combine in-band prevention and out-of-band detection to be scalable. In this Research Track, the notion of security service refers to this type of essential security building blocks that are not embedded in system and network layers; and should neither be part of the application layer as one needs reliable and reusable building blocks – security services – to be successful.

The corresponding research domains evolve rapidly and expand drastically as the exploitation of vast data sources demands for data centric controls, for finely grained and attribute based evaluation of security policies and for advanced support to prove compliance in terms of policies and regulations.

Also new and advanced data protection mechanisms become gradually available, enhancing the data protection capabilities through e.g., homomorphic encryption, and enabling controlled data sharing through MPC (Multi-party computation). These techniques deliver encouraging results but demand for further improvement to be scalable and widely applicable. Meanwhile, strategic solutions tend to combine advanced mechanisms with (data) access middleware to deliver solutions in real-world digital platforms.

Based on the strategic needs related to these security services, three technical research themes follow in further detail: (I) Identity Management and Authentication, (II) Authorization and Audit, and (III) Advanced encryption techniques and data access middleware. In addition, this Research Track includes a fourth theme that addresses the important domain of (IV) policy and regulations.

The research on policies and regulations is an important ingredient to ensure that technical results can be applied in a context that respects new legal provisions that directly affect business (such as the well-known example of the GDPR). It definitely will fit very well in the context of topics such as authentication, authorization, audit and data sharing. Yet this work is not unique to Research Track II, it will certainly affect and interact with other Tracks of the cybersecurity research program.

3.1. Identity Management and Authentication

The creation of authentication systems that can work seamlessly with a growing palette of applications and platforms remains a continuous and major challenge. The costs associated with trustworthy management of users, identity and the associated platforms for access control are very high, both for users and for the parties providing online services and offering applications. The required research therefore seeks solutions that, on the one hand, provide security, increase transparency and privacy while at the same time limit the threshold and burden (friction) for the user.

Authentication relies on the use of elements that prove the identity, also known as authentication factors. Three types of authentication factors typically come into play, namely knowledge factors (e.g., passwords), possession factors (e.g., authentication tokens), and inherent factors (e.g., biometrics). From a security point

of view, authentication solutions based on a single factor suffer from different vulnerabilities, which depend on the type of employed authentication factor. For instance, passwords can be forgotten or eavesdropped, authentication tokens can be lost or stolen, and some biometrics can be spoofed. The use of multi-factor authentication, where several types of authentication factors are combined, is a sensible recommendation in security practices[3]. However, the use of multiple factors presents other challenges, since if not carefully designed it can lead to increased user friction, availability issues, etc.

In the case of biometrics, nowadays there are different mature and accurate enough modalities, from both main types (physiological, such as face, fingerprint, vein patterns, text-independent speaker, gait, electrocardiogram, photoplethysmogram, or hand geometry recognition; and behavioral, such as text-dependent speaker, or handwritten signature). However, the biometric information in the biometric templates carries critical information, and its disclosure may lead not only to security, but also to considerable privacy losses. Biometric authentication schemes must therefore rely on mechanisms that provide protection for the biometric information. There are different technologies that can be used for this purpose.such as the use of biometric protection techniques, focused on using biometric templates which are usable for authentication, but do not disclose biometric information. The use of template protection techniques usually reduces to some extent the accuracy of the authentication. There exist alternative approaches, based on MPC or HE, which can replicate the accuracy of the unprotected biometric schemes, with some additional computational burden. These alternatives will be ideal in different use cases, depending on factors such as the available computational power, or bandwidth restrictions.

An important research theme in this context is the use of multiple sensors, which are nowadays prevalent in smart devices and environments. These solutions have the ability to identify users based on behavioral patterns (behaviometrics): the way the users interact with the environment and devices is modelled. This enables more transparent authentication mechanisms. Some examples of behaviometrics are in fact behavioral biometrics, such as keystroke dynamics, mouse dynamics, touchscreen dynamics, or gait authentication. However, this denomination also includes context-based information describing the interactions between the user and its environment or devices, such as location, applications usage, or system logs. As in the case of biometrics, behaviometrics do comprise critical private information, and it therefore also becomes paramount to consider data protection requirements and solutions when integrating these technologies.

A) Industry Needs – Use Cases and Technology Outlook

This area constantly shows a strong interest and demand in industry as virtually any application domain in the digital revolution, and many vertical industry segments, are constantly aiming for improvement in this space. This is illustrated below with some obvious examples.

In the E-commerce space, there is an obvious need for user-friendly authentication, while the state of practice reveals several mishaps due to lack of authentication. Without proper authentication, adversaries would get access to personal data, including the overview of the user's online purchases, and even to financial information such as credit cards. Even worse, an adversary would be able to impersonate the user in online transactions. In addition, essential concerns for data protection arise in many deployment scenarios that utilize cloud-based resources: novel authentication techniques often come with the extra data collection and cloud processing.

More than ever the fin-tech and banking sector needs strong authentication, for example because of risk-adaptive authentication for higher-risk payment transactions, because of the Open Banking/PSD2 standard. Important challenges in this industry include in-app authentication for mobile apps, challenged by the rise of mobile malware, and the shift towards behavioral biometrics[4].

Meanwhile, specialized security providers (so-called ISVs) deliver solutions to be integrated in new applications. From the perspective of a business owner/application architect, there is an inherent need to

[3] Guidelines for SMEs on the security of personal data processing, ENISA, 2016

[4] https://blog.malwarebytes.com/101/2018/04/securing-financial-data-of-the-future-behavioral-biometrics-explained/, https://www.gemalto.com/financial/inspired/behavioral-biometrics

integrate third party providers, be it private identity or national identity providers, enablers for mobile devices and apps, etc. More and more of the available solutions are cloud-based (e.g., IDaaS). The need for robust integration is a challenge in its own right.

While identity management and authentication is a crowded ICT security market with many technical offerings, it also is evolving rapidly and in need of contributions that rely on advanced strategic research. Highlights of the scientific/technical state-of-the-art present some of the directions.

B) State-of-the-art: Highlights

Collaborative authentication solutions reduce the security risk of single-factor authentication schemes by relying on multiple devices (or factors) to carry out an authentication instance. The underlying assumption is that a certain number of the user's devices – larger than a predefined threshold – is not compromised and actively present to perform the authentication process. This concept has been deployed in the Pico solution proposed by Stajano [STAJ11]. The evolution of this (in principle) generic approach is driven by new possibilities in the development of additional factors. We sketch recent and relevant work in the space of biometric authentication factors, and in the space of context-based authentication.

Novel research in multi-factor authentication often builds upon ingredients based on biometry. We illustrate this work by sketching related work on cardiac, gait, and online signature modalities. These are of special interest in wearables and smartphone ecosystems. Biometric authentication algorithms are usually classified as discriminative or generative approaches. Generative approaches learn a joint probability distribution model by estimating the underlying data distribution, whereas discriminative methods learn a conditional probability distribution model directly on the observed data. Among the generative approaches, one of the most accurate and versatile is based on Hidden Markov Models, which are very well suited for signals with different states, such as the ones provided by cardiac [ODIN10][ANDB06], speech [REQD00], gait signals [SEGU17], or online signature [ARAL12]. HMMs have demonstrated high accuracy, obtaining Equal Error Rates (EER) in the order of 1% for all these modalities, which grant a good compromise between usability and security for many applications. In the case of discriminative approaches, Deep Learning techniques are nowadays attracting a lot of attention, due to the ability to provide remarkable performance figures in different domains, including also gait recognition [HOTR17] and ECG-based authentication [ANDB06][PYKM17]. Siamese architectures have been proposed for online signature recognition [TVFO18]. Furthermore, eigen-model representations derived from Universal Background Model (UBM) systems provide fixed length templates that have also been successfully used for authentication [ARGO12]. In the case of the Deep Learning approaches, the highest-level representation layers of Siamese networks also provide fixed length representations that can also be used for authentication outside the network context. These fixed length representations are specially convenient for template protection within cryptosystems, such as fuzzy extractors [VAPJ18] based on a fuzzy commitment [DODI04][JUWA99], as shown in [ARGO12][VAPJ18].

One of the promising research directions in the development of frictionless authentication schemes is context-based authentication. Context is used as an additional authentication factor. For example, the authentication process between two devices may only succeed if they share the same context. Although context can be defined in various ways, many define the context as the location of the device and/or the proximity between the authenticator and the verifier. To verify the proximity between two devices, multiple approaches can be used. Obviously, one could rely on GPS signals to determine the distance between two devices. However, GPS does not work well indoor, and GPS signals can be spoofed. One could rely on auxiliary information to verify that two devices are within proximity. Halevi et al. [HMSX12] use ambient audio to detect the proximity of two devices to thwart relay attacks in NFC payment systems. Truong et al. [TGSS14] propose a framework that detects co-location of two devices comparing features from multiple sensors, including GPS, Bluetooth, WiFi and audio. Schürmann and Sigg [SCSI13] use ambient audio to derive a pairwise cryptographic key between two co-located devices. Karapanos et al. [KMSC15] verify the proximity of two devices by comparing the ambient noise recorded by their microphones. Anand and Saxena [ABSA17] propose to use acoustic noises to mask the sounds of vibration in vibrational authentication and pairing schemes. Also, visual channels can be used within the authentication process. For example, Sturgess and Martinovic [STMA17] presented VisAuth - authentication over a visual Channel Using an embedded image. Mayrhofer and Welch [MAWE07] propose a human-verifiable authentication protocol using visible laser light. Yet another approach is to use distance

bounding protocols. Based on the concept of integrity regions [CCKT10], Singelée and Preneel proposed [SIPR07] a key agreement solution that relies on distance bounding protocols.

The proposed research agenda sketched below starts from the promises and challenges that come with the themes of context-based authentication and collaborative multi-factor authentication as sketched above.

C) Main Areas of Work

The execution of this Research Track includes three work research activities: a first one related to identity management, a second one related to frictionless authentication and a third one addressing privacy preservation on authentication solutions.

(RA 2.1.1) Identity

The first activity deals with the management of identities and cryptographic keys, associated with both users and devices, in decentralized deployment environments. Multiple personal mobile and wearable devices should collaborate to establish a trustworthy identity. The first goal of this activity is to define protocols for distributed identity and key management where the presence of these devices as well as their strength to establish the identity of a user can continuously change.

A second goal is to raise the level of abstraction of the above protocols for risk-adaptive applications by means of trust brokering middleware complemented with scalable deployment and configuration tools. The objective is to address evolving authentication needs (e.g., initiate new combinations of authentication factors for, step-up authentication) while minimizing the effort required to modify the risk-adaptive applications.

(RA 2.1.2) Frictionless Authentication: Collaborative and Continuous

The second activity will investigate frictionless authentication on top of multi-factor and multi-modal knowledge, possession and inherence factors, including revocable biometrics. A first goal is to research the strengths and weaknesses of data versus decision fusion for collaborating authentication factors. This collaborative authentication should be context-adaptive such that the right set of devices is leveraged to offer the best accuracy and user experience in any given situation. Furthermore, the collaborative authentication should be resilient in adversarial environments in which not only each participating authentication factor can be attacked (e.g., spoofing, replay, poisoning, evasion, mimicry, liveness and other attacks), but also the collaboration itself, as well as the context in which they are deployed.

Regular two-factor authentication solutions do not reverify a user's identity after an application session is started, widening the window of opportunity for session hijacking. Another goal of this research activity is to extend the frictionless authentication with support for continuous authentication factors. The objective is to enable continuous verification of a person's identity by means of his or her behavioral traits. By leveraging behavioral biometrics (a.k.a. behaviometrics), continuous authentication aims to fill this security gap in a user-friendly manner. Similar to traditional biometrics, any proposed behavioral biometrics solution must be resilient and robust in adversarial environments. The envisioned outcome is to support continuous authentication on top of state-of-practice and emerging authentication standards (e.g., FIDO2, WebAuthn) to offer a flexible integration with service providers.

(RA 2.1.3) Privacy-preserving Identity and Authentication

Continuous authentication leverages amongst others behavioral biometrics which inherently hold sensitive information. For example, a gait authentication factor leveraging the accelerometer of a smartphone or smartwatch may reveal the location of the individual. This research activity will carry out a systematic privacy threat analysis of state-of-the-art identity and authentication techniques. The focus will be on threats in collaborative authentication schemes and the leakage of sensitive information through side channels or secondary use of information.

Complementing the privacy threat analysis, this research activity will design, implement and evaluate enabling technologies that aim to maintain the confidentiality of irrevocable biometrics and behaviometrics with template protection schemes. This task will also consider the case where the execution of biometric and behaviometric matching algorithms is delegated to resource rich but untrustworthy computational environments (e.g., the cloud). We will investigate the applicability and trade-offs, and further enhance privacy-preserving machine learning techniques, multi-party computation and homomorphic encryption schemes for privacy-preserving authentication.

D) Expected Outcomes and Road Map

In year 1, the above three research activities will yield initial results on collaborative authentication and identity management for mobile and wearable devices, with foundational support for context-dependent multi-modal data and decision fusion for multi-factor authentication. Another expected outcome is the privacy threat analysis of state-of-the-art authentication solutions.

Initial tangible results on privacy enhancing enabling technologies are expected in year 2. The main outcomes include algorithms and proof-of-concept implementations for biometric template protection and revocable biometrics, as well as attacks and countermeasures for continuous authentication in adversarial contexts. We will also report on the impact of privacy-preserving machine learning algorithms on the accuracy of behavioral authentication.

From year 3 to 5, we will further enhance and strengthen the previous building blocks, by exploring emerging threats and implementing new defenses, while evaluating trade-offs between security, privacy and usability for continuous and collaborative authentication.

3.2. Authorization and Audit

Authorization and audit are both essential to many security and privacy solutions, typically in combination with authentication. When a user has been correctly identified by means of an authentication mechanism, a digital platform will/must subsequently assess whether the user is allowed to execute the intended actions and operations, or to access the requested data items. This evaluation to grant or deny permission is the goal of an authorization service.

It remains a continuing challenge to offer authorization that delivers security and performance, while being robust in an environment with many types of attackers. In addition, authorization solutions must empower multiple stakeholders: end users, platform managers and operators, as well as the owners of the digital applications. The trade-off between performance and security often leads to incomplete decision making that is compensated by the analysis of history that is typically recorded in log files. This is where an audit process complements authorization. Analysis of the audit trail can detect incidents that have not been prevented by an inline authorization engine.

Audit based solutions cover much more ground than the prevention based ones sketched above. A growing number of cybersecurity requirements demands for extensive data recording and analysis. A basic cornerstone is the thorough registration of defined policies, the history of the application, the actions of users, operators etc. The resulting audit trail allows internal stakeholders and possibly external parties to assess the quality of a digital service, and especially its achievements in terms of security and privacy. Additional value has to be harvested for example in the context of forensics, or in the context where new policies can be derived from history, etc. Audit based technologies will in the mid-term play a key role in a gradual process towards modular certification: in the future, specific facets of a security solution could in principle be certified on the basis of a robust audit trail. Such solutions will demand for effective automation to become scalable and widely applicable.

A) Industry Needs – Use Cases and Technology Outlook

A large segment of business and organizations that depend on digital platforms and services rely on popular technologies such as cloud platforms, web-based clients and services and mobile apps. Some industry interest groups organize information sharing and training to enable software developers to prevent some of the most common and prioritized threats while developing and deploying the software system. OWASP is a well-known example. Their “top 10” enumerates such threats in the context of web application development. It is important to notice that 4 threats out of 10 in this ranking are related to weaknesses in authorization and to problems with access control. For example, Threat 4 (labeled “insecure direct object references”) emerges from lacking enforcement of access control rules in the application logic.

Many advances in the state of practice for industry have improved the enforcement of access control rules. Externalized access control enforcement, such as support for RBAC in application servers, or policy-based decision points enable externalization and modularization of authorization logic into dedicated, specialized security services. As such, security specialists and domain experts can focus on the access control rules, rather

than imposing this responsibility on (for example) web developers. Yet the notion of externalized authorization still is a high potential opportunity that often demands for significant improvements and investment in application development. Still correct placement and configuration in the application code of annotations, or the policy enforcement point remains a challenging development concern causing mistakes easily.

Moreover, externalized security policies also require expressive power to assert complex conditions and also require broad application-level context information about the ongoing operation that must be verified. Such context information can include the execution history of a given user. This overall challenge limits the state of practice and still demands for research, e.g., to systematically and effectively determine the right balance between expressive power and efficient evaluation of security policies.

Application-level audit solutions are needed for out of band analysis and reporting, as well as for in-band authorization based on history. In such audit solutions, execution trails can be analyzed for appropriate access control in all possible situations. Such audit solutions also need to be able to analyze precisely which operations were executed on behalf of a specific identity, the permissions that were presented, and the decision made by the security service (permission or denial). Moreover, in such execution trails, the use of privileged accounts, typically identities for back-end services with a broad set of permissions, should be analyzed and accidental, dangerous, misuse of such accounts should be detected. As such, access control relies upon inline verification before execution as well as on out-of-band verification in audit trails.

It appears there also is a need for synergy between authorization and audit. For example, synergy between audit and authorization is relevant in the case of History-Based Access Control (HBAC), when access decisions are based on the real-time evaluation of a history of activities in the audit trail. Keeping the real-time information up to date such that the access decision can occur instantly is hard. This can be approached using in-band evaluation of the history, or out-of-band evaluation using frequent intervals when the analysis of the audit trail cannot be executed in real-time, e.g., when using complex statistical analysis. The synergy is also relevant in the opposite direction. For the purpose of security analytics, one needs to analyze the history of access control decisions in the audit trail to assess correctness of the current implementation of access control rules.

The relevance of the challenges above has been widely recognized in industry. Relevant cases and strong demand appear for example in e-commerce, in finance and the banking sector, in content provisioning, and many other online services.

B) State-of-the-art: Highlights

Access control has been an important research subject for a long time. Amongst others, a lot of effort is spent on designing models to efficiently and correctly specify the permissions of users in a system, e.g., lattice-based access control [LATH1985], role-based access control [FERR2001], attribute-based access control [HU2014] and more recent advances such as usage control [PARK2004] and relationship-based access control [GIUN2008, FONG2001]. These models have then led to formal definitions of their properties (e.g., [BELL1973, BIBA1977]), to supporting administrative models (e.g., [SAND1999]) and methods such as role mining [KUHI2003].

Significant research efforts have been spent to support the reliable incorporation of access control in application code. Amongst others, this has led to the approach of policy-based access control [Slom1994, Sama2001] in which the access rules are declaratively specified in so-called policies. This in turn has led to research on languages for expressing these policies (e.g, Ponder [PONDER], XACML [GODIK2003] and SecPAL [BECK2010]), on combining policies of multiple parties (e.g., [BONA2002]) and on expressing specific rules such as separation-of-duty [Brew1989]. In addition, access control research also encompasses automatic placement of access control code in application code (e.g., [MUTH2012]). Modular implementation of application-level access control is another important Track, for example based on the use of AOP [DEWIN2002]. This research continued in implementing adaptable access control policies and domain-specific access policies supporting rich application context [VH2005]. Applied research further evolved into federated access control over organizations and efficiency aspects of decentralized evaluation of authorization polices. Performance of access policies was further extended for distributed multi-tenant SaaS architectures [DECAT2015]. The supported complexity of access policies has been further addressed in entity-based domain models and data search operations (EBAC [BOG2015] and Sequoia [BOG2018]).

C) Main Areas of Work

The focus of this research theme is on application level authorization and audit in complex distributed systems and applications (e.g., based on IOT, cloud, micro-services). The execution of this research includes three research activities: (1) enhancing the core authorization capabilities to achieve more expressive, correct, and well performing externalized authorization in distributed applications (in micro-service architectures, in IOT, etc.). (2) Enhancing the intelligence of audit solutions that can leverage on techniques to analyze and verify logs and execution histories (in-band and out-of-band) and (3) Creating synergy between audit and authorization by automating feedback loops that can perform static and dynamic verification of both authentication and authorization flows.

(RA 2.2.1) Enhancing Authorization Capabilities

The first activity tackles the goal to support for more complex, and expressive policies. Externalized security policies also require expressive power to express complex conditions and typically require broad application-level context information about the ongoing operation that must be utilized. Such context information can include more domain-specific concepts of specific business domains such as *breaking the glass* in healthcare, or general access control principles such as separation of duty. These requirements obviously have performance implications too.

This leads to a second research goal on performance engineering. In complex distributed systems such as micro-service architectures, there is a need for managing the security-performance trade-off using a systematic approach and method, and this requires support in popular frameworks and platform(s)). This trade-off applies to the complex user and execution context and this requires expressive power and configuration at the level of the policies to reason on and configure the performance, completeness and freshness of attributes in complex context models.

(RA 2.2.2) Intelligent Audit

There is a need for intelligence and analytics on application-level audit trails. This should enable security analytics for common use cases, by creating a robust audit trail and supporting its interpretation. For example, in execution trails, the use of privileged accounts (typically identities for back-end services with a broad set of permissions) should be analyzed and accidental, dangerous, potential abuse/misuse of such accounts should be detected. In general, privileged accounts and related attributes must be considered important attack vectors (e.g., hidden attribute attacks) when attackers aim for escalation of permissions/privilege.

The above example of course requires audit and analysis of authentication flows, authenticated interactions and identity flows in complex distributed systems like micro-service architectures. Hence, there is a need to create consistent and protected audit trails in heavily decentralized systems like micro-service architectures and IOT that can serve as a base for compliance verification and even legal prosecution. Creating such an audit trail in a well-performing, non-blocking way in such a decentralized system, while ensuring consistency and correctness is an important challenge.

(RA 2.2.3) Synergy between Audit and Authorization

There is a need for the verification of correctness and completeness of externalized authorization. Externalized authorization (such as support for RBAC in application servers, or policy-based decision points) enables externalization and modularization of policies and their enforcement into dedicated, specialized security services. Yet correct placement of controls in the application code, and configuration of annotations (or the policy) still is a challenging task. Static code-verification methods should be complemented with dynamic verification based on distributed execution trails to ensure correctness and completeness of both the specification and enforcement of access policies. First, there is indeed this need to assess the correctness of the actual execution trail. Secondly, also suggestions for improvement and performance optimizations will be generated from the analysis of the audit trail.

This synergy also enables the delivery of more advanced audit analytics information on application and user history towards the access control system. The goal is to achieve real-time in-band access control based on in-band and out-of-band analytics of audit trails containing user and application history. This major goal will leverage the results of both RA1 and RA2: the efficient collection of distributed audit trails, the support for statistical analysis and advanced analytics on the audit trail, and the efficient encoding and configuration as contextual information towards the access policy and access decision point.

D) Expected Outcomes and Road Map

Research Activity 1 will yield important results by the end of Y1, including enhanced support for domain-specific policies. The Consortium will focus for example on supporting expressive power for regulatory frameworks. In addition, the programme will yield a prototype of efficient policy execution environments with performance guarantees (Y1). These initial results will be further extended into Year 2. Y2 will add focus on audit though. Initial results on audit include fine-grained protected audit trails and efficient analytics of application-level audit trails in cloud architectures (Y2). As of Year 3, the work will elaborate on synergy between audit and authorization, aiming for two main results: (1) runtime correctness assessment of authorization policies and enforcement points, and (2) an integrated framework for real-time in-band authorization based on out-of-band and in-band analytics on audit trails.

3.3. Advanced Encryption Techniques and Data Access Middleware

Over the last years, outsourcing of data storage and processing has become increasingly popular and at the same time, the levels of trust in the third parties responsible for storing and processing this data has decreased. Moreover, in order to use the data we store, we need to have a reliable and secure way to access and operate these data. This requires both advanced cryptographic solutions to operate on data in untrusted environments, as well as advanced data access firewalls towards applications that support fine-grained security policies in data access middleware.

Classical cryptographic solutions enable to secure data at rest (data storage) or data in transit (secure communication). A novel strategy is secure data processing, where the data are guaranteed to remain private even during processing, either by splitting the data over several servers and performing a joint computation (secure collaborative data processing), or to process the information in encrypted form (secure outsourced computation). This novel strategy relies on recent developments of advanced cryptographic techniques, a research area known as COED (computing on encrypted data).

Secure collaborative data processing can be implemented using secure multiparty computation (MPC). MPC allows a group of parties to compute some arbitrary function f on the parties' private inputs, while preserving a number of security properties such as privacy and correctness. The former property implies data confidentiality, namely, nothing leaks from the protocol execution but the computed output. The latter requirement implies that the protocol enforces the integrity of the computations made by the parties, namely, honest parties are not led to accept a wrong output.

Secure outsourced computation corresponds to a secure version of computation as a service (CaaS) and allows to process data in encrypted form using a technology called homomorphic encryption (HE). The main difference with MPC is that it only involves one processing server (which does not need to be trusted). HE enables outsourcing of encrypted data to a single processing server, which will perform computations on a user's behalf and return the result in encrypted form, without gaining any knowledge about the data it is processing. Despite a decade of efficiency improvements, HE is still rather limited to very specific application scenarios.

The above cryptographic techniques are orthogonal to and can be complemented with data access middleware (DAM), which can also support data filtering and query rewriting: i.e. the fine-grained selection of data elements depending on the access rights of the user and the security policy of the owner of the data. These techniques also require further research to move from practical feasibility results, to a stage in which we have fully practical systems.

A) Industry Needs – Use Cases and Technology Outlook

The tension between the value of data in all aspects of our society (economy, social protection, health, security, etc.) and the protection of fundamental rights is ever increasing. The birth of regulatory frameworks such as the GDPR shows that privacy-preserving technologies will be a crucial element in guaranteeing that we can simultaneously benefit from the data and computational power that we have available and safeguard the right to privacy of EU citizens.

Cloud services have become an integral part of how businesses operate, and it is estimated that by 2020 more than 83% of enterprises workloads are driven by the cloud. As such, there is already a huge amount of

sensitive data residing on cloud platforms such as health data, election data, business trade secrets etc., which makes cloud platforms ideal targets for both internal and external hackers. In such scenario, the data owners lack control over what happens with their data and they either have to rely on the security measures put in place by the cloud provider or somehow secure the data themselves. A standard solution to ensure data confidentiality is to encrypt the data by a key only known to the data owner. The problem with this approach is that it voids any extra services that the cloud provider offers, even simple operations such as search or statistical processing, which would require intermediate decryption of the data. Cloud providers are therefore looking for methods to enhance their services with privacy preserving capabilities.

Another approach commonly seen to mitigate the risk of data breaches is to adopt a distributed storage architecture, where a large logical database is split over several smaller databases managed by different departments. While this approach avoids super large data breaches, it complicates processing of the data which is now spread over multiple databases. A related problem is pooling data from different data sources to perform value adding statistical computations, e.g., gathering data from different hospitals, government organizations or even across national borders. There are obvious privacy and security concerns with such pooling, but in many cases, due to regulatory restrictions, data that was gathered for one purpose cannot be reused for a different purpose. Various national agencies, and indeed the United Nations, as part of the United Nations Global Data Platform, are now looking at using systems to help increase the value of collected statistics, whilst maintaining privacy.

While the long term vision on data protection is to encrypt the data and perform the processing on the encrypted format, the state-of-practice will require intermediate decryption for a while. This practical limitation demands for data access control that limits the visibility and access to data when accessing a variety of database technologies. Ideally, a practical solution would combine advanced encryption when feasible, and compensate for limitations by using traditional protection layer. This sketches the combined role that is to be full-filled by data access middleware.

Data access middleware (DAM) is the software layer between a distributed application (e.g., a web application or web API) and a distributed data system (NoSQL, RDBMS). While the historic role of DAM has mainly been the one of object-relational data mapping (e.g., Hibernate, JPA), there is strong a demand to apply data security tactics in this middleware. This demands relates to two security concerns: (1) the lack of trust of the end-users by the application service provider and organizational customers and (2) the lack of trust of the execution environment by the application service provider and its customers. There is an important need to easily apply and configure application-level data security tactics like searchable encryption, homomorphic encryption or advanced data filtering to tackle the attack vectors from both the end-user side as well as the hosting side.

B) State-of-the-art: Highlights

The research related to this theme addresses the development of practical solutions for processing encrypted data in different settings. We will focus on three promising data protection tactics, each of these having its unique features, advantages and shortcomings: (i) multi-party computation, (ii) (fully) homomorphic encryption and (iii) policy driven-query rewriting in data access middleware.

Secure multi-party computation (MPC) started with the foundational paper of Yao [YAO86], who presented a protocol which will securely compute any functionality involving two parties using so-called garbled circuits. Yao's protocol however is limited to two parties and is not secure against active malicious adversaries, and overcoming these limitations has been the focus of the existing literature. The work by Beaver et al. [BMR90] looked at how one can take the key advantage of Yao like protocols (namely the fact they are constant round) and extend this to the case of more than two players. A long large body of work has led to very efficient protocols in both the two-party [WRK17A] and multiparty [HSS17, WRK17B] settings.

The pioneering work of Ben-Or, Goldwasser, Wigderson, [BGW87] and Chaum, Crepeau and Damgård [CCD87] showed how one can build MPC protocols for three or more players from secret sharing schemes. This is called the "secret-sharing" or "BGW/CCD" approach. Unlike Yao's protocol these schemes are secure in an unconditional sense, but the round complexity is dependent on the multiplicative depth of the circuit. The most famous success of this line of work has been in conducting an auction of sugar beet production certificates for Danish farmers, the world's first real life deployment of MPC technology.

An important technique in secret-sharing MPC was introduced by Beaver [BEA91], involving a so-called offline, or pre-processing phase, where the players generate correlated randomness that is independent of the secret data to be computed on. This work has received a lot of attention recently because of two advancements: On one hand, homomorphic encryption allows one to generate the required randomness for computations in any finite field and any number of players. This is the core of the so-called SPDZ protocol [SPDZ12]. On the other hand, batching oblivious transfer (OT) provides a particular efficient way of generating the randomness for binary circuits ("TinyOT" [NNOB12]). With both protocols types we obtain fully malicious security with aborts, a protocol which scales linearly with the number of players, and a highly efficient online protocol.

These protocols and many improvements have been integrated into the SCALE/MAMBA [SCALE19] project developed at KU Leuven, which implements the FHE and OT variants of SPDZ in the dishonest majority setting as well as these honest majority protocols.

(Fully) Homomorphic encryption. The notion of homomorphic encryption was introduced by Rivest et al. [RIVEST1978] and was first instantiated for restricted classes of functions. In 2009, Gentry published his seminal thesis [GEN09] introducing the first plausible FHE scheme based on hard problems in ideal lattices. Even though this scheme is recognized as a profound theoretical breakthrough, it is highly inefficient due to large parameters. Following Gentry's blueprint, many researchers have been trying to develop more practical FHE schemes.

Currently, the most efficient FHE schemes are Brakerski-Gentry-Vaikuntanathan [BGV12], Fan-Vercauteren [FV12] and HEAAN [CKKS17] schemes. These schemes exploit state-of-the-art optimization techniques such as modulus switching and relinearization that efficiently suppress the noise growth after homomorphic multiplication. In addition, the ciphertext-to-plaintext expansion ratio of these schemes is significantly decreased by encrypting several messages with a single ciphertext [SV14, CKKS17]. In practice, these schemes are considered the most efficient for arithmetic circuits, which is supported by their success in public competitions [IDASH]. However, bootstrapping (which is needed to evaluate arbitrary functions) in these schemes is still slow despite various improvements [GHS12, HS15].

Another interesting family of FHE schemes [GSW13] is characterized by the use of matrix arithmetic. Although these schemes are less efficient than the above schemes for arithmetic circuits, their bootstrapping functions are much faster. The most efficient FHE scheme of this type, TFHE [CGGI16], supports a look-up table functionality that allows to evaluate non-polynomial homomorphic functions and finite automata.

Data access middleware. Recent research on data access middleware has addressed performance aspects, e.g., [REN2017] [RAF2018B], the plug-ability and configurability of data protection tactics and data filtering tactics using advanced application policies. The Persist middleware [RAF2018A] supports policy-driven application and configuration of data protection tactics such as encryption, sharding and replication based on execution context, application-specific metadata and data-specific metadata. Sequoia is a data access middleware that supports a posteriori data filtering on queries, as well as secure a-priori query rewriting to apply user-defined security policies. Query rewriting involves manipulating database queries before execution and allows enforcing data management policies without having an invasive impact on the end-to-end architecture. Recent work, Sequoia [BOG2018] presents an extensible architecture for both relational databases and document stores. It discusses the rewriting approach, and provides a formal verification of equivalency and an extensive evaluation that shows that this approach scales better than the current state of practice. This defines an important path for future research.

C) Main Areas of Work

(RA 2.3.1) Secure Outsourced Data Processing

Despite a decade of improvements, full blown FHE, where arbitrary functions can be evaluated efficiently, is currently still out of reach. However, for specific applications, practical solutions are available and the overall goal of this Research Activity is to push the applicability and practicality to the next level. More in particular, we will consider the following research topics. (1) We will investigate the relation and conversions between the different existing schemes and investigate whether it is possible to port the advantages of one scheme to another, e.g., schemes based on TFHE allow efficient table lookup and thus evaluation of highly non-linear functions, which are currently lacking in the other schemes. (2) We will catalogue existing and develop new efficient data representations for a particular basic operation and data type, which will mitigate the mismatch

between the native plaintext space and real life data types. (3) We will extend existing software libraries with above improvements and develop functionalities that allow statistical data processing on encrypted data, and illustrate these with demonstrators manipulating encrypted financial data.

(RA 2.3.2) Secure Collaborative Data Processing

MPC technology has progressed immensely in the last few years and one can now claim that the practical feasibility of MPC has been established. The next task is to scale up the potential application areas, as well as the performance and data throughput of MPC systems. Thus whilst feasibility is established a lot of basic research still needs to be done, in particular: (1) we will examine ways to combine low round complexity and high round complexity protocols in an efficient manner. (2) We will examine a number of applications of MPC techniques to areas such as post-quantum threshold cryptography, statistical calculations, data processing, and side-channel protection. (3) We will be looking at both new theoretical improvements to protocols, as well as implementing and testing the most promising using our existing SCALE-MAMBA MPC system.

This research effort will be supported by practical demonstrators in the area of secure collaborative processing of shared data, as required by applications in health, e.g., genomic data analysis and finance, e.g., fraud detection.

(RA 2.3.3) Data Access Middleware

First, decentralized middleware architectures for secure data access must be further explored, i.e., architectures with multiple federated parties that contribute to a data query or data analysis while using a decentralized data set. Such an approach typically requires advanced encryption like homomorphic encryption or multi-party computation.

Second, we will investigate more client-centric decentralization and data access models in which the client specifies policies on which data can be shared and analyzed by the client in community-driven data networks.

Third, secure query rewriting currently only focuses on the fine-grained selection of data items based on security constraints. Obligations, i.e. actions that need to be undertaken in case of permitted access, are typically not supported, as such specification capabilities are generally not provided by the query languages. Support for fine-grained audit trails is still needed for secure data access.

Fourth, the rewriting approach is also limited with regard to support for several database models. Some database models, particularly wide-column stores (e.g. Cassandra) and key-value stores (e.g. Redis), typically enhance scalability through constraining the capabilities (expressive power) of the supported queries. This limits the applicable rewriting approach for those systems. Additional research should investigate how the rewriting approach can be partially applied and complemented with an ”a posteriori” filter to alleviate this matter.

D) Expected Outcomes and Road Map

For (RA 2.3.1), the first expected outcome (Y1) is a catalogue of combinations of common operations and data types and the best HE method to deal with this combination. Furthermore, we will also provide detailed estimates of the costs incurred by moving from one HE scheme to another. The second expected outcome (Y2) is a software library providing all common statistical processing subroutines, but operating on encrypted data.

For (RA 2.3.2), the first expected outcomes (Y1) are the integration of garbled circuit and LSSS operations within SCALE-MAMBA and the development of threshold variants for a number of NIST PQC submissions. The expected outcomes for the second year (Y2) are to combine garbled circuit and LSSS operations to perform efficient scientific operations, such as floating point calculations. And to support end users to provide MPC solutions in a number of real world situations using SCALE-MAMBA.

For (RA 2.3.3), the first expected outcome (Y1) is the development of a decentralized data access middleware for advanced data encryption tactics. The expected outcomes for the second year (Y2) are the development of data access middleware for client-centric decentralized systems and the combination with policy-driven data protection tactics.

3.4. Policy and Regulation

In the last ten to fifteen years, the European Union has adopted several instruments to harmonize obligations and responsibilities in relation to information and network security, and to support an EU-wide coordinated response to cyberattacks. The EU is urging its Member States to adopt the necessary legislative and organizational measures to address cyber-attacks and to restore trust in connected products, IoT devices and cloud services. This is witnessed by a series of directives ranging from the 2002 ePrivacy Directive and the 2008 Directive on European Critical Infrastructures, over the 2013 Directive on attacks against information systems (the 'EU Cybercrime Directive') and the 2016 Directive on Security of Network and Information Systems ('EU NIS Directive'), to the recently adopted Cybersecurity Act (April 9, 2019). The Cybersecurity Act reinforces the mandate of the EU Agency for Cybersecurity, (European Union Agency for Network and Information and Security, ENISA) so as to better support Member States with tackling cybersecurity threats and attacks. The Act also establishes an EU framework for cybersecurity certification, boosting the cybersecurity of online services and consumer devices. Moreover, the EU legislator in previous years introduced a series of (new or revised) rules with a considerable impact on data processing, sharing and reuse activities of private and public sector actors. Not only the well-known GDPR, but also the Free Flow of (non-personal) Data Regulation, the B2B Platform Regulation, the revised Public Sector Information Directive, and the new Payment Services Directive (cf. the Open Banking/PSD2 standard), have set legal standards in relation to access to data, liability and reliable data portability.

The policy and regulation Research Track is interdisciplinary and transversal, in the sense that its topics cut across the various Research Tracks. However, as the initial focus will be on security services, it has been embedded in Research Track 2, while maintaining also strong links with Research Track 3 on system and infrastructure security.

A) Industry Needs – Use Cases and Technology Outlook

Manufacturers and service providers across sectors will need to understand the legal requirements for, and implications of, cybersecurity, as their business relies more heavily on data processing, data sharing and cloud solutions. A recently published study carried out by KU Leuven CiTiP and LINC shows that businesses in Belgium – from a wide variety of economic sectors – are confronted with at least one type of cybercrime every year, with some of them suffering serious harm from these incidents [PAOLI18]. During the 9th annual Internet-of-Things European Summit on 15-16 May 2018 in Brussels, the European Consumer Organisation (BEUC) underlined the threats IoT devices pose to consumers in terms of security and liability. According to BEUC, during a test hackers achieved almost a 50% success rate in their attempts to hack smart-home devices (e.g. alarm, vacuum cleaner, GPS watch to locate children). Many of today's IoT devices are rushed to market with little consideration for basic security and privacy protections, inducing various levels of risk to both users and the Internet itself (from unwitting surveillance and data compromise to physical risk, e.g., smart locks, to security cameras used as part of a botnet to attack the Internet). In its IoT Trust Framework, the Internet Society's OTA (Online Trust Alliance) elaborates on the notion of "Trust by Design" – an umbrella term that includes Privacy by Design and Security by Design – as an essential component of a healthy IoT ecosystem [OTA17]. Although such initiatives are extremely valuable to make existing legislation more effective, they cannot replace the legislation itself. It therefore remains key to understand the relationship between adhering to multi-stakeholder guidelines and legal compliance, and to ensure consistency with legislation and regulatory decisions, in order to increase legal certainty for companies.

Existing and novel forms of security attacks raise, on the one hand, questions relating to *compliance* with existing laws and regulations (and liability issues in case of non-compliance), and, on the other hand, *policy* questions (i.e. questions about the appropriate policy responses at government and company level). The risk-based approach to data protection and privacy introduced by the GDPR makes compliance less straight-forward. It urges companies to carefully assess the practical impacts and challenges associated with implementing GDPR provisions on 'risk', 'high risk', 'risk assessments' and 'DPIAs'. Moreover, GDPR requires a continuous dialogue between legal and technical experts, by making data-protection-by-design/default ('DPbD'; a holistic approach to embedding principles in technical and organizational measures undertaken by data controllers) a qualified duty for data controllers. It is equally important to understand GDPR's relation with other legislative instruments, in particular the ePrivacy framework and EU NIS Directive, e.g., in relation to incident reporting in case of a personal data breach or cyberattack. The appropriate response by companies

under existing reporting obligations is not "set in stone", as it will depend on the likely risk to individuals resulting from the breach of confidentiality, availability or integrity [WPGUIDELINES17]. Hence, it may not always be obvious to companies what they need to do in order to comply in practice with their legal duties imposed in the context of ePrivacy, GDPR and/or NIS.

B) State-of-the-art: Highlights

The interdisciplinary research theme on policy and regulation aims to support and steer the development of security solutions. It therefore encompasses three types of analysis: legal compliance analysis, policy analysis and legal engineering analysis. Each of those will zoom in on a number of (sub)themes as priority areas, which have been selected in light of the technical research proposed and building further on pre-existing legal expertise and involvement in EU policy debates.

Legal compliance analysis

As highlighted above (in this section), a first research need arises from mapping the legal requirements for new security and authentication solutions in order to comply with GDPR and other relevant rules (such as ePrivacy, NIS, etc.), outlining the interaction between those rules, and clarifying the role of technical standards. This entails an analysis of existing legal standards in relation to (liability for) data access, sharing, and portability, in light of upcoming challenges. It was, for instance, noted (under 'Identity Management and Authentication') that the use of biometrics may lead not only to enhanced security solutions, but also to considerable privacy losses, given the critical information enclosed in biometric templates. Enhanced data protection capabilities are also challenging fundamental concepts under GDPR, like 'data controller', 'data processor', or 'consent'. Can you, – for example – in the context of MPC, still be considered a 'data controller' when you are – using the terminology of the GDPR – 'determining the purposes and means of the processing of personal data', without having gained any knowledge about the data underlying the processing…? It has also been noted that, especially in relation to IoT, the use of sensors for (multi-factor) authentication brings a multitude of legal issues in particular with regard to profiling and valid consent [WACHTER18].

Previous fundamental research on legal aspects of biometric applications [KINDT13], and basic concepts in data protection law – in particular the notions of consent [KOSTA13], data controller-processor [VAN ALSENOY16] and the purpose specification principle [COUDERT19] have paved the way. Relevant focal areas can be on the one hand, the use of biometrics (short term) and on the other hand, and data protection and responsibility in case of multi-factor authentication (long term).

Short term: use of biometrics

A tailored approach to the use of biometric traits (such as the use of facial images or fingerprints) in security solutions is needed in order to strike a balance between security and privacy characteristics. Using reduced biometric information by deploying biometric templates is already a step forward in protecting critical information contained in these traits. Also the use of so-called 'protected templates', proposed about a decade ago, is a step in the good direction [BREEBAART08]. The qualities of such protected biometric information has been subject of standardization in the meantime but take-up remains slow. This could be overcome with best practices guidelines [KINDT10]. One of the most urgent security risks nowadays, is the creation of virtual identities by the mixing of face or fingerprint of various persons in one identifier, which is then fit for use by several persons/perpetrators [OTHMAN11]. Biometric systems should in general have adequate detection measures in place which reveal such morphing and when unauthorized persons spoof the system. Presentation attack detection (PAD) is therefore increasingly relevant and required in for example the domain of (mobile) payments [KINDT19]. Another main risk is the central storage of biometric data in databases for one purpose and then the re-use for other purposes. [KINDT18]. Such central storage is apt to the risk of 'function creep' as mentioned also in the recent technical analysis [HERMANS19]. The need to address this risk of central storage was also urged by several European Data Protection Supervisory Authorities over the last 15 years but is not clearly addressed in the GDPR or any Belgian legislation. The research shall also unveil risks of new and imminent modalities, such as speech. Speech offers convenience and opens up new deployments in the interconnected world, but will also require 'legal engineering', whereby protection is embedded into the technical solutions [NAUTSCH19]. This angle will allow to bridge the legal compliance analysis Track with the legal engineering Track.

Mid term: data protection and responsibility in case of multi-factor authentication

For many legal acts on the internet, e.g., electronically signing a document or buying something, it is often necessary to identify a person. Different systems for authentication exist, including the use of social media accounts, which bring their own legal questions [SCHROERS18]. Increasingly multi-factor authentication is developed to authenticate a person, including the use of a variety of sensors. This development needs to be further researched, especially in relation to privacy and data protection but also regarding the allocation of responsibility. A central problem to be considered is the fact that it will be increasingly difficult for the user to know which information is used for authentication, and the possibility of function creep [TSORMPATZOUDI15]. Under current authentication systems users are usually aware that they are in the process of authentication and have at least a vague idea which authentication factors are used, e.g., username and password or the use of a smartcard and PIN. Especially in case of a 'frictionless' authentication, even though it can be desirable for usability purposes, the user might not be aware of the processing that takes place. An authenticated person is in general considered being responsible for what is done 'on his or her behalf'. An important legal question is therefore who is held responsible in case there is a failure of the authentication system [SCHROERS&TSORMPATZOUDI2016]. No system is perfect, and failures can occur due to technical or human factors. However, in increasingly complex systems it becomes often difficult to identify the cause for a failure, especially when the user has less knowledge and control over the process. Similarly, to autonomous driving cars, this leads to the consideration whether for normal users a shift in the allocation of liability or the burden of proof would increase trust in the system. It is important to identify which kind of obligations the different parties in an authentication system usually have, but also upon whom generally the burden of proof rests. The eIDAS Regulation, which was enacted in 2014 and is fully in force since 2018, includes certain provisions on electronic identification and authentication [EIDAS REGULATION]. However, this only relates to cross-border national electronic identification schemes [CUIJPERS&SCHROERS2014]. Whereas in some countries national regulation on electronic identification exists (albeit often only in relation to national electronic identification systems), the majority of obligations derive from the terms and conditions of the identification providers.

Policy analysis

A second research need lies in the pro-active monitoring of, and contribution to, ongoing policy debates at EU and national level, for instance in relation to data-driven intelligent cybersecurity tools. Also with regard to the implications of the EU Cybersecurity Act in relation to certification schemes and the development of a coordinated approach to software vulnerability disclosure, discussions are ongoing and require our close attention. The focus of this research theme may shift as policy discussions become more prominent, or are being put on hold (to give an example: the ePrivacy reform was subject of intense discussions until recently, but those have been halted as a result of the EU elections). In a first phase, the emphasis will lie on two areas where we can build on previous experience, *software vulnerability disclosure* (short term) and *framework principles for data-driven intelligent cybersecurity tools* (long term).

Short term: A coordinated approach to software vulnerability disclosure

In the area of software vulnerability disclosure, there is – despite ENISA's Good Practice Guide on Vulnerability Disclosure of January 2016 [ENISA16] – still no solid Belgian policy yet on how to deal with ethical hacking and coordinated vulnerability disclosure [CEPS18A], nor a common European approach [CEPS18A, CEPS18B]. Software vulnerabilities patching is gaining a growing role for the security of our ever-connected systems and infrastructures. At the moment, a great number of European Member States (including Belgium) have still not implemented any transparent policy on the coordinated disclosure of software vulnerabilities [FANTIN18]. The EU Cybersecurity Act, recently approved by the EU Parliament, timidly addresses coordinated disclosure processes by mandating them to both Member States' voluntary initiatives and the potential coordination role of ENISA [FANTIN19]. These elements concur to create a fragmented landscape, since, as a matter of fact, it is still governed by non-binding standards or industry-specific practices. Additionally, the role of independent security researchers ('ethical hackers') is often not waived in national criminal laws, making them in principle liable for their potential intrusions into proprietary software and applications, even when done in bona fide. For these reasons, research on the main vulnerabilities disclosure practices adopted by industry and system providers would be of added value for a twofold reason: firstly, because it would assess and recommend what principles such self-regulatory frameworks share in common; secondly, because it would clarify how the relationship between software producer and ethical hacker could be improved at the liability level, in order to clarify when ethical hackers' activities would not fall under the violation of cybersecurity provisions.

Mid term: Framework Principles for Data-Driven Intelligent Cybersecurity Tools
The fast pace of big data applications and data analytics has also a valuable potential in the cybersecurity domain. Data driven intelligent threat detection operations can optimize and significantly reduce the response time to an imminent attack. But whilst the benefits of such a use seem widely recognized in the threat intelligence domain, policy and regulation concerns still pervade the future of these tools in information security applications with broader questions and from different angles. To name a number of areas, first and foremost, questions arise on how shall we keep our cybersecurity algorithms unbiased, impartial and effective (also, shall algorithmic discrimination be a criterion for the development of algorithms merely used for cybersecurity purposes?). From a cyber-forensics perspective, concerns may raise on the validity and the admissibility of data generated by the tools before a court of law, since questions on what standards shall such applications embed in order to produce court-proof information are still unresolved. Lastly, from a dual use standpoint, many argue that legal dilemmas may come out with respect to the categorization of the data driven intelligent cybersecurity tools and products in the list of those items subject to strict scrutiny and export control authorizations by national authorities and regulations. All such questions highlight the lack of established and consolidated legal and ethical principles in this context, particularly with reference to the use in cybersecurity products and practices.

Legal Engineering Analysis
Thirdly, the trend towards 'techno-regulation' (in the sense of embedding norms in the technology [LEENES11]) has instigated a growing need for interdisciplinary research that bridges legal and technical standards. An illustration of such 'techno-regulation' relates to data-protection-by-design ('DPbD'), which is now a qualified duty in the GDPR. DPbD entails a holistic approach to embedding principles in technical and organizational measures undertaken by data controllers. However, practitioners tend to see DPbD less holistically, instead framing it through the confidentiality-focused lens of privacy enhancing technologies (PETs) only, and omitting the aspect of privacy-as-control [VEALE18]. This may leave data re-identifiable by capable adversaries while heavily limiting data controllers' ability to provide data subject rights, such as access, erasure and objection, to manage this risk.

Embedding legal values into the design of technical systems has also been put forward in relation to security ('security-by-design', see the aforementioned example of OTA's IoT Trust Framework [OTA17]) and ethics ('ethics-by-design', which has become quite prominent in ongoing discussions on artificial intelligence [IEEE19]). Challenges involved in such interdisciplinary research have been described in the literature [KOOPS11, LEMÉTAYER&COUDERT17] and are gradually overcome [EMANUILOV18], as can be illustrated by recent results obtained in the framework of the PRiSE project (Privacy-by-design Regulation in Software Engineering, KU Leuven, 2017-2021) [SION19]. Legal engineering research not only leads to increased legal compliance, but also to more complete and more efficient compliance.

C) Main Areas of Work
The execution of this research theme comprises of legal compliance analysis (RA 2.4.1), policy analysis (RA 2.4.2), and legal engineering analysis (RA 2.4.3) in the areas outlined above and following the methodology described in the subsequent paragraphs.

(RA 2.4.1) Legal Compliance Analysis
The main goal of this research activity is to offer a legal analysis of EU and national laws and regulations in relation to cybersecurity and cybercrime. It will produce a deepened understanding of the duties for, and liabilities of, different actors (from industry, government and law enforcement) with regard to the prevention and prosecution of cyberattacks. For each of the themes outlined above, it consists of the following methodological steps:

- 1° mapping existing legal requirements in relation to cybersecurity;
- 2° studying the role of technical standards – such as ISO/IEC 24760 and ISO/IEC 29146 on identity and access management) and voluntary initiatives (like OTA's IoT Trust Framework) in legal compliance (in close interaction with technical researchers);
- 3° identifying areas which are open for interpretation, which are contradictory, or which are not covered;
- 4° applying (mainly descriptive and evaluative) legal methods in order to outline the different interpretation options and list their pros and cons for different actors.

The research will take into account laws that apply horizontally (i.e. across all sectors, such as GDPR, EU Cybercrime Directive, NIS Directive), as well as sector-specific rules (e.g., ePrivacy in telecommunications, Radio Equipment Directive, medical devices, car safety, product liability).

(RA 2.4.2) Policy Analysis

The main goal of this research activity is to monitor and steer ongoing policy discussions in relation to cybersecurity and cybercrime, adopting not merely a descriptive, but also normative approach. In terms of methodology, it entails:

- 1° identifying market failures and regulatory hurdles that hinder progress in cybersecurity (e.g., fragmentation across EU Member States in the area of software vulnerability disclosure);
- 2° assessing the implications of recent and upcoming policy initiatives (e.g., EU Cybersecurity Act) for Belgian/Flemish stakeholders;
- 3° analyzing the specificities of the Belgian/Flemish legal and economic eco-system;
- 4° contributing to policy discussions at national and EU level.

(RA 2.4.3) Legal Engineering Analysis

The main goal of this interdisciplinary Research Track, situated at the intersection of law and engineering, is to zoom in on those legal and ethical requirements in the field of data and cybersecurity which entail the embedding of legal values into the design of technical systems (cf. GDPR's data-protection-by-design and by-default principles, 'security-by-design'). It will be carried out in a multidisciplinary context and it will deepen insights into legal engineering techniques by adopting a bidirectional approach:

- breaking down existing legal requirements (e.g., resulting from GDPR) from a logical, machine-executable perspective to assist data scientists and engineers in developing automated compliance tools ('studying law through the lens of the engineer');
- assessing the robustness and appropriateness of security technologies for compliance with legal obligations ('studying technology through the lens of the lawyer').

It will entail:

- 1° examining current model risk management frameworks and studying practical, implementable best practices;
- 2° assessing the legal robustness of specific security technologies and practices through small-scale empirical tests with technical researchers (cf. Facebook investigation for Belgian Privacy Commission);
- 3° bridging technical standards with legal requirements and converting regulatory context to tools, methods, actionable guidelines and best practices (e.g., on how to make inevitable trade-offs in DPbD more explicit and transparent through Data Protection Impact Assessments);
- 4° developing new models and templates where appropriate (building on experience gained in e.g., the PriSE project).

D) Expected Outcomes and RoadMap

For RA1, the first expected outcome (Y1) is a map of legal requirements outlining mutual interdependencies and interactions between actors (under the form of a written report with tables). The second deliverable (Y2) will incorporate the map and elaborate on interpretation issues resulting in liability gaps and overlaps. For RA2, the first expected outcome (Y1) is an overview of recent and upcoming policy initiatives with a relevance score for Flemish stakeholders; a second outcome (Y2) will consist of a white paper reflecting specific concerns of Flemish stakeholders with a view to contributing to ongoing policy discussions. Finally, for RA3, the first expected outcome (Y1) is the selection of an appropriate case study to test legal engineering solutions in practice. The second expected outcome (Y2) is the development of a tool kit with actionable guidelines and best practices.

3.5. Connections with other Research Tracks

Connection with Track 1: there is a connection between (RA 2.2.1) in this Research Track and (RA 1.2.3) of Research Track 1. The Track addresses security services offering enforcement of access control policies in an externalized and modular way. Track 1 (Theme 2) addresses runtime verification of application-level security policies in general: from this perspective, access control and information flow control policies are two cases.

As policy enforcement is crosscutting through the software stack, both efforts will turn out to be relevant and beneficial. Obviously, we expect synergies and opportunities for collaboration.

Connection with Track 4: Research in Theme 2 on authentication will interact with the research on cryptographic protocols from Track 4 (Theme 3), for example in the case of privacy preserving authentication. Furthermore, the research on cryptographic algorithms and protocols of Track 4 (Themes 2 and 3) is at the basis of the advanced encryption techniques for data protection (Theme 3 of this track, featuring MPC and homomorphic encryption).

The impact of Theme 4 (Policies and Regulation) on other tracks is obvious as this research is an important element to ensure that technical results can be applied in a context that respects new legal provisions that directly affect business. Clearly this value is not unique to Research Track 2, it will certainly affect and interact with other tracks, definitely including Track 1 (think of the management of vulnerabilities) and Track 3 (think for instance of the management of monitoring data).

3.6. References

Identity Management and Authentication

[AbSa17] S. Abhishek Anand, Nitesh Saxena: Coresident evil: noisy vibrational pairing in the face of co-located acoustic eavesdropping. WISEC 2017: 173-183

[AnDB06] R. Andreão, B. Dorizzi, J. Boudy. (2006). ECG Signal Analysis through Hidden Markov Models, IEEE transactions on bio-medical engineering. Volume 53.

[ArAl12] E. Argones Rúa and J. L. Alba Castro, "Online Signature Verification Based on Generative Models," in IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics), vol. 42, no. 4, pp. 1231-1242, Aug. 2012.

[Argo12] E. Argones Rúa et al. (2012), Biometric Template Protection Using Universal Background Models: An Application to Online Signature, in IEEE Transactions on Information Forensics and Security, volume 7, number 1, pp. 269-282.

[CCKT10] Srdjan Capkun, Mario Cagalj, Ghassan Karame, Nils Ole Tippenhauer: Integrity Regions: Authentication through Presence in Wireless Networks. IEEE Trans. Mob. Comput. 9(11): 1608-1621 (2010)

[Dodi04] Y. Dodis et al. (2004), Fuzzy Extractors: How to Generate Strong Keys from Biometrics and Other Noisy Data, in Proceedings of Advances in Cryptology - EUROCRYPT 2004, Interlaken, Switzerland, 2-6 May 2004.

[HMSX12] Tzipora Halevi, Di Ma, Nitesh Saxena, Tuo Xiang: Secure Proximity Detection for NFC Devices Based on Ambient Sensor Data. ESORICS 2012: 379-396.

[HoTr17] H. V. Hoang and M. T. Tran. (2017), DeepSense-Inception: Gait Identification from Inertial Sensors with Inception-like Architecture and Recurrent Network, 13th International Conference on Computational Intelligence and Security (CIS), Hong Kong, pp. 594-598.

[JuWa99] Ari Juels and Martin Wattenberg. (1999), A fuzzy commitment scheme. In Proceedings of the 6th ACM conference on Computer and communications security (CCS '99). ACM, New York, NY, USA, pp. 28-36.

[KMSC15] Nikolaos Karapanos, Claudio Marforio, Claudio Soriente, Srdjan Capkun: Sound-Proof: Usable Two-Factor Authentication Based on Ambient Sound. USENIX Security Symposium 2015: 483-498

[MaWe07] Rene Mayrhofer, Martyn Welch: A Human-Verifiable Authentication Protocol Using Visible Laser Light. ARES 2007: 1143-1148

[Odin10] I. Odinaka et al. (2010), ECG biometrics: A robust short-time frequency analysis, IEEE International Workshop on Information Forensics and Security, Seattle, WA.

[PyKM17] B Pyakillya, N Kazachenko, N Mikhailovsky. (2017), Deep Learning for ECG Classification, Journal of Physics: Conference Series. Volume 913.

[ReQD00] Douglas A. Reynolds, Thomas F. Quatieri, Robert B. Dunn, Speaker Verification Using Adapted Gaussian Mixture Models, Digital Signal Processing, Volume 10, Issues 1–3, 2000, Pages 19-41

[ScSi13] Dominik Schürmann, Stephan Sigg: Secure Communication Based on Ambient Audio. IEEE Trans. Mob. Comput. 12(2): 358-370 (2013)

[Segu17] R. San-Segundo et al. (2017), I-vector Analysis for Gait-based Person Identification using Smartphone Inertial Signals, Pervasive and Mobile Computing, volume 38, pp. 140-153.

[SiPr07] Dave Singelée, Bart Preneel: Key Establishment Using Secure Distance Bounding Protocols. MobiQuitous 2007: 1-6

[Staj11] Frank Stajano: Pico: No More Passwords! Security Protocols Workshop 2011: 49-81.

[StMa17] Jack Sturgess, Ivan Martinovic: VisAuth: Authentication over a Visual Channel Using an Embedded Image. CANS 2017: 537-546

[TGSS14] Hien Thi Thu Truong, Xiang Gao, Babins Shrestha, Nitesh Saxena, N. Asokan, Petteri Nurmi: Comparing and fusing different sensor modalities for relay attack resistance in Zero-Interaction Authentication. PerCom 2014: 163-171

[TVFO18] R. Tolosana, R. Vera-Rodriguez, J. Fierrez and J. Ortega-Garcia, "Exploring Recurrent Neural Networks for On-Line Handwritten Signature Biometrics," in IEEE Access, vol. 6, pp. 5128-5138, 2018.

[VAPJ18] Tim Van hamme, Enrique Argones Rúa, Davy Preuveneers, Wouter Joosen, Gait template protection using HMM-UBM, BIOSIG, Proceedings of the Special Interest Group on Biometrics and Electronic Signatures, IEEE, Lecture Notes in Informatics (LNI), October 10, 2018

Authorization and Audit

[Beck2010] Becker, M. Y., Fournet, C., and Gordon, A. D. SecPAL: Design and Semantics of a Decentralized Authorization Language. J. Comput. Secur. 18, 4 (Dec. 2010), 619–665.

[Bell1973] Bell, D. E., and LaPadula, L. J. Secure computer systems: Mathematical foundations. Tech. rep., DTIC Document, 1973.

[Biba1977] Biba, K. J. Integrity considerations for secure computer systems. Tech. rep., DTIC Document, 1977.

[Bog2015] Bogaerts, J., Decat, M., Lagaisse, B., & Joosen, W. (2015, December). Entity-Based Access Control: supporting more expressive access control policies. In Proceedings of the 31st Annual Computer Security Applications Conference (pp. 291-300). ACM.

[Bog2018] Bogaerts, J., Lagaisse, B., & Joosen, W. I. (2018). SEQUOIA: a middleware supporting policy-based access control for search and aggregation in data-driven applications. IEEE Transactions on Dependable and Secure Computing.

[Bona2002] Bonatti, P., De Capitani di Vimercati, S., and Samarati, P. An Algebra for Composing Access Control Policies. ACM Trans. Inf. Syst. Secur. 5, 1 (Feb. 2002), 1–35.

[Brew1989] Brewer, D., and Nash, M. The Chinese Wall security policy. In IEEE Security and Privacy (May 1989)

[Decat2015] Decat, M., Lagaisse, B., & Joosen, W. (2015, December). Scalable and secure concurrent evaluation of history-based access control policies. In Proceedings of the 31st Annual Computer Security Applications Conference (pp. 281-290). ACM.

[Dewin2002] De Win, Bart, Bart Vanhaute, and Bart De Decker. "Security through aspect-oriented programming." Advances in Network and Distributed Systems Security. Springer, Boston, MA, 2002. 125-138.

[Ferr2001] Ferraiolo, D. F., Sandhu, R., Gavrila, S., Kuhn, D. R., and Chandramouli, R. Proposed NIST Standard for Role-based Access Control. ACM Trans. Inf. Syst. Secur. 4, 3 (Aug. 2001), 224–274.

[Fong2001] Fong, P. W. Relationship-based Access Control: Protection Model and Policy Language. In Proceedings of the First ACM Conference on Data and Application Security and Privacy (New York, NY, USA, 2011), CODASPY '11, ACM, pp. 191–202

[Giun2008] Giunchiglia, F., Zhang, R., and Crispo, B. RelBAC: Relation Based Access Control. In Semantics, Knowledge and Grid, 2008. SKG '08. Fourth International Conference on (Dec 2008), pp. 3–11.

[Godik2003] Godik, S., et al. eXtensible Access Control Markup Language (XACML) 1.0. OASIS Standard (2003).

[HU2014] Hu, V., Ferraiolo, D., Kuhn, R., Schnitzer, A., Sandlin, K., Miller, R., and Scarfone, K. Guide to Attribute Based Access Control (ABAC) Definition and Considerations. NIST Special Publication (2014).

[KUHI2003] Kuhlmann, M., Shohat, D., and Schimpf, G. Role Mining - Revealing Business Roles for Security Administration Using Data Mining Technology. In Proceedings of the Eighth ACM Symposium on Access Control Models and Technologies (New York, NY, USA, 2003), SACMAT '03, ACM, pp. 179–186.

[LATH1985] Latham, D. Department of Defense Trusted Computer System Evaluation Criteria. Tech. rep., US Department of Defense, 1985.

[MUTH2012] Muthukumaran, D., Jaeger, T., and Ganapathy, V. Leveraging "Choice" to Automate Authorization Hook Placement. In Proceedings of the 2012 ACM Conference on Computer and Communications Security (New York, NY, USA, 2012), CCS '12, ACM, pp. 145–156

[PARK2004] Park, J., and Sandhu, R. The UCONABC Usage Control Model. ACM Trans. Inf. Syst. Secur. 7, 1 (Feb. 2004), 128–174.

[PONDER] Damianou, N., Dulay, N., Lupu, E., and Sloman, M. The Ponder Policy Specification Language.

[SAMA2001] Samarati, P., and de Vimercati, S. Access Control: Policies, Models, and Mechanisms. In Foundations of Security Analysis and Design, R. Focardi and R. Gorrieri, Eds., vol. 2171 of Lecture Notes in Computer Science. Springer Berlin Heidelberg, 2001, pp. 137–196.

[SAND1999] Sandhu, R., Bhamidipati, V., and Munawer, Q. The ARBAC97 Model for Role-based Administration of Roles. ACM Trans. Inf. Syst. Secur. 2, 1 (Feb. 1999), 105–135.

[SLOM1994] Sloman, M. Policy driven management for distributed systems. Journal of Network and Systems Management 2, 4 (1994), 333–360.

[VH2005] Verhanneman, T., Piessens, F., De Win, B., & Joosen, W. (2005, December). Uniform application-level access control enforcement of organizationwide policies. In 21st Annual Computer Security Applications Conference (ACSAC'05) (pp. 10-pp). IEEE.

Advanced Encryption Techniques and Data Access Middleware

[BEA91] Efficient Multiparty Protocols Using Circuit Randomization. D. Beaver. In CRYPTO 1991, 420-432, Springer LNCS 576, 1991.

[BGV12] Brakerski, Z., Gentry, C., and Vaikuntanathan, V. (Leveled) fully homomorphic encryption without bootstrapping. In ITCS 2012: 3rd Innovations in Theoretical Computer Science (Jan. 2012), S. Goldwasser, Ed., Association for Computing Machinery, pp. 309–325.

[BGW87] Completeness theorems for non-cryptographic fault-tolerant distributed computation. M. Ben-Or, S. Goldwasser and A. Wigderson. In 20th STOC, 1-10, 1988.

[BOG2018] Bogaerts, J., Lagaisse, B., & Joosen, W. I. (2018). SEQUOIA: a middleware supporting policy-based access control for search and aggregation in data-driven applications. IEEE Transactions on Dependable and Secure Computing.

[BMR90] Donald Beaver, Silvio Micali, and Phillip Rogaway. The round complexity of secure protocols (extended abstract). In 22nd ACM STOC, pages 503–513. ACM Press, May 1990.

[CCD87] Multiparty Unconditionally Secure Protocols. D. Chaum, C. Crepeau and I. Damgård. In 20th STOC, 11-19, 1988.

[CKKS17] Cheon, J. H., Kim, A., Kim, M., and Song, Y. S. Homomorphic encryption for arithmetic of approximate numbers. In Advances in Cryptology – ASIACRYPT 2017, Part I (Dec. 2017), T. Takagi and T. Peyrin, Eds., vol. 10624 of Lecture Notes in Computer Science, Springer, Heidelberg, pp. 409–437.

[CGGI16] Chillotti, I., Gama, N., Georgieva, M., and Izabachène, M. Faster fully homomorphic encryption: Bootstrapping in less than 0.1 seconds. In Advances in Cryptology – ASIACRYPT 2016, Part I (Dec. 2016), J. H. Cheon and T. Takagi, Eds., vol. 10031 of Lecture Notes in Computer Science, Springer, Heidelberg, pp. 3–33.

[GEN09] Gentry, C. A Fully Homomorphic Encryption Scheme. PhD thesis, Stanford University, 2009.

[GHS12] Gentry, C., Halevi, S., and Smart, N. P. Better bootstrapping in fully homomorphic encryption. In PKC 2012: 15th International Conference on Theory and Practice of Public Key Cryptography (May 2012), M. Fischlin, J. Buchmann, and M. Manulis, Eds., vol. 7293 of Lecture Notes in Computer Science, Springer, Heidelberg, pp. 1–16.

[GSW13] Gentry, C., Sahai, A., and Waters, B. Homomorphic encryption from learning with errors: Conceptually-simpler, asymptotically-faster, attribute-based. In Advances in Cryptology – CRYPTO 2013, Part I (Aug. 2013), R. Canetti and J. A. Garay, Eds., vol. 8042 of Lecture Notes in Computer Science, Springer, Heidelberg, pp. 75–92.

[HS15] Halevi, S., and Shoup, V. Bootstrapping for HElib. In Advances in Cryptology – EUROCRYPT 2015, Part I (Apr. 2015), E. Oswald and M. Fischlin, Eds., vol. 9056 of Lecture Notes in Computer Science, Springer, Heidelberg, pp. 641–670.

[HSS17] Carmit Hazay, Peter Scholl, and Eduardo Soria-Vazquez. Low cost constant round MPC combining BMR and oblivious transfer. In Tsuyoshi Takagi and Thomas Peyrin, editors, ASIACRYPT 2017, Part I, volume 10624 of LNCS, pages 598–628. Springer, Heidelberg, December 2017.

[IDASH] iDASH: secure genome analysis competition. http://www. humangenomeprivacy.org/2018/.

[FV12] Fan, J., and Vercauteren, F. Somewhat practical fully homomorphic encryption. Cryptology ePrint Archive, Report 2012/144, 2012. http: //eprint.iacr.org/2012/144.

[NNOB12] A New Approach to Practical Active-Secure Two-Party Computation. J.B. Nielsen, P.S. Nordholt, C. Orlandi, and S.S. Burra. In CRYPTO 2012, 681-700, Springer LNCS 7417, 2012.

[REN2017] Reniers, V., Rafique, A., Van Landuyt, D., & Joosen, W. (2017). Object-NoSQL Database Mappers: a benchmark study on the performance overhead. Journal of Internet Services and Applications, 8(1), 1.

[RAF2018A] Rafique, A., Van Landuyt, D., & Joosen, W. (2018). Persist: Policy-based data management middleware for multi-tenant saas leveraging federated cloud storage. Journal of Grid Computing, 16(2), 165-194.

[RAF2018B] Rafique, A., Van Landuyt, D., Lagaisse, B., & Joosen, W. (2018). On the performance impact of data access middleware for nosql data stores a study of the trade-off between performance and migration cost. IEEE Transactions on Cloud Computing, 6(3), 843-856.

[SCALE19] SCALE-MAMBA MPC Software. https://homes.esat.kuleuven.be/~nsmart/SCALE/

[SPDZ12] Multiparty Computaton from Somewhat Homomorphic Encryption. I. Damgård, V. Pastro, N.P. Smart and S. Zakarias. In CRYPTO 2012, 643-662, Springer LNCS 7417, 2012.

[SV14] Smart, N. P., and Vercauteren, F. Fully homomorphic SIMD operations. Des. Codes Cryptography 71, 1 (Apr. 2014), 57–81.

[WRK17A] Xiao Wang, Samuel Ranellucci, and Jonathan Katz. Authenticated garbling and communication-efficient, constant-round, secure two-party computation. IACR Cryptology ePrint Archive, 2017:30, 2017.

[WRK17B] Xiao Wang, Samuel Ranellucci, and Jonathan Katz. Global-Scale Secure Multiparty Computation. IACR Cryptology ePrint Archive, 2017:189, 2017.

[YAO86] Protocols for Secure Computations. A. Yao. In 23rd FOCS, 160-164, 1982.

Policy and Regulation

[BREEBAART08] Breebaart, J., Busch, C. , Grave, J. and Kindt, E., "A reference architecture for biometric template protection based on pseudo identities", in Brömme, A. (ed.), BioSig 2008. Proceedings of the Special Interest Group on Biometrics and Electronic Signatures, Bonn, Gesellschaft für Informatik, 2008, 25-37.

[CEPS18A] CEPS, "Software Vulnerability Disclosure in Europe – Technology, Policies and Legal Challenges", Report of a CEPS Task Force (Chair: M. Schaake – Rapporteurs: L. Pupillo, A. Ferreira, G. Varisco), June 2018, https://www.ceps.eu/publications/software-vulnerability-disclosure-europe-technology-policies-and-legal-challenges

[CEPS18B] CEPS, "Strengthening the EU's Cyber Defence Capabilities", Report of a CEPS Task Force (Chair: J. de Hoop Scheffer – Rapporteurs: L. Pupillo, M. Griffith, S. Blockmans, A. Renda), November 2018; https://www.ceps.eu/system/files/CEPS_TFR%20on%20Cyber%20Defence_1.pdf.

[COUDERT19] Fanny Coudert, The purpose specification principle in the Area of Freedom, Security and Justice: towards renewed data protection pinciples for information-based practices in the field of security, PhD thesis KU Leuven, Faculty of Law (ongoing, sup. P. Valcke)

[CUIJPERS&SCHROERS14] Cuijpers, Colette; Schroers, Jessica; 2014. eIDAS as guideline for the development of a pan European eID framework in FutureID. Open Identity Summit 2014; 2014; Vol. 237; pp. 23 - 38 Publisher: Gesellschaft für Informatik; Bonn.

[EIDAS REGULATION] Regulation (EU) No 910/2014 of the European Parliament and of the Council of 23 July 2014 on electronic identification and trust services for electronic transactions in the internal market and repealing Directive 1999/93/EC, OJ L 257, 28.8.2014.

[EMANUILOV18] Ivo Emanuilov; Kim Wuyts; Dimitri Van Landuyt; Natalie Bertels; Fanny Coudert; Peggy Valcke; Wouter Joosen, "Navigating Law and Software Engineering Towards Privacy by Design: Stepping Stones for Bridging the Gap." In: Data Protection and Privacy: The Internet of Bodies, Hart Publishing, 2018, pp. 123 – 140

[ENISA16] ENISA, Good Practice Guide on Vulnerability Disclosure: From challenges to recommendations, January 2016, https://www.enisa.europa.eu/publications/vulnerability-disclosure

[FANTIN18] Stefano Fantin, "Discussing Vulnerabilities in the Policy Arena"; Conference Freedom AND Security – Killing the zero sum process, Europol - The Hague, 22-23 November 2018

[FANTIN19] Stefano Fantin, "Weighting the EU Cybersecurity Act: progress or missed opportunity?", CITIP KU Leuven blog, March 2019

[HERMANS19] J. Hermans, and R. Peeters, "Vingerafdrukken op de Belgische eID - Technische analyze," COSIC internal report, 22 pages, 2019

[IEEE19] IEEE Global Initiative on Ethics of Autonomous and Intelligent Systems, Ethically Aligned Design: A Vision for Prioritizing Human Well-being with Autonomous and Intelligent Systems, First Edition, 2019, https://ethicsinaction.ieee.org/

[KINDT10] Els Kindt, "The use of privacy enhancing technologies for biometric systems analyzed from a legal perspective" in Bezzi, M., Duquenoy, P., Fischer-Hübner, S., Hansen, M. and Zhang, G. (eds.), Privacy and Identity Management for Life, Berlin - New York, Springer, 2010, pp. 134-145.

[KINDT18] Els Kindt, 'Having Yes, Using No ? About the new legal regime for biometric data', in Computer Law and Security Report, 2018, 523-538, available at http://authors.elsevier.com/sd/article/S0267364917303667

[KINDT13] Els Kindt, Privacy and Data Protection Issues of Biometric Applications. A Comparative Legal Analysis, PhD thesis KU Leuven, Faculty of Law (sup. J. Dumortier), published by Springer, 2013, 975 p.

[KINDT19] Els Kindt, 'Chapter 21. A legal perspective on the relevance of biometric presentation attack detection (PAD) for payment services under PSDII and the GDPR' in S. Marcel, M. Nixon, J. Fierrez and N. Evans (eds.), Handbook of Biometric Anti-Spoofing – Presentation Attack Detection, Second edition, Advances in Computer Vision and Pattern Recognition, Springer, 2019, pp. 481-501, also available at https://www.springer.com/la/book/9783319926261

[KOOPS11] Bert-Jaap Koops, "The (In)Flexibility of Techno-Regulation and the Case of Purpose-Binding", Legisprudence, Vol. 5, No. 2, pp. 171-194, 2011. Available at SSRN: https://ssrn.com/abstract=1953967

[KOSTA13] Eleni Kosta, Consent in European Data Protection Law, PhD thesis KU Leuven, Faculty of Law (sup. J. Dumortier and P. Valcke), published by Brill, 2013, 434 p., DOI: https://doi.org/10.1163/9789004232365

[LEENES11] Ronald Leenes, "Framing Techno-Regulation: An Exploration of State and Non-State Regulation by Technology" Legisprudence, Vol. 5, No. 2, pp. 143-169, October 2011; Tilburg Law School Research Paper No. 10/2012. SSRN: https://ssrn.com/abstract=2182439 or http://dx.doi.org/10.2139/ssrn.2182439

[LEMÉTAYER&COUDERT17] Daniel Le Metayer, Bossuet M., Coudert F., Gayrel C., Jaime F., Jouvray C., Kung A., Ma Z., Mana A. "Interdisciplinarity in practice: Challenges and benefits for privacy research", Computer Law & Security Review, (2017) 33 (6), 864-869. doi: 10.1016/j.clsr.2017.05.020.

[NAUTSCH19] A. Nautsch, A. Jiménez, A. Treiber, J. Kolberg, C. Jasserand, E. Kindt, H. Delgadof, M. Todisco, M. Amine Hmani, A. Mtibaa, M. Ahmed Abdelraheem, A. Abad, F. Teixeira, D. Matrouf, M. Gomez-Barrero, D. Petrovska-Delacrétaz, G. Chollet, N. Evans, Th. Schneider, J.F. Bonastre, B. Raj, I. Trancoso, Ch. Busch, 'Preserving Privacy in Speaker and Speech Characterization', Computer Speech and Language (2019) 1 – 38 (in print).

[OTA17] Online Trust Alliance (Internet Society), Internet of Things (IoT) Trust Framework v2.5, 2017, https://otalliance.org/initiatives/internet-things (see also: https://www.internetsociety.org/blog/2019/04/the-internet-of-things-why-trust-by-design-matters/)

[OTHMAN11] Othman, A., Ross, A.; 2011, 'Mixing Fingerprints for Generating Identities', 2011 IEEE International Workshop on Information Forensics and Security, available at https://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=6123152

[PAOLI18] Paoli, L., Visschers, J., Verstraete, C. & van Hellemont E. (2018). *The Impact of Cybercrime on Belgian Businesses*, KU Leuven Centre for IT & IP Law Series, Intersentia

[SCHROERS18] Jessica Schroers; 2018. I have a Facebook account, therefore I am – authentication with social networks.International Review of Law, Computers and Technology; 2018; Vol. 32; iss. 2; pp. 211-223.

[SCHROERS&TSORMPATZOUDI16] Jessica Schroers and Pagona Tsormpatzoudi; 2016. Identity-Theft Through e-Government Services – Government to Pay the Bill? In: Privacy and Identity Management. Time for a Revolution? pp.253-264 Publisher: Springer; Belgium.

[SION19] Laurens Sion; Pierre Dewitte; Dimitri Van Landuyt; Kim Wuyts; Ivo Emanuilov; Peggy Valcke; Wouter Joosen. An Architectural View for Data Protection by Design. 2019 IEEE International Conference on Software Architecture (ICSA); 2019, Publisher: IEEE

[TSORMPATZOUDI15] Tsormpatzoudi, Pagona; Dimitrova, Diana; Schroers, Jessica; Kindt, Els; 2015. Privacy by Design – The Case of Automated Border Control. In: Privacy and Identity Management for the Future Internet in the Age of Globalisation; 2015; Vol. 457; pp. 139 - 152 Publisher: Springer; Dordrecht.

[VAN ALSENOY16] Brendan Van Alsenoy, Regulating data protection: The allocation of responsibility and risk among actors involved in personal data processing, PhD thesis KU Leuven, Faculty of Law (2016, sup. P. Valcke and E. Kindt), published by Intersentia (Data Protection Law in the EU: Roles, Responsibilities and Liability) 2019, 694 p.

[VEALE18] Michael Veale; Reuben Binns; Jef Ausloos, "When Data Protection by Design and Data Subject Rights Clash", International Data Privacy Law, Vol. 8, Issue 2, 2018, p.105–123,

https://doi.org/10.1093/idpl/ipy002. Available at SSRN: https://ssrn.com/abstract=3081069 or http://dx.doi.org/10.2139/ssrn.3081069

[WACHTER18] Sandra Wachter, "Normative Challenges of Identification in the Internet of Things: Privacy, Profiling, Discrimination, and the GDPR." Computer Law & Security Review 34.3 (2018): 436-49.

4. Research Track 3: System and Infrastructure Security

*Contributing Authors: Bart Preneel, Benedikt Gierlichs, Bjorn De Sutter, Cyprien Delpeche, Danny De Cock, Danny Hughes, Dave Singelee, **Frank Piessens**, Ingrid Verbauwhede, Jan Tobias Mühlberg, Lieven Desmet, Pieter Maene, Sam Michiels, Stijn Volckaert, Vincent Naessens, Wouter Joosen*

Scope

Security guarantees often rely upon security foundations delivered by the lower layers of an architecture, and need to be realized across the system stack, as well as across systems and networks. Strategic basic research must therefore invest in these security foundations and in methods and techniques to preserve the security guarantees across the system stack and across network nodes.

The Research Track on system-level security includes the *security of networks* and the *security of systems* (lower levels in the software stack), as well as the growing need to deal with the *operational management of security in infrastructures*.

First, the research theme on system security consists of research activities to secure individual nodes in the infrastructure, and to preserve security guarantees across the system stack from hardware to application. The second research theme focuses on analyzing and securing the communication protocols across systems, both for Internet communication protocols as well as for communication protocols used in the IoT. Finally, the third research theme proposes techniques to monitor the security posture of interconnected systems, and to manage and deploy system security technology.

These three research themes follow in further detail.

4.1. System Security

Current ICT systems are among the most complex systems ever built. One of the key engineering techniques that enables the construction of such complex systems is the use of layered abstractions: the system is built as a stack of layers where each layer hides implementation details of lower layers. The lower layers in this stack are hardware and the higher layers are software, but the boundary between hardware and software has been blurred over time as hardware has become more and more programmable (e.g., microcode, FPGAs).

The importance and success of this engineering technique is undeniable and has enabled the exponential growth of the ICT field for decades.

However, from a security point of view, the use of abstraction layers can introduce significant vulnerabilities and hence risks for the resulting ICT systems. So-called "layer-below" attacks, where an attacker exploits implementation details of lower layers to attack one of the upper layers have been common and have been among the most dangerous attacks over the history of computing. Important examples include:

(1) *memory corruption attacks* that can take over software programmed in C, by relying on implementation details of compiler, operating system and hardware
(2) *side-channel attacks* that can leak secrets from a computation by relying on physical or logical implementation details of the algorithm implementing the computation,
(3) *protocol attacks* that break the secure session abstraction offered by cryptographic protocols by relying on implementation errors of the protocol

One particularly dangerous aspect from the point of view of security is that implementations that refine an abstract specification are known not to preserve security properties of the specification. The wave of software-controlled, micro-architectural side-channel attacks that started in 2018 has been a frightening reminder of the practical relevance of this fact that has been known by the security community for quite a while.

In the context of system security, there is a need to further develop basic security techniques that support finely grained isolation, remote attestation and integrity of software, data and control flow. Yet these techniques must scale well, be able to operate in system-critical contexts and remain robust to new and emerging attacks (such as software-controlled, micro-architectural side-channel attacks). Moreover, security properties leveraging on these security foundations must be preserved across the system stack, so that strong application-level guarantees can be offered.

Finally, since the humans developing the layers make mistakes, as is obvious from the high rate at which software and hardware vulnerabilities and remote exploits thereof surface, there is a need for system layer tools that mitigate the exploitation of vulnerabilities in other layers either by preventing them completely or by raising the bar such that investing in the identification of vulnerabilities does not yield a return on investment for attackers.

A) Industry Needs – Use Cases and Technology Outlook

To deliver secure applications and devices, strong security properties are needed throughout the full software/hardware stack. Important examples of such properties include isolation of security-sensitive data or operations, data and control flow integrity and confidentiality, and the ability to remotely attest the genuine deployment, execution and preservation of security properties.

In order to achieve such strong security properties, appropriate security primitives and abstractions need to be provided at the right level of abstraction, and the security guarantees offered by these primitives and abstractions need to be preserved throughout the full stack, ranging in principle from application-level code and deployment, over system software to hardware.

An emerging class of attacks that require additional attention and mitigation are the software-controlled, micro-architectural side-channel attacks. As recently demonstrated with practical attacks such as Spectre, Meltdown and ForeShadow, existing security guarantees such as the confidentiality of security-sensitive data can be violated if confidentiality properties are not well preserved in lower-layer refinements.

To mitigate remote exploits such as code reuse attacks and non-control data attacks, research is needed to increase the efficiency, the effectiveness, and the technology readiness level of academic solutions such that they become easier and cheaper to use, and compatible with industrial software development lifecycle requirements.

Secure communication requires cryptography protocols; these protocols are evolving from simple entity authentication over authenticated key agreement to offer ever more complex security features and goals, such as the connection of more than two devices or users. In practice one finds two types of protocols: publicly documented and proprietary protocols.

Publicly documented protocols have the advantage of being (formally) verifiable (cf. the research on cryptographic protocols in Technology Building Blocks (Research Track 4)). Having a formal proof for the security of a protocol has obvious advantages but it cannot guarantee the security of said protocol in a practical setting (e.g., the KRACK attack on the formally verified WPA2 protocol, cf. infra). Insecure formally verified protocols can be the result of an inadequate attacker model, a new type of attack or a protocol implementation which does not strictly adhere to the protocol specification.

It has been shown time and time again that most proprietary protocols lack adequate security analysis during their design. Once the protocol specification is revealed (leaked or reverse engineered) security issues are often identified (e.g., insecure protocols in medical devices and vehicles).

B) State-of-the-art: Highlights

In recent years, research in software security has identified and realized a wide range of hardware security features that system software can then use to offer security services. Examples include: virtual memory and privilege levels for process isolation, fine-grained capabilities [WOODRUFF14], trusted computing primitives for secure boot, attestation and sealing of secrets, and hardware-accelerated cryptography or hardware security modules for offering fast and secure cryptographic services [PEREZ06, MCKEEN13, STRACKX10, NOORMAN13, NOORMAN17A]. By now, a range of these security features are readily available [MAENE18], and applications vary from cloud applications [SCHUSTER15] to embedded control systems [VANBULCK17]. Ongoing research aims

towards closing (micro-)architectural side channels and guaranteeing the absence of vulnerabilities in isolated code [AGTEN15].

Recent research on software-controlled micro-architectural attacks, with Rowhammer [SEABORN15], Meltdown [LIPP18], Spectre [KOCHER18], Foreshadow [VANBULCK18A], Nemesis [VANBULCK18B], and TLBleed [GRAS18] being prominent examples, revealed fundamental flaws in commodity hardware. These flaws are generally exploitable through software and allow attacker code to observe and manipulate victim code, effectively bypassing established security mechanisms such as virtual address spaces, process isolation and hierarchical protection rings. The findings range from plain design errors to intricate side-channels. An ongoing body of research explores how to protect programs against such attacks through compiler-based program transformations [COPPENS09], fully abstract compilation [AGTEN12], and approaches to formally reason about information flow in hardware [ZHANG15].

Separation logic can describe ownership transfer, where concurrent program threads move ownership of heap cells into and out of shared resources, such as semaphores or critical regions [OHEARN07, OHEARN19]. Static verification tools such as VeriFast [JACOBS08] rely on separation logic to verify statically that a program contains no data races and no undefined behavior [PHILIPPAERTS14]. Object capabilities are a technique for fine-grained privilege separation in programming languages and systems [DEVRIESE16], with important applications in security. Capabilities have been used to enforce ownership and at-most-once consumption of unique references at run-time and with a flexible notion of borrowing and uniqueness [HALLER10], and recent prototypes show that capabilities can be implemented and evaluated efficiently in hardware [Woodruff14]. Ongoing research identifies approaches to automate program compartmentalization to capability systems [TSAMPAS17], reasoning about evolving invariants on shared data structures [DEVRIESE16], and the relationship between separation logic and capability-based security.

Design-time techniques to detect vulnerabilities include manual, static, and dynamic program analyzes [SONG19]. Run-time security policies to mitigate the exploitation of vulnerabilities rely on hardware-supported isolation (to shield off security-critical components), entropy-based diversification (to reduce the amount of a priori or leaked knowledge useful to an attacker), and redundancy (to enforce correct behavior). A wide range of policies and mechanisms exist [LARSEN18], most of which focus on memory exploits, but all of which leave major vulnerabilities exploitable [VANDERVEEN17], including data-only attacks that are receiving growing interest [ISPOGLOU18].

Multi-Variant eXecution (MVX) is recognized as one of the most effective protection mechanisms [VOLCKAERT16A], as also witnessed by the US$ 65M CFAR research programme of DARPA that started in 2015 for protecting code by means of MVX and that is in its final stages. The core idea of MVX is to execute and monitor multiple variants of the same program in parallel on the same input. The variants are constructed in such a way that no remote exploit can ever simultaneously compromise all variants without causing an observable divergence in their execution behavior. Several major hurdles have been taken to enable MVX for security-critical software, including support for parallel applications [VOLCKAERT17] and mechanisms that can trade-off efficiency for effectiveness [VOLCKAERT16B]. Other issues remain unsolved, however, such as automated support for common software artefacts such as diverging benign behavior and shared memory accesses.

C) Main Areas of Work

This research theme includes four research activities. The first one aims for the preservation of security properties across layers and the protection against software-driven side-channel attacks (RA 3.1.1). The second one aims for the inception and development of processor extensions to support new system security models (RA 3.1.2). (RA 3.1.3) aims for the achievement of security and safety properties in mixed-criticality systems and the last research activity in this theme targets diversity-based multi-variant execution techniques for system defense (RA 3.1.4).

(RA 3.1.1) Protection Against Software-Controlled Side-Channel Attacks (on general purpose hardware)

Without significant extensions to the specifications of the various abstraction layers in current systems, there can be no strong assurance of the confidentiality of data handled by higher layers. It is always possible that a functionally correct implementation of the underlying layers leaks the confidential data of the upper layer.

The underlying problem is that confidentiality properties are so-called 2-safety properties that can only be specified by talking about the relation between two executions of the system, whereas functional correctness is usually a 1-safety property that can be specified as a predicate on single executions.

A challenging and interesting research problem is (1) how to enhance the specifications of the various abstraction layers in ICT systems to also cover the security properties that implementations of these layers must have or preserve, and then (2) to develop implementations of higher layers in terms of lower layers that satisfy these security properties with high assurance.

The research in this activity will focus on how security specifications should be formatted, what their impact is on performance and resource consumption, and how to balance tradeoffs between security and performance concerns in implementations. Moreover, the activity will investigate techniques to check the compliance of an implementation to such security specifications – which is notoriously harder for 2-safety properties. Formal verification of these does not yet scale to practical systems.

The purpose of this research activity is to investigate these topics in a number of particular stacks of abstraction layers. More specifically, we will investigate different instances of such stacks, including at least the following abstraction layers: (1) separation logic specifications, (2) source code in safe and unsafe programming languages, (3) machine code in a variety of instruction set architectures (ISAs) and (4) ISA implementations described in hardware description languages.

(RA 3.1.2) Processor Extension to Support New System Security Models

A first unit of work within the process-extension research activity is the exploration of capability-based processor models. Capability systems provide fine-grained notion of memory protection and ownership, and enable the enforcement of programming languages' memory models and fault isolation in hardware rather than software. A capability-based ISA supplies protection primitives to the compiler, language runtime, and operating system. Recent research and prototypes illustrate the effectiveness and efficiency of the approach.

To enable the adoption of capability systems, and to provide rigorous security guarantees for these systems, we aim to introduce a capability-based security model for the lower layers of the system and hardware stack and develop a formal understanding of software execution under this model. We will investigate the challenges of capability-based/ownership-based model of security for software development and hardware design, and relate that model to more established notions of trusted execution (e.g., protected module architectures). We will further investigate the performance impact of implementations of capability systems and develop approaches to translate notions of capabilities and ownership in higher-level programs and specifications to lower-layer abstractions and the execution model. Future work in this field will extend and accelerate ongoing efforts towards secure compilation and fully abstract compilation of high-level abstractions to capability models and trusted execution environments.

Hardware security extensions, such as the capability extensions, have a negative effect on the performance of the processor. It reduces the throughput of the processor or increases the area as isolation techniques reduce sharing of resources. To support the processor extensions, we aim at integrating results of the hardware building blocks from Research Track 4. More specifically, we aim at integrating low latency crypto modules and protocols into the processor pipeline or on the memory architecture. Moreover, the provided solutions should resist physical side-channel attacks such as power, EM and fault attacks. Side-channel security evaluation and countermeasures will be coordinated with Research Track 4.

(RA 3.1.3) Security and Safety in Mixed Criticality Systems

Mixed-criticality systems combine components with different levels of criticality, a paradigm that becomes more and more prevalent in smart infrastructure. An example of such systems is smart vehicles, that integrate safety-critical functions (steering, braking) with infotainment functions. The software components that implement the different functions have diverging requirement w.r.t. safety, security and performance, and their development is governed by entirely different standards. Yet, the components interact and may even share communication and processing resources, which can lead to unintended and potentially disastrous consequences when components malfunction or become compromised.

A key research objective in this field is to devise technology that can provide rigorous safety, security, and availability guarantees for mixed-criticality software that executes in potentially malicious environments. We will investigate how advanced security architectures in modern processors, e.g., trusted execution with protected module architectures and capability systems, can be used and extended to satisfy these requirements. Specifically, we will work towards providing strong availability guarantees for critical software, in combination with attestable isolation and integrity protection, while relying on a minimal Trusted Computing Base.

Besides hardware support, we will also work on compilation and verification techniques that allow for safety and security requirements to be propagated from high-level languages to the execution platform, where these requirements must be enforced. Such a close integration of software development tools and the execution platform is difficult to achieve but bears the potential of alleviating developers from the burden of target-specific development. It also opens up possibilities for the automated application of diversification approaches such as Multi-Variant Execution.

(RA 3.1.4) Diversity-based Multi-Variant Execution Mitigation Techniques for System Defense

Multi-variant execution (MVX) is ready for deployment on safety/mission-critical software. Broader deployment for security-critical applications (such as web browsers and web servers) by the general public and across industries, is not yet feasible, however. The reasons are that (1) costly manual patching of source code is required to make software MVX-compatible code; (2) current approaches are all-or-nothing, thus often incurring too much overhead; (3) commonly used software artefacts (e.g., shared memory) are not supported; (4) the potential to mitigate non-memory-errors (e.g., integer overflow) with acceptable overhead has not been explored yet.

An interesting research direction to overcome these hurdles is to explore Partial Multi-Variant eXecution (PMVX), in which security-critical parts are isolated from the rest of the software, and only the security-critical partition is executed in multiple variants. This promises to reduce the overhead and to limit the effort needed to make the different software partitions MVX-compatible. Our research will focus on developing the best mechanisms to switch between single and multi-variant mode at run time, and on the co-design of static and dynamic software analysis and rewriting techniques with hardware-enforced protection and isolation to enable effective and efficient PMVX that mitigates memory vulnerabilities (both code pointer and non-control data) as well as non-memory errors.

The most promising designs will be evaluated in proof-of-concept tools on real-world software, for which we will build on the MVX prototype already publicly available from UGent, compiler technology that now scales to large applications, such as LLVM, and increasing hardware support for memory protection and compartmentalization such as Intel MPK.

D) Expected Outcomes and Road Map

For (RA 3.1.1), we will investigate how architecture and micro-architecture optimization techniques such as caching and speculative execution, enable side-channel attacks. We aim to develop modular hardware designs to demonstrate attacks and defenses (Y1). Formal specifications of this lowest layer in the system stack will be used as building blocks to model higher-level interactions and to verify security properties (Y2-3).

The Cybersecurity Strategic Research Programme will accelerate work towards program compartmentalization for capability systems in (RA 3.1.2), ultimately paving the way to a deeper understanding of the relationship between ownership in high-level abstractions and the use of capabilities in lower-level abstractions (Y1).

In (RA 3.1.3), we will initially work on secure interrupt handling and secure scheduling for protected modules with Sancus and Intel SGX (Y1). Based on these mechanisms, we aim to develop a notion of generalized trusted availability, where critical software is guaranteed to adhere to a scheduling policy regardless of the system state and with an untrusted operating system performing resource management and allocation (Y2).

For (RA 3.1.4), we will initially work on a proof of concept implementation of a PMVX engine and compiler techniques to automatically rewrite applications for PMVX compatibility (Y1). The second major objective will be shared memory support in the engine and compiler tools and the use of Intel MPK to isolate partitions, as

well as support for detecting non-memory exploits to maximize the offered protection (Y2). On the longer run (Y3-Y5), we aim for a feature-complete proof-of-concept implementation that supports an efficiency vs. effectiveness trade-off such that it can be deployed in the widest range of scenarios, i.e., with tight or loose performance and resource constraints.

4.2. Network Security

A challenging aspect of securing infrastructure is the securing the communication. In the context of Internet communication, there is a growing need to protect critical Internet components and protocols, such as DNS, BGP, etc. Similarly, there is a need to improve support for service providers and network operators to roll-out secure communication networks for IoT and Industry 4.0.

This research theme focuses on analyzing and securing the communication protocols across systems, both for Internet communication protocols as well as for communication protocols used in the IoT.

A) Industry Needs – Use Cases and Technology Outlook

There is a clear need for secure and reliable communication infrastructure, as this is the backbone for all business-to-consumer, business-to-business, on-premise and cloud communication. This critical communication infrastructure includes routing protocols (such as BGP), DNS (to translate IP addresses to names and vice versa), wireless protocols such as WiFi and 5G, and TLS and HTTPS as secure application channels. Particularly the older but still widely used Internet protocols exhibit inherent security challenges with insufficient security provisions and insecure defaults.

Moreover, the emerging need for secure IoT communication and industrial wireless networks drives further investment in tackling the security challenges and performance trade-offs in networks such as LoRa, NB-IOT, SigFox, 5G, LTE and WiFi.

Finally, as new communication protocols arise, it is important to quickly assess and improve the security of these protocols in industry-critical contexts. In particular, the implementation analysis of emerging communication protocols will enable securing new industrial deployments of these protocols from the early stages of adoption.

B) State-of-the-art: Highlights

Over the past few years, several attacks have been incepted on critical communication infrastructure and security protocols. Vanhoef et al. improved earlier work [AlFardan13] to obtain a practical method to break RC4 keystreams in WPA-TKIP and TLS [Vanhoef14]. More recently, they discovered serious weaknesses in the WPA2 protocol by using key reinstallation attacks (dubbed KRACK attacks) [Vanhoef17]. Shortly after, Poddebniak et al. showed the practical feasibility of breaking S/MIME and OpenPGP email encryption [Poddebniak18]. A set of origin-leaking vectors was revealed to bypass Cloud-based Security Providers (such as CloudFlare, ProLexic and Incapsula), and CloudPiercer was able to bypass more than 7 out of 10 cloud-protected websites at the time [Vissers15]. While SSL/TLS has been the most successful security protocol, many flaws have been identified in the specification [Mavrogiannopoulos12, Meyer13] and in implementations [Albrecht16]. After many years of research, TLS 1.3 has been published in mid 2018 [Rescorla18].

At the same time, several proposals have been incepted to enhance the security and privacy of critical infrastructure, such as DNS and BGP. DNSSEC was published in 2005 [Arends05] yet its deployment has been slow [Yang10]. RFC 7816 proposed DNS query name minimization to improve DNS privacy, so that the full query name is no longer sent to upstream name servers [Bortzmeyer16], and the use of 0x20 DNS query encoding strongly reduces the impact of DNS cache poisoning attacks [Dagon08]. Recently, CloudFlare and Google started to adopt the DNS over TLS protocol [Dickinson18] to enhance the privacy of queries to their public resolvers.

The vulnerability of BGP to hijacking is well understood [Vervier15]; after decades of research on solutions (e.g., [Kent00]), an RFC was published [Lepinski17]. However, as with DNSSEC, there are many deployment challenges [Gilad18]. One of the main challenges remains a robust and scalable PKI; a promising approach is ARPKI [Basin18].

In recent years, numerous works on the security of IoT devices and protocols have been published, each of these discovering novel security vulnerabilities. This leads to the scenario where millions of IoT devices are deployed yet consumers may know very little about the capabilities and security of these devices [LUND14]. Most of the discovered security issues are caused by not having security implemented in the IoT system, by using weak cryptography and by making incorrect security assumptions. Sometimes the result can be quite devastating. Ronen et al. found several security vulnerabilities in Philips Hue smart lamps and their implementation of the ZigBee Light Link protocol and showed how these could be exploited to spread a worm wirelessly over a large area [RONEN18]. Tellez et al. focused on WSN (wireless sensor networks) and elements of their security, and found a security problem in a popular password-based bootloader protocol [TELLEZ16]. Also, automotive systems could be vulnerable to IoT hacks. Wouters et al. reverse-engineered the Passive Keyless Entry and Start (PKES) system of Tesla and discovered that it relies on the outdated proprietary DST40 cipher [WOUTERS19]. They have shown that a genuine key fob can be cloned in seconds by performing a brute-force attack. Verdult et al. executed multiple cryptanalytic attacks on the Hitag2 stream cipher used in NXP transponders, widely deployed in key fobs [VERDULT12]. Sometimes IoT security vulnerabilities could even lead to physical harm of the end-user. Marin et al. have shown that Implantable Cardiac Devices (ICDs) do not have proper security protocols implemented and are vulnerable to various wireless attacks [MARIN16]. These security vulnerabilities clearly demonstrate the need for security-by-design and privacy-by-design and the development of lightweight cryptographic algorithms and protocols for IoT.

C) Main Areas of Work

The execution of this research comprises a security study of critical Internet component (RA 3.2.1), the assessment and improvement of secure communication protocols for IoT and Industry 4.0 (RA 3.2.2), and inception of tools and methods for advanced protocol analysis (RA 3.2.3).

(RA 3.2.1) Study of Critical Internet Components and Protocols

The foundations of the Internet infrastructure as we know it today date from the past four decades and have evolved incrementally over time. Essential building blocks such as routing protocols, DNS name translation, the HTTP protocol and public key infrastructures (PKIs), as well as more recent communication channels such as WiFi, 5G and TLS, are used in security-critical and safety-critical transactions on a continuous basis, but often lack the necessary security characteristics or are configures insecurely.

In this research activity, these core building blocks will be assessed from a security point of view. Our focus hereby are the availability of the services, as well as the confidentiality and integrity characteristics. New attack vectors will be identified, and potential mitigation techniques will be proposed and validated.

(RA 3.2.2) Secure Communication Protocols for the IoT

Multiple communication standards are being developed specifically for IoT and Industry 4.0. These protocols are composed of one or more negotiation phases that typically include authenticated key agreement and data transfer phases. They are often optimized for an IoT or industrial setting, for example having a low data rate and long communication range while still guaranteeing low energy consumption. Typical examples include LoRa, SigFox, Dash7, DART, Nwave, 5G, LTE, NB-IOT, … and industrial networks such as SCADA systems and CANBUS. In this research activity, we will assess the security of these new and emerging communication protocols and improve them, taking into account their specific characteristics and constraints (such as limited payload size, maximum bitrate, etc.).

The research in this activity will focus on how to provide lightweight, end-to-end security, and how to deal with low-latency constraints, which are typical for networks such as CANBUS. Moreover, the activity will investigate how to offer data authentication with payload lengths of only a few bits (e.g., in SigFox), and study energy trade-offs between communication and computation costs when securing IoT communication protocols.

As part of the security assessment of particular communication protocols, the activity will study to what extent the security specifications of the communication standard are clear, unambiguous and complete, and what can be improved. In addition, the activity will assess how to use the communication protocol, or a combination of protocols, securely in a given application (e.g., V2X, smart grids, etc.).

Another research topic that will be covered in this research activity, is key management. We will particularly focus on highly distributed (IoT) networks and networks where the topology dynamically changes over time (e.g., V2X, industrial control systems, etc.).

(RA 3.2.3) Analysis of Protocol Implementations

Our research will focus on the study of the implementation security of both publicly documented and proprietary protocols. This work may require reverse engineering parts of the protocol, the implementation and the physical communication medium. Both reverse engineering and verifying protocols is often a manual and time-consuming task. Therefore, tools will be developed to automate parts of these tasks.

The following list of tools would be useful in a practical reverse engineering and security evaluation setting.

- Automatically recovering physical layer parameters for RF communications
- Automatically detecting recurring fields in packets (e.g., counters, checksums and high entropy data segments)
- Generating test inputs based on protocol specifications
- Fuzzing protocol implementations and state machines by providing random or specially formed inputs
- A reactive jammer, applicable in multiple scenarios

It goes without saying that responsible disclosure procedures will be followed if new flaws are discovered.

D) Expected Outcomes and Road Map

For (RA 3.2.1), the first expected outcome (Y1) is a study of DNS and PKI improvement proposals. The second expected outcome (Y2) incorporates a study on improvements for secure routing.

For (RA 3.2.2), the first expected outcome (Y1) is a detailed analysis of relevant IoT protocols; the second expected outcome (Y2) is a set of improvements for concrete constraints (including low energy and low latency).

Finally, for (RA 3.2.3), the first expected outcome (Y1) is a set of tools to support reverse engineering and security evaluation of IoT protocols; the second outcome will be a report on the application of these to several widely used implementations.

4.3. Security Monitoring and Management

The dynamic evolution of systems and adversaries (attackers) demands for capabilities to continuously monitor and dynamically adjust configurations, versions and deployment settings to deal with the overall rapid evolution on the space. This third research theme proposes techniques to monitor security posture of interconnected systems, and to manage and deploy system security technology.

A) Industry Needs – Use Cases and Technology Outlook

From an operational perspective, systems require a 24/7 operational security model. In the current state of practice, this might include monitoring for intrusion and anomaly detection, and (semi-)automated incident and response handling. An important aspect of these defense techniques, is a continuous intelligence feed: known vulnerabilities and zero-days, known bad actors (e.g., botnets, malicious IPs, domain names used for phishing and spam), existing and novel attack scenario's, … From this angle, there is a continuous need toward gathering and absorbing more actionable security intelligence. Moreover, mature results of such ecosystem observations and measurements studies can be absorbed in self-assessment tools for security maturity, as is for instance already the case for Qualys SSL Labs, the Mozilla Observatory or the HttpHeaders.io assessment website.

At the same time, operational management becomes more and more complex, and a central, classic operational security strategy runs into its limits. In large IoT deployments for instance, the communication overhead to centralize security events and the decision-making urges us to investigate decentral detection and response solutions.

Finally, the growing complexity in security technologies also impacts the operational management. As such, several challenges need to be tackled in secure deployment and updating strategies, to fully embrace the

latest state-of-the-art defense techniques. For example, integrating software diversification as part of the software deployment incurs multiple challenges, such as operational feedback (e.g., bug reporting) as well as in updating deployed software instances.

B) State-of-the-art: Highlights

To quantify the state of security, researchers have performed a broad scale of (longitudinal) ecosystem studies in the last decade to measure the vulnerability landscape and observe new trends in attacks. For example, in the context of web application security, some measured the widespreadness of including (untrusted) third-party JavaScript [NIKIFORAKIS12], some investigated the prevalence of DOM-based XSS on the most frequently used websites[LEKIES13], and others studied the effectiveness of third-party cookie policies in browsers [FRANKEN18]. A longitudinal analysis of the use of web security mechanisms has been performed on the European Web [CHEN16] and the malicious ecosystem of free live streaming services was analyzed [RAFIQUE16]. Similarly, research methods have been established for internet-wide scanning of IPv4 [DURUMERIC13], for enumerating active IPv6 hosts [BORGOLTE18], as well as for assessing the abuse of domain names [HAO13, VISSERS17].

Due to their heterogeneity, the use of a multitude of proprietary protocols and components, and the common involvement of stringent real-time and safety requirements, embedded control systems (CPS, IoT, ICS) are an inherently difficult domain of application for trust assessment, intrusion detection and intrusion response. Still, many of these systems are interconnected with internet technology and cloud integration [NICHOLSON12], which necessitates the integration of network-wide security mechanisms [ZARPELAO17]. These mechanisms either aim to validate communication behavior according to a set of rules [MITCHELL15], or aim to detect anomalies automatically [WRESSNEGGER18]. Techniques to detect [MUEHLBERG15] and prevent [NOORMAN17B] software attacks in a decentralized manner have been presented and typically rely on light-weight isolation [DANIELS17] and trusted computing [NOORMAN17A] primitives. Approaches towards comprehensive but decentralized intrusion detection and trust assessment in embedded control systems are an ongoing area of research.

Software diversification, both temporal and spatial, has matured over the last decade [LARSEN14]. By diversifying the representation of installed and even running software frequently even when there is no functional need to, many attack vectors become ineffective, such as the development of exploits based on patches [COPPENS13] and the exploitation of information leaks [WILLIAMSKING16]. A major issue remains how to integrate software diversification into industrial software development life cycles (SDLCs), i.e., how to make diversification compatible with industrial SDLC requirements. For some aspects, progress has been made, such as bug reporting [ABRATH18]. Other aspects still need to be tackled, however.

To protect against so-called man-at-the-end attacks (including reverse engineering and tampering), a wide range of pure software-based protection techniques has been developed that complement remote attestation techniques. These range from obfuscations and watermarking techniques [COLLBERG09] to strong anti-debugging techniques [ABRATH16]. As all of these techniques come with overhead and each one only fends off certain attacks [SCHRITTWIESER16] they have to be combined. Decision support to select and combine protections is a long-standing challenge, however. Only recently, the first hurdles towards building the necessary attacks models and knowledge bases have been taken [CECCATO19, BASILE19].

C) Main Areas of Work

The execution of this research theme comprises in intelligence gathering (RA 3.3.1), methods and tools for secure deployment (RA 3.3.2), and decentralized detection and response (RA 3.3.3).

(RA 3.3.1) Intelligence Gathering and Identification of Security State

In this research activity, we will study and observe trends in the attacker landscape, as well as measure the state of practice in defense and mitigation techniques. These continuous observations provide a complementary intelligence gathering, and fuel the identification of the security state of a (set of) systems.

In ecosystem studies, we will assess the widespreadness of existing and novel security weaknesses, as well as the attacker landscape. To this extent, security measurement methods will be developed and these security metrics will be used to Track longitudinal trends. The use of security metrics and a measured baseline will help to develop maturity trajectories for different industry verticals to gradually improve their best practices

compared to the current initiatives of peers and the envisioned security end goals. The initial focus of the ecosystem studies will be on (client-side) web security and domain name security.

Similar intelligence gathering will be performed in industrial control systems, in which we will investigate how intrusion detection techniques can be set up in such industrial networks, and how automated responses can be rolled out.

(RA 3.3.2) Methods and Tools for Secure Deployment

In this activity, we will evaluate the effectiveness of software-based protections against man-at-the-end attacks, with respect to the varying deployment contexts and established as well as emerging attack techniques. We will research models and evaluation techniques to measure and predict the potency, resilience, and stealth that various protections offer vis-à-vis different attacks. The goal is not only to study formal models, but also to focus on practical methods to aid users of such protections in the form of decision support systems. Empirical research as well as the use of many online information sources (e.g., hacker blogs) and academic papers presenting novel offensive and defensive methods will be the main starting points.

Moreover, we will research techniques to integrate diversity-based mitigation techniques into industrial software development life cycles. We will first assess the use of basic techniques in the Flemish industry and the (perceived) constraints that block the take-up of more advanced techniques. We will then study options to improve the practical viability of existing techniques by making diversification compatible with the identified constraints, e.g., by developing novel schemes for distributing diversified application versions. Other options include novel schemes for diversifying code at install time, at boot time, at load time, or at run time, with white-box or black-box access to the tools in the developer's toolbox; novel schemes to collect and interpret crash reports, etc. Finally, we will develop proof-of-concept tool implementations to evaluate the automated deployment of the most interesting schemes.

(RA 3.3.3) Detection and Response for IoT and Industrial Control Systems

A growing number of systems is centralizing security information in a Security Information and Event Management (SIEM) system to support the detection of and response to security incidents. While this paradigm is highly effective today, it creates a single point of failure and comes with its own security and privacy risks. Moreover, for large scale systems such as the IoT with billions of devices the communication overhead of such an approach may be too large. Moreover, the inherent latency in this approach is not acceptable for autonomous systems (e.g., transportation) that require an instantaneous response. Research is needed to study to which extent security-related information can be managed at a local level and compressed in an optimal way such that only relevant information and reporting are centralized. Moreover, information can be aggregated at a higher level using secure computation techniques (c.f. secure multi-party computation, MPC) which offers a combination of robustness and privacy. The research will explore trade-offs between satisfying safety requirements, security, privacy and efficiency; and validate these on realistic cases.

In addition, we will explore mixed-criticality scenarios (e.g., autonomous systems with remote connectivity), where security functions are sharing the same execution platform with functions of different criticality. In these domains, we seek to integrate software self-protection, local and remote trust assessment, and incident detection and response.

D) Expected Outcomes and Road Map

For (RA 3.3.1), the first expected outcome (Y1) is to perform at least one new ecosystem studies in the context of web security or domain name security, as well as a baseline measurement on client-side web security techniques. In Y2, two new ecosystem measurements will be set up, including at least one longitudinal study.

For (RA 3.3.2), the first objective is to define a concrete model and evaluation methodology, including metrics, for the most common step of reverse engineering of protected applications, being the search for the relevant assets and code fragments under attack (Y1). The second objective is to validate that methodology (Y2). The longer-term objective (Y3-Y5) is to develop similar methodologies to model the relation between other protections and attacks, with a focus on tampering, thus enabling complete decision support.

Regarding diversity, the first objective is to obtain an overview of concrete technical hurdles in the Flemish industry regarding the uptake of diversification-based protection (Y1). Based on the inputs we will collect in

industry, we will then define a plan to develop the necessary adaptations and extensions to state-of-the-art techniques to make them compatible with the industry's software development lifecycle requirements.

For (RA 3.3.3), the first expected outcome (Y1) is a study of the trade-offs between localized and centralized detection and response systems. The second expected outcome includes secure and privacy-friendly aggregation of localized detection and response systems (Y2).

4.4. Connections with other Research Tracks

Research Theme 1 on system security investigates the preservation of security guarantees across the system stack. These activities build upon foundations of secure hardware and secure implementations of cryptographic algorithms and protocols, as researched in Research Track 4 (Technology Building Blocks).

Moreover, the security guarantees provided by the security extensions of processors enable the research activities on secure compilation in Research Track 1 (Application and Software Security) to preserve guarantees across the stack. It is expected that the activities of Research Track 1 and Theme 1 (System Security) will positively impact each other, strengthen each other's results, and create opportunities for collaboration.

The research activities on architectural side channels for general-purpose hardware in theme 1 (System security) have similarities to the hardware-based side channels (Theme 4 of this Track 4). While they share some techniques to obtain information, they operate at different abstraction levels; hence the required countermeasures are very different.

The secure communication activities in an IoT or Industry 4.0 context (Theme 2) require efficient cryptographic algorithms, as studied in Theme 2 of Track 4, to achieve high bandwidth or low latency. Moreover, the two tracks also complement each other on protocol evaluation: this Track evaluates protocol implementations in Theme 2, whereas Theme 3 of Track 4 focuses on the analysis of cryptographic protocols.

4.5. References

[ABRATH16] Bert Abrath, Bart Coppens, Stijn Volckaert, Joris Wijnant, Bjorn De Sutter. Tightly-coupled self-debugging software protection. Proceedings of the 6th Workshop on Software Security, Protection, and Reverse Engineering. Art. 07, pp. 7:1-7:10, December 2016

[ABRATH18] Bert Abrath, Bart Coppens, Mohit Mishra, Jens Van den Broeck, Bjorn De Sutter. ΔBreakpad: Diversified Binary Crash Reporting. IEEE Transactions on Dependable and Secure Computing, 2018, DOI: 10.1109/TDSC.2018.2823751

[AGTEN12] Agten, P., Strackx, R., Jacobs, B. and Piessens, F., 2012, June. Secure compilation to modern processors. In 2012 IEEE 25th Computer Security Foundations Symposium (pp. 171-185). IEEE.

[AGTEN15] Agten, P., Jacobs, B. and Piessens, F., 2015, January. Sound modular verification of C code executing in an unverified context. In ACM SIGPLAN Notices (Vol. 50, No. 1, pp. 581-594). ACM.

[ALBRECHT16] Martin R. Albrecht, Kenneth G. Paterson, Lucky Microseconds: A Timing Attack on Amazon's s2n Implementation of TLS, EUROCRYPT (1) 2016: 622-643

[ALFARDAN13] Nadhem J. AlFardan, Daniel J. Bernstein, Kenneth G. Paterson, Bertram Poettering, Jacob C. N. Schuldt, On the Security of RC4 in TLS, USENIX Security Symposium 2013: 305-320

[ARENDS05] Roy Arends, Rob Austein, Matt Larson, Dan Massey, Scott Rose, DNS Security Introduction and Requirements. RFC 4033 (Proposed Standard), March 2005. Updated by RFC 6014.

[BASILE19] Cataldo Basile, Daniele Canavese, Leonardo Regano, Paolo Falcarin, Bjorn De Sutter. A meta-model for software protections and reverse engineering attacks. Journal of Systems and Software, Volume 150, 2019

[BASIN18] David A. Basin, Cas Cremers, Tiffany Hyun-Jin Kim, Adrian Perrig, Ralf Sasse, Pawel Szalachowski, Design, Analysis, and Implementation of ARPKI: An Attack-Resilient Public-Key Infrastructure, IEEE Trans. Dependable Sec. Comput. 15(3): 393-408 (2018)

[BORGOLTE18] Kevin Borgolte, Shuang Hao, Tobias Fiebig, Giovanni Vigna, Enumerating active IPv6 hosts for large-scale security scans via DNSSEC-signed reverse zones, IEEE Symposium on Security and Privacy (SP'18), pages 770-784, May 20, 2018

[BORTZMEYER16] Stephane Bortzmeyer, DNS Query Name Minimisation to Improve Privacy, IETF RFC 7816, March 2016

[CECCATO19] Mariano Ceccato, Paolo Tonella, Cataldo Basile, Paolo Falcarin, Marco Torchiano, Bart Coppens, and Bjorn De Sutter. Understanding the Behaviour of Hackers while Performing Attack Tasks in a Professional Setting and in a Public Challeng. Empirical Software Engineering, vol. 24, no. 1, pp. 240-286, Feb 2019.

[CHEN18] Ping Chen, Christophe Huygens, Lieven Desmet, Wouter Joosen, Longitudinal study of the use of client-side security mechanisms on the European web, Workshop on Empirical Research Methods in Information Security, Proceedings of the 25th International Conference Companion on World Wide Web, pages 457-462, Montreal, Canada, April 11-15, 2016

[COLLBERG09] C. Collberg and J. Nagra. Surreptitious Software: Obfuscation, Watermarking, and Tamperproofing for Software Protection. Pearson Education, 2009

[COPPENS09] Coppens, B., Verbauwhede, I., De Bosschere, K. and De Sutter, B., 2009, May. Practical mitigations for timing-based side-channel attacks on modern x86 processors. In 2009 30th IEEE Symposium on Security and Privacy (pp. 45-60). IEEE.

[COPPENS13] Bart Coppens, Bjorn De Sutter, Koen De Bosschere. Protecting your software updates. IEEE Security & Privacy. Vol. 11 No. 2, pages 47-54, March-April 2013

[DAGON08] David Dagon, Manos Antonakakis, Paul Vixie, Tatuya Jinmei, Wenke Lee, Increased DNS forgery resistance through 0x20-bit encoding: security via leet queries. In Proceedings of the 15th ACM conference on Computer and communications security (CCS '08), ACM, New York, NY, USA, 211-222.

[DANIELS17] Daniels, W., Hughes, D., Ammar, M., Crispo, B., Matthys, N. and Joosen, W., 2017, December. S µ V-the security microvisor: a virtualisation-based security middleware for the internet of things. In Proceedings of the 18th ACM/IFIP/USENIX Middleware Conference: Industrial Track (pp. 36-42). ACM.

[DEVRIESE16] Devriese, D., Birkedal, L. and Piessens, F., 2016, March. Reasoning about object capabilities with logical relations and effect parametricity. In 2016 IEEE European Symposium on Security and Privacy (EuroS&P) (pp. 147-162). IEEE.

[DICKINSON18] Sara Dickinson, Daniel Kahn Gillmor, Tirumaleswar Reddy, Usage Profiles for DNS over TLS and DNS over DTLS, IETF RFC 8310, March 2018.

[DURUMERIC13] Zakir Durumeric, Eric Wustrow, and J. Alex Halderman, ZMap: Fast Internet-wide scanning and its security applications, 22nd USENIX Security Symposium (USENIX Security '13), pp. 605-620. 2013

[FRANKEN18] Gertjan Franken, Tom van Goethem, Wouter Joosen, Who Left Open the Cookie Jar? A Comprehensive Evaluation of Third-Party Cookie Policies, 27th USENIX Security Symposium , USENIX Security Symposium, pages 151-168, BALTIMORE, MD, USA, 2018

[GILAD18] Yossi Gilad, Tomas Hlavacek, Amir Herzberg, Michael Schapira, Haya Shulman, Perfect is the Enemy of Good: Setting Realistic Goals for BGP Security, HotNets 2018: 57-63

[GRAS18] Gras, B., Razavi, K., Bos, H. and Giuffrida, C., 2018. Translation Leak-aside Buffer: Defeating Cache Side-channel Protections with TLB Attacks. In 27th USENIX Security Symposium (USENIX Security 18) (pp. 955-972).

[HALLER10] Haller, P. and Odersky, M., 2010, June. Capabilities for uniqueness and borrowing. In European Conference on Object-Oriented Programming (pp. 354-378). Springer, Berlin, Heidelberg.

[HAO13] Shuang Hao, Matthew Thomas, Vern Paxson, Nick Feamster, Christian Kreibich, Chris Grier, Scott Hollenbeck, Understanding the domain registration behavior of spammers, the 2013 conference on Internet measurement conference (IMC 2013), pages 63-76, Oct 23, 2013

[ISPOGLOU18] Kyriakos K. Ispoglou, Bader AlBassam, Trent Jaeger, and Mathias Payer. 2018. Block Oriented Programming: Automating Data-Only Attacks. In Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security (CCS '18). ACM, New York, NY, USA, 1868-1882. DOI: https://doi.org/10.1145/3243734.3243739

[JACOBS08] Jacobs, B. and Piessens, F., 2008. The VeriFast program verifier (Vol. 7, pp. 3-2). Technical Report CW-520, Department of Computer Science, Katholieke Universiteit Leuven.

[JUGLARET16] Juglaret, Y., Hritcu, C., De Amorim, A.A., Eng, B. and Pierce, B.C., 2016, June. Beyond good and evil: Formalizing the security guarantees of compartmentalizing compilation. In 2016 IEEE 29th Computer Security Foundations Symposium (CSF) (pp. 45-60). IEEE.

[KENT00] Stephen T. Kent, Charles Lynn, Karen Seo,Secure Border Gateway Protocol (S-BGP), IEEE Journal on Selected Areas in Communications 18(4): 582-592 (2000).

[KOCHER18] Kocher, P., Genkin, D., Gruss, D., Haas, W., Hamburg, M., Lipp, M., Mangard, S., Prescher, T., Schwarz, M. and Yarom, Y., 2018. Spectre attacks: Exploiting speculative execution. arXiv preprint arXiv:1801.01203.

[LARSEN14] Per Larsen, Andrei Homescu, Stefan Brunthaler, and Michael Franz. 2014. SoK: Automated Software Diversity. In Proceedings of the 2014 IEEE Symposium on Security and Privacy (SP '14). IEEE Computer Society, Washington, DC, USA, 276-291. DOI: https://doi.org/10.1109/SP.2014.25

[LARSEN18] Per Larsen and Ahmad-Reza Sadeghi (Eds.). 2018. The Continuing Arms Race: Code-Reuse Attacks and Defenses. Association for Computing Machinery and Morgan & Claypool, New York, NY, USA.

[LEKIES13] Sebastian Lekies, Ben Stock, Martin Johns, 25 Million Flows Later - Large-scale Detection of DOM-based XSS, in 20th ACM Conference on Computer and Communications Security (ACM CCS'13), November 2013

[LEPINSKI17] Matthew Lepinski, Kotikalapudi Sriram, BGPsec Protocol Specification, IETF RFC 8205, September 2017.

[LIPP18] Lipp, M., Schwarz, M., Gruss, D., Prescher, T., Haas, W., Fogh, A., Horn, J., Mangard, S., Kocher, P., Genkin, D. and Yarom, Y., 2018. Meltdown: Reading kernel memory from user space. In 27th USENIX Security Symposium (USENIX Security 18) (pp. 973-990).

[LUND14] D. Lund, C. MacGillivray, V. Turner, and M. Morales, "Worldwide andregional internet of things (iot) 2014–2020 forecast: A virtuous circleof proven value and demand,"International Data Corporation (IDC),Tech. Rep, 2014.

[MAENE18] Maene, P., Goetzfried, J., De Clercq, R., Mueller, T., Freiling, F. and Verbauwhede, I., 2018. Hardware-based trusted computing architectures for isolation and attestation. IEEE Transactions on Computers, 67(3), pp.361-374.

[MARIN16] E. Marin, D. Singelée, F. D. Garcia, T. Chothia, R. Willems, and B. Preneel, "On the (in)security of the Latest Generation Implantable Cardiac Defibrillators and How to Secure Them," In Annual Computer Security Applications Conference (ACSAC), ACM, pp. 226-236, 2016.

[MAVROGIANNOPOULOS12] Nikos Mavrogiannopoulos, Frederik Vercauteren, Vesselin Velichkov, Bart Preneel, A cross-protocol attack on the TLS protocol, ACM Conference on Computer and Communications Security 2012: 62-72

[MCKEEN13] McKeen, F., Alexandrovich, I., Berenzon, A., Rozas, C.V., Shafi, H., Shanbhogue, V. and Savagaonkar, U.R., 2013. Innovative instructions and software model for isolated execution. Hasp@ isca, 10(1).

[MEYER13] Christopher Meyer, Jörg Schwenk, SoK: Lessons Learned from SSL/TLS Attacks, WISA 2013: 189-209

[MITCHELL15] Mitchell, R. and Chen, R., 2015. Behavior rule specification-based intrusion detection for safety critical medical cyber physical systems. IEEE Transactions on Dependable and Secure Computing, 12(1), pp.16-30.

[MUEHLBERG15] Muehlberg, J.T., Noorman, J. and Piessens, F., 2015, September. Lightweight and flexible trust assessment modules for the Internet of Things. In European Symposium on Research in Computer Security (pp. 503-520). Springer, Cham.

[NICHOLSON12] Nicholson, A., Webber, S., Dyer, S., Patel, T. and Janicke, H., 2012. SCADA security in the light of Cyber-Warfare. Computers & Security, 31(4), pp.418-436.

[NIKIFORAKIS12] Nick Nikiforakis, Luca Invernizzi, Alexandros Kapravelos, Steven Van Acker, Wouter Joosen, Christopher Kruegel, Frank Piessens, Giovanni Vigna, You are what you include: Large-scale evaluation of remote JavaScript inclusions, ACM Conference on Computer and Communications Security (CCS 2012), pages 736-747, Raleigh, NC, USA, October 16-18, 2012

[NOORMAN13] Noorman, J., Agten, P., Daniels, W., Strackx, R., Van Herrewege, A., Huygens, C., Preneel, B., Verbauwhede, I. and Piessens, F., 2013. Sancus: Low-cost trustworthy extensible networked devices with a zero-software trusted computing base. In Presented as part of the 22nd USENIX Security Symposium (USENIX Security 13) (pp. 479-498).

[NOORMAN17A] Noorman, J., Bulck, J.V., Muehlberg, J.T., Piessens, F., Maene, P., Preneel, B., Verbauwhede, I., Goetzfried, J., Mueller, T. and Freiling, F., 2017. Sancus 2.0: A low-cost security architecture for IoT devices. ACM Transactions on Privacy and Security (TOPS), 20(3), p.7.

[NOORMAN17B] Noorman, J., Muehlberg, J.T. and Piessens, F., 2017, September. Authentic execution of distributed event-driven applications with a small TCB. In International Workshop on Security and Trust Management (pp. 55-71). SPRINGER, CHAM.

[OHEARN07] O'Hearn, P.W., 2007. Resources, concurrency, and local reasoning. Theoretical computer science, 375(1-3), pp.271-307.

[OHEARN19] Peter O'Hearn, Separation Logic, Communications of the ACM, February 2019, Vol. 62 No. 2, Pages 86-95

[PEREZ06] Perez, R., Sailer, R. and van Doorn, L., 2006, July. vTPM: virtualizing the trusted platform module. In Proc. 15th Conf. on USENIX Security Symposium (pp. 305-320).

[PHILIPPAERTS14] Philippaerts, P., Muehlberg, J.T., Penninckx, W., Smans, J., Jacobs, B. and Piessens, F., 2014. Software verification with VeriFast: Industrial case studies. Science of Computer Programming, 82, pp.77-97.

[PODDEBNIAK18] Damian Poddebniak, Christian Dresen, Jens Müller, Fabian Ising, Sebastian Schinzel, Simon Friedberger, Jurak Somorovsky, Jörg Schwenk, Efail: Breaking S/MIME and OpenPGP Email Encryption using Exfiltration Channels, Proceedings of the 27th USENIX Conference on Security Symposium (USENIX Security 2018), August 15-17, 2018, Baltimore, MD, USA

[RAFIQUE16] M Zubair Rafique, Tom Van Goethem, Wouter Joosen, Christophe Huygens, Nick Nikiforakis, It's free for a reason: Exploring the ecosystem of free live streaming services, Proceedings of the 23rd Network and Distributed System Security Symposium (NDSS 2016) , Proceedings of the 23rd Network and Distributed System Security Symposium (NDSS 2016) , pages 1-15, San Diego, USA, February 21-24, 2016

[RESCORLA18] Erick Rescorla, The Transport Layer Security (TLS) Protocol Version 1.3. RFC 8446 (Proposed Standard), August 2018.

[RONEN18] Eyal Ronen, Adi Shamir, Achi-Or Weingarten, Colin O'Flynn: IoT Goes Nuclear: Creating a Zigbee Chain Reaction. IEEE Security & Privacy 16(1): 54-62 (2018)

[SCHRITTWIESER16] Sebastian Schrittwieser, Stefan Katzenbeisser, Johannes Kinder, Georg Merzdovnik, and Edgar Weippl. 2016. Protecting Software through Obfuscation: Can It Keep Pace with Progress in Code Analysis? ACM Comput. Surv. 49, 1, Article 4 (April 2016), 37 pages

[SCHUSTER15] Schuster, F., Costa, M., Fournet, C., Gkantsidis, C., Peinado, M., Mainar-Ruiz, G. and Russinovich, M., 2015, May. VC3: Trustworthy data analytics in the cloud using SGX. In 2015 IEEE Symposium on Security and Privacy (pp. 38-54). IEEE.

[SEABORN15] Seaborn, M. and Dullien, T., 2015. Exploiting the DRAM Rowhammer bug to gain kernel privileges. Black Hat, 15.

[SONG19] Dokyung Song, Julian Lettner, Prabhu Rajasekaran, Yeoul Na, Stijn Volckaert, Per Larsen, and Michael Franz. SoK: Sanitizing for Security. In IEEE Symposium on Security and Privacy (S&P'19). Accepted. To Appear.

[STRACKX10] Strackx, R., Piessens, F. and Preneel, B., 2010, September. Efficient isolation of trusted subsystems in embedded systems. In International Conference on Security and Privacy in Communication Systems (pp. 344-361). Springer, Berlin, Heidelberg.

[TELLEZ16] M. Tellez, S. El-Tawab, and H. M. Heydari, “Improving the security of wireless sensor networks in an IoT environmental monitoring system,”in Systems and Information Engineering Design Symposium (SIEDS),2016 IEEE, pp. 72–77, IEEE, 2016

[TSAMPAS17] Tsampas, S., El-Korashy, A., Patrignani, M., Devriese, D., Garg, D. and Piessens, F., 2017. Towards automatic compartmentalization of C programs on capability machines. In Workshop on Foundations of Computer Security (pp. 1-14).

[VANBULCK17] Van Bulck, J., Muehlberg, J.T. and Piessens, F., 2017, December. VulCAN: Efficient component authentication and software isolation for automotive control networks. In Proceedings of the 33rd Annual Computer Security Applications Conference (pp. 225-237). ACM.

[VANBULCK18A] Van Bulck, J., Minkin, M., Weisse, O., Genkin, D., Kasikci, B., Piessens, F., Silberstein, M., Wenisch, T.F., Yarom, Y. and Strackx, R., 2018. Foreshadow: Extracting the keys to the intel SGX kingdom with transient out-of-order execution. In 27th USENIX Security Symposium (USENIX Security 18) (pp. 991-1008).

[VANBULCK18B] Van Bulck, J., Piessens, F. and STrackx, R., 2018, October. Nemesis: Studying micro-architectural timing leaks in rudimentary CPU interrupt logic. In Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security (pp. 178-195). ACM.

[VANDERVEEN17] Victor van der Veen, Dennis Andriesse, Manolis Stamatogiannakis, Xi Chen, Herbert Bos, and Cristiano Giuffrdia. 2017. The Dynamics of Innocent Flesh on the Bone: Code Reuse Ten Years Later. In Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security (CCS '17). ACM, New York, NY, USA, 1675-1689. DOI: https://doi.org/10.1145/3133956.3134026

[VANHOEF14] Mathy Vanhoef, Frank Piessens, All Your Biases Belong to Us: Breaking RC4 in WPA-TKIP and TLS, USENIX Security Conference 2015

[VANHOEF17] Mathy Vanhoef, Frank Piessens, Key reinstallation attacks: Forcing nonce reuse in WPA2, ACM SIGSAC Conference on Computer and Communications Security (CCS 2017), pages 1313-1328, Dallas, TX, October 30 - November 3, 2017

[VERDULT12] R. Verdult, F. D. Garcia, and J. Balasch, "Gone in 360 Seconds: Hijacking with Hitag2," In 21st USENIX Security Symposium 2012, Usenix, pp. 237-252, 2012.

[VERVIER15] Pierre-Antoine Vervier, Olivier Thonnard, Marc Dacier, Mind Your Blocks: On the Stealthiness of Malicious BGP Hijacks. NDSS 2015

[VISSERS15] Thomas Vissers, Tom Van Goethem, Wouter Joosen, Nick Nikiforakis, Maneuvering around clouds: Bypassing cloud-based security providers, ACM SIGSAC Conference on Computer and Communications Security 2015 (CCS 2015), pages 1530-1541, Denver, Colorado, USA, October 12-16, 2015

[VISSERS17] Thomas Vissers, Jan Spooren, Pieter Agten, Dirk Jumpertz, Peter Janssen, Marc Van Wesemael, Frank Piessens, Wouter Joosen, Lieven Desmet, Exploring the ecosystem of malicious domain registrations in the .eu TLD, Research in Attacks, Intrusions, and Defenses (RAID 2017), Atlanta, USA, September 18-20, 2017

[VOLCKAERT16A] Stijn Volckaert, Bart Coppens, and Bjorn De Sutter. Cloning your Gadgets: Complete ROP Attack Immunity with Multi-Variant Execution. IEEE Transactions on Dependable and Secure Computing, Vol. 13, Nr. 4, pp. 437-450, 2016

[VOLCKAERT16B] Stijn Volckaert, Bart Coppens, Alexios Voulimeneas, Andrei Homescu, Per Larsen, Bjorn De Sutter, Michael Franz. Secure and Efficient Application Monitoring and Replication. Proceedings of the 2016 USENIX Annual Technical Conference (USENIX ATC '16), pp. 167-179, 2016

[VOLCKAERT17] Stijn Volckaert, Bart Coppens, Bjorn De Sutter, Koen De Bosschere, Per Larsen, Michael Franz. Taming Parallelism in a Multi-Variant Execution Environment. Proceedings of the Twelfth European Conference on Computer Systems (EuroSys '17), pages 270-285, April 2017

[WILLIAMSKING16] David Williams-King, Graham Gobieski, Kent Williams-King, James P. Blake, Xinhao Yuan, Patrick Colp, Michelle Zheng, Vasileios P. Kemerlis, Junfeng Yang, and William Aiello. 2016. Shuffler: fast and deployable continuous code re-randomization. In Proceedings of the 12th USENIX conference on Operating Systems Design and Implementation (OSDI'16). USENIX Association, Berkeley, CA, USA, 367-382.

[WOODRUFF14] Woodruff, J., Watson, R.N., Chisnall, D., Moore, S.W., Anderson, J., Davis, B., Laurie, B., Neumann, P.G., Norton, R. and Roe, M., 2014, June. The CHERI capability model: Revisiting RISC in an age of risk. In 2014 ACM/IEEE 41st International Symposium on Computer Architecture (ISCA) (pp. 457-468). IEEE.

[WOUTERS19] L. Wouters, E. Marin, T. Ashur, B. Gierlichs, and B. Preneel, "Fast, Furious and Insecure: Passive Keyless Entry and Start Systems in Modern Supercars," 21 pages, to appear in TCHES 2019.

[WRESSNEGGER18] Wressnegger, C., Kellner, A. and Rieck, K., 2018, June. ZOE: Content-based Anomaly Detection for Industrial Control Systems. In 2018 48th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN) (pp. 127-138). IEEE.

[YANG10] H. Yang; E. Osterweil; D. Massey; S. Lu; L. Zhang (8 April 2010). "Deploying Cryptography in Internet-Scale Systems: A Case Study on DNSSEC". IEEE Transactions on Dependable and Secure Computing. 8 (5): 656–669.

[ZARPELAO17] Zarpelao, B.B., Miani, R.S., Kawakani, C.T. and de Alvarenga, S.C., 2017. A survey of intrusion detection in Internet of Things. Journal of Network and Computer Applications, 84, pp.25-37.

[ZHANG15] Zhang, D., Wang, Y., Suh, G.E. and Myers, A.C., 2015. A hardware design language for timing-sensitive information-flow security. ACM SIGARCH Computer Architecture News, 43(1), pp.503-516.

5. Research Track 4: Technology Building Blocks: Secure Hardware, Cryptography and Secure Implementations

Contributing Authors: Bart Preneel, Benedikt Gierlichs, Claudia Diaz, Cyprien Delpech de Saint Guilhem, Dave Singelée, Elena Andreeva, Emmanuela Orsini, Frederik Vercauteren, ***Ingrid Verbauwhede****, Nele Mentens, Nigel Smart, Svetla Nikova, Vincent Rijmen, Wouter Castryck*

Scope

This Research Track studies core cybersecurity technologies that underpin all security solutions: cryptographic algorithms and protocols are essential for protection of data at rest and in transit; an emerging trend is the protection of data while it being processed, also known as Computing on Encrypted Data (COED). Research challenges are related to increasing the robustness against post-quantum threats, developing more lightweight solutions and optimizing building blocks for Multi-Party Computation.

There are several reasons why cryptography and security benefit from hardware approaches: first, hardware offers a strong root of trust, to securely store key material, to generate random numbers and to protect sensitive computations; second, hardware implementations offer better performance in terms of throughput, latency and power or energy consumption.

The research on hardware roots of trust will study solutions anchored in technology foundations: Physical Unclonable Functions, true hardware random number generators and silicon odometers.

The last piece of the puzzle are secure implementations: this research will study efficient implementation of novel cryptographic algorithms, the development of countermeasures against implementation attacks and the study of white-box crypto.

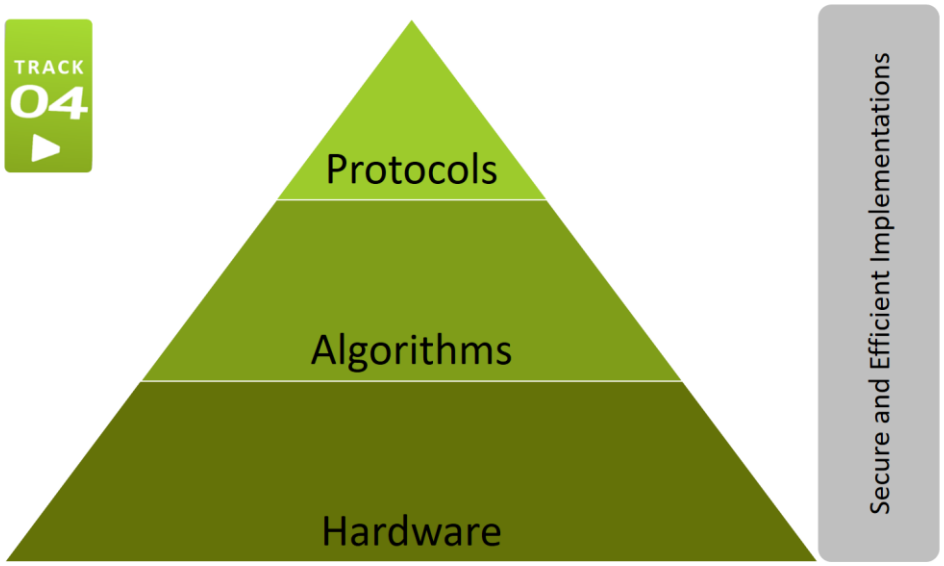

Figure 5-1: Technology Building Blocks: Secure Hardware, Cryptography and Secure Implementations

The following research themes follow in further detail: (I) secure hardware: roots of trust anchored in technology foundations, (II) cryptographic algorithms (III) cryptographic protocols and (IV) secure implementations.

5.1. Secure Hardware: Roots of Trust Anchored into Technology Foundations

In classic hardware security research, the influence of semiconductor technology is not well covered mostly due lack of knowledge of each other's field. The security expertise of COSIC in collaboration with the

semiconductor expertise of imec is therefore a unique proposition. We have identified several research topics that require strong interaction between technology and security.

Physically unclonable functions (PUFs) and true random number generators (TRNGs) are two essential roots-of-trust that are directly impacted by technology development. A PUF provides a chip with a unique static feature, i.e. a chip fingerprint, which cannot be remanufactured and hence can be used to identify this specific chip. Conventionally, a PUF harvests the static entropy from random process variation of the devices or interconnects, to generate the unique chip fingerprint. On the other hand, a TRNG can dynamically provide random bits to the system. The entropy is harvested from the stochastic process that happens in operations, such as thermal noise and clock jitters. In both cases, we need an in-depth study on technology to understand the fundamental cause of process variations and the stochastic process.

A) Industry Needs – Use Cases and Technology Outlook

Customer story 1: Unique identification and key generation is essential in almost any security application, whether it is for a secure web link or card key access to a door or a bank. The security of the application relies on the secrecy of the key. Both PUFs and TRNGs are essential building blocks in this process.

Customer story 2: For equipment that remains in use for decades (e.g., airplanes or trains), new spare electronic parts might no longer be available. Semiconductor companies are not able to still provide electronic chips in old technologies. The equipment owner/manufacturer therefore needs to look for second hand or recycled parts. At this moment, it is extremely difficult if not impossible to derive the usage, i.e. the odometer reading, of a silicon chip. Some chips in older technologies might only have collected dust, others might have been overused, e.g., by overclocking, or by overpowering. Therefore, on-chip silicon odometers that are resistant to attacks need to be developed.

B) State-of-the-art: Highlights

PUFs are typically used to replace secure key-storage and to enable device authentication. Since PUF circuits are affected by noise, the digital output will typically have errors. Early solutions to the problem of error-rates relied on error-correcting codes and fuzzy extractors. Unfortunately, these solutions require costly embedded non-volatile memories (NVM) for storing the helper data. A modern approach in PUF design aims towards high-stability, ideal statistical properties and reconfigurability. Novel PUF types based on gate-oxide breakdown [WU18,CHUANG18:2] and RRAM [CHUANG18] achieve 0% native bit error rate, thereby eliminating the need for NVM for helper data storage. In addition, state-of-the-art implementations of these two PUF types achieve ideal statistical properties without any bias or correlation between generated bits. RRAM PUFs are a promising approach if the application requires reconfiguration. So far, we gained an in-depth understanding of the physical mechanisms and limitations of reconfigurability. We have the quantitative analysis to show the severity of the flaws in existing solutions. The logical next step is to investigate and apply post processing methods to combat these flaws. Another topic that is missing in PUF state-of-the art is the correct way to evaluate PUFs based on physical models rather than statistical tests.

A reliable age monitoring solution is needed to prevent chip recycling and reselling. Several implementations of "silicon odometers" have been proposed using degradation mechanisms such as Biased Temperature Instability (BTI), Hot-carrier Injection (HCI) or Electromigration (EM) [KEANE10, WANG14]. These implementations, however, do not provide resilience against tampering. Since most of the degradation mechanisms can be recovered to some extent, there may exist loopholes for tampering the recorded aging information. One example is the odometer based on ring oscillators (RO), the oscillation frequency of an RO can reflect the amount of hot-carrier injection (HCI) degradation caused by operation. Normally the HCI degradation only gets worse with time, but there is an effect called thermal curing, which means the degradation might be reversed by baking the chip at higher temperature. An adversary can therefore use this mechanism to cheat. A more comprehensive study on the reliability physics is needed to design an accurate and tamper-resistant silicon odometer.

TRNGs used in security applications have to comply with latest security standards such as German AIS-31 [KILLMANN11] and American NIST SP 800-90B [TURAN18]. Unlike the past approaches which relied on statistical testing of the generated bits, a modern design approach requires a formal evaluation of TRNG security based on the stochastic model of the randomness generating process. This approach requires deep understanding

of the way that the noise is generated in the circuit and accurate measurement of technology parameters such as noise strength. In addition, TRNGs require on-the-fly test modules to monitor the health of the entropy source [Yang16] and post-processing modules that improve the entropy of the generated random data. One interesting option for implementing post-processing modules is using entropy extractors. Entropy extractors are Boolean functions that formally guarantee the full-entropy output as long as the input entropy is above some minimal prescribed limit. Entropy extractors [BARAK06,DODIS04] have been studied in computer science for many years but only a few hardware implementations are available in open literature [MATHEW16,ROZIC18].

C) Main Areas of Work

As the silicon technology follows Moore's law, transistor dimensions have scaled down to the size of tens of atoms. The variability in these advanced technologies increased drastically, since a small difference can make a big impact on this tiny scale. Even though this increased variability is one of the major challenges for technology scaling, it can give a positive impact on PUFs as well as on other security primitives. On the other hand, since the end of CMOS technology scaling can almost be seen, more so-called beyond CMOS solutions are being studied. In the pathfinding practices of these new technologies, there are plenty of opportunities to be explored from a security perspective. In collaboration with the reliability group of imec and the hardware security group of COSIC, three main areas of research have been identified.

(RA 4.1.1) Developing PUFs

Technology insights for PUF evaluation: By understanding the fundamental causes of the time-zero and time dependent variability, we can both improve the performance of PUF and have a good design-time estimation on the entropy. Although a well-defined statistical model is provided by the foundry, the chip designers or end users typically do not know where the randomness comes from. A better understanding of the variability in these process steps will help the designers to develop and evaluate PUF circuits at design time. Second, it can help us to select a technology option that favors PUFs. Finally, it helps on developing methods to examine if enough entropy is provided after the chip has been fabricated.

Developing PUFs in emerging technologies: This direction is focused on advanced CMOS technology nodes and emerging memory technologies. For example, in finFET technologies, both transistors' vertical spacers are affected by variability of the manufacturing process. Moreover, the 3D stacking technique, which is widely considered to realize the so-called system-in-package, also provides a large amount of variability to be explored. Regarding resistive memory devices such as RRAM and MRAM, in addition to the device-to-device variation, there is another term called cycle-to-cycle variation, which is caused by the inevitable stochastic switching behavior. By properly using these variations, it is possible to make a big advance on PUF designs – realizing truly reconfigurable PUFs.

(RA 4.1.2) True Random Number Generators

Self-tuning TRNG circuits. TRNGs are by design technology-dependent building blocks. To integrate such building blocks in a system-on-chip, one typically has to first characterize the target technology using a test chip, tune some design parameters and then implement the building block. This cycle requires time, expertise and costs money. We will investigate and develop self-tuning TRNG circuits. That is, circuits that can self-adjust their parameters on startup to provide good performance for the technology and environment parameters. We will develop self-tuning TRNGs and quantify the design trade-offs that they can provide.

Entropy extractors. Our goal is to identify entropy extractors that are suitable for hardware implementation, to explore the design trade-offs under the hardware limitations (e.g., interface, area) and to develop hardware prototypes.

Designing efficient entropy sources. Based on imec's expertise on semiconductor physics we will obtain better in-depth understanding on how the noise in electronic circuits is generated. In addition, imec has expertise on circuit design, which helps to design more efficient and reliable devices. Combining these, we aim to realize better TRNGs that are compliant with the latest security standards.

(RA 4.1.3) Technology Solutions to Secure Circular Economy

Silicon odometers. Similar to vehicles, a silicon odometer is needed to record the actual usage of a chip to mitigate the threats of IC counterfeiting and abusing of device lifetime. A silicon odometer needs to reflect

the amount of device degradation, and should be tamper-free, i.e. the odometer cannot be turned-off at any time, and the record cannot be modified. With our understanding on device reliability, we will investigate different types of degradations, to design more precise silicon odometers and robust countermeasures for physical tampering.

D) Expected Outcomes and Road Map

For (RA 4.1.1), the first result (Y1) will be the security analysis of PUFs based on physical and stochastic models, rather than just black box statistical models. We will explore the PUF types that can provide the reconfiguration ability, such as RRAM PUF. Current simulation results and physical models [CHUANG18] show that re-configuration is not ideal. Therefore, we need to use some type of digital post-processing to compensate for these non-idealities to assure that the circuit is truly reconfigurable. The second result (Y2) will be a prototype containing a PUF followed by the chosen post-processing module, integrated into an application (such as device authentication).

For (RA 4.1.2), the first result (Y1) will be developing reconfigurable, self-tuning entropy sources and testing their efficiency on FPGA prototypes. The second result (Y2) will be efficient implementations of entropy extractors. The long-term goal (Y5) is to fabricate a TRNG prototype using novel entropy sources and the proposed post processing modules.

For (RA 4.1.3), the long-term goal (Y5) is to develop a prototype of an accurate and tamper-resistant silicon odometer.

5.2. Cryptographic Algorithms

Cryptology is the science that studies mathematical techniques in order to provide secrecy, authenticity and related properties for digital information. The goal is to protect data while in transit, in storage and during computation; it lies at the core of the protection of digital data and processes. It also allows establishing trust relationships over open networks and enables the collaboration of mutually distrusting parties towards achieving a common goal. Cryptology is a fundamental enabler for security, privacy and trust. Today cryptographic techniques are widely deployed at the core of computer and network security, and for applications including finance, digital transactions, secure authentication However, there are a number of important challenges that are not addressed by the current state-of-the-art deployed cryptographic algorithms:

- For the Internet of Things, and in particular for applications with low-cost embedded devices, there is a need for novel cryptographic designs for authenticated encryption that improve the current tradeoffs between power/speed/energy/area/latency; in this context reducing latency and energy consumption are two problems that so far have received insufficient attention.
- For cloud environments there is a need for advanced research on Computing on Encrypted Data (COED); the most promising techniques are Multi-Party Computation (MPC) and Somewhat Fully Homomorphic Encryption (SFHE). For specific problems, more efficient dedicated protocols can be developed.
- As Computing on Encrypted Data (COED) becomes increasingly important, symmetric cryptographic algorithms should be designed to optimally fit the constraints of these new contexts; this includes minimizing the total number of AND gates and/or the AND depth.
- The progress on developing large quantum computers poses a major threat to most public-key deployments and a moderate threat to many symmetric key systems. While some progress has been made in the past decade on the development of post-quantum cryptography, a substantial amount of research is needed to develop robust solutions that would be suitable replacements for the current solutions in the next 5-7 years.

A) Industry Needs – Use Cases and Technology Outlook

There is a strong industry demand for novel cryptographic solutions. As the developments are mostly driven by standardization, it is important to align the work with current and future standardization initiatives in standardization bodies such a NIST, IETF, ISO/IEC and ETSI. Currently NIST has ongoing efforts on postquantum cryptography, lightweight cryptography and threshold cryptography. Moreover, one can anticipate that in the next 3-5 years standardization efforts in Computing on Encrypted Data (COED) will pick up.

Customer story 1: company in finance or insurance needs to migrate towards cryptographic algorithms and protocols secure against attacks on quantum computers; this is essential for security of 15 years and more.

Customer story 2: company developing implantable medical devices requires a cryptographic algorithm that can work on a very low energy budget in order to avoid early exhaustion of the battery (or usage of harvested energy) and in order to avoid local heating of the human body.

Customer story 3: Two or more companies want to perform some computation on their joint private data, without revealing to each other the exact nature of the data. This could for example be a set of companies organizing a complex supply chain in which demand and supply of capacity is a price sensitive variable, or companies wishing to engage in an auction (such as financial institutions trading stocks or currencies), or could be organizations wishing to extract data from joint records on individuals without compromising privacy. For these purposes methods to compute on encrypted data are necessary.

B) State-of-the-art: Highlights

Symmetric-key algorithms form the backbone of the security architectures that protect critical infrastructures such as banking, mobile networks and the Internet. For these applications there exist portfolios of secure algorithms (including the widely deployed AES algorithm [DAEMEN02] and the algorithms of 3GPP [ETSI17,ETSI19]), however the search continues to improve the trade-offs between security, cost and performance and in particular long-term security.

Secondly, with the advent of smart devices, there comes a growing need for new primitives that are better suited for low-end processors and nodes with tight energy constraints [NIST17L]. Additionally, the increasing demands on secure computing platforms (studied in Research Track 3) require primitives that can protect the memory and the communication lines inside computing platforms; in this context, latency is important. Examples of such designs are Chaskey [MOUHA14], PRINCE [BORGHOFF12] and SKINNY [BEIERLE16]. The CAESAR competition has yielded new insights in the robustness of these primitives and novel constructions for authenticated encryption but has also generated new design approaches and open problems [CAESAR19]. A novel design in this domain has been the forkcipher [ANDREEVA18], that is suitable for efficient lightweight authenticated encryption of short messages [ANDREEVA19].

Thirdly, the deployment of platforms for Computing on Encrypted Data (e.g., multi-party computation and homomorphic encryption) requires new primitives that can be implemented efficiently on these platforms. In recent years the advancements in these areas have enabled a growing number of practical solutions. In MPC linear operations come almost for free but non-linear operations are extremely expensive. This step incurs communication cost and is often performed by using a so-called ”Beaver triple”, which is precomputed in an offline phase. There is emerging research area that studies primitives for symmetric key encryption that are particularly suitable for MPC and FHE-based systems. There are a first set of designs and also some novel attacks [DOBRAUNIG18,GRASSI16,RECHBERGER18]. The ongoing activities of NIST in the standardization of lightweight cryptographic algorithms [NIST19L] form just one illustration of the liveliness of this topic. Most of the schemes for arbitrary length data (modes of operation) used in the classical two-party setting in symmetric cryptography come with security proofs only over binary fields. Only a few works [ROTARU17,GRASSI16] have so far researched the suitability of the classical symmetric modes of operation for MPC. Due to the different cost metrics in MPC and the fact that many practical MPC protocols work over large prime fields, when working with non-binary fields in the MPC setting one cannot directly apply the classical modes designed for binary fields. The design principles for modes over such MPC native fields are not well understood.

Finally, most applications that involve the protection of digital rights require primitives that are suitable for white-box or gray-box cryptography – the research on the security in these models is included in Theme 4 on secure implementations).

Following the current insights, symmetric-key algorithms offer a better resistance against attacks using quantum computers, hence proposals for applications such as electronic signatures based on symmetric-key techniques attract more and more attention (e.g., [BERNSTEIN15]).

Since the foundational work of Diffie and Hellman (DH) [DIFFIE76], public-key cryptographic algorithms have become ubiquitous in communication technology. Indeed, algorithms such as DH key exchange (DHE)

[Diffie76], RSA encryption [RIVEST78] and the DSA [NIST94] – together with their more efficient elliptic curve-based variants such as ECDHE and ECDSA – are used every day by billions to securely communicate over the Internet. The most recent TLS 1.3 standard [IETF18], completed last year and intended to bring a significant improvement to the security of Internet communications, still makes use of these public-key algorithms.

However, these widely deployed algorithms are not resistant to quantum computers and in the last decade there has been a strong drive of research to develop replacement post-quantum primitives. This culminated in the ongoing NIST Post-Quantum Cryptography (PQCrypto) competition [NIST16] which has received submissions based on different mathematical assumptions. Several are constructed from lattice-based assumptions (CRYSTALS [CRYSTALS17], FrodoKEM [FRODOKEM17], etc.) but others use different assumptions such as elliptic curve isogenies (SIKE [SIKE17]) or multivariate quadratic (MQ) systems – with COSIC's SABER [SABER17] and LUOV [LUOV17] submissions having advanced to Round 2 of the competition in January 2019. It is still unclear which assumption(s) will provide the best security against quantum computers and much active research is currently directed at understanding the limits of the proposed primitives. COSIC's work is well represented here as well with recently published attacks [Beullens18] and numerous contributions to the discussion forums.

Another significant line of research is that of threshold computation of public-key algorithms. Recent works have proposed constructions for the classical ECDSA algorithm as it is currently the most widely deployed [DOERNER18, GENNARO18, LINDELL18]. Furthermore, NIST has also issued communications indicating that further research should be conducted in this direction [NIST19]. COSIC is also present in this area, notably with the upcoming publication of a threshold construction for a lattice-based post-quantum algorithm [KRAITSBERG19].

Secure pseudorandom number generators (PRNGs) (also known as Deterministic Random Bit Generators or DRBGs) underpin the vast majority of cryptographic applications. Relative to its importance, this area is under-researched. Moreover, multiple issues have been identified with deployments. For example, the NIST SP 800-90A [SHUMOW07] has had a troubled history, particularly in light of the now infamous DualEC-DRBG. Recently, Woodage and Shumow [WOODAGE19] analyzed the security properties of the three novel DRBG mechanisms (HASH-DRBG, HMAC-DRBG, CTR-DRBG) in the NIST SP 800-90A [BARKER15] standard.

C) Main Areas of Work

The research is divided into three activities: one on symmetric cryptographic algorithms, one on public-key cryptographic algorithms and one on validation and proofs.

For each of these research lines a combination may be required of mathematical research on hard problems to base the security on and of adapted security tools and models for evaluation.

(RA 4.2.1) Symmetric-key Algoritms

Our research on the security of novel authenticated encryption algorithm will build on the results of the CAESAR competition [CAESAR19]. Further study is required to increase the assurance in these constructions; moreover, the know-how built up during this competition is a useful starting point to develop novel designs that offer improved security/performance tradeoffs. Our approach will cover both the design and analysis of new building blocks; the corresponding security reductions based on formal assumptions on the building blocks will be studied in (RA 4.2.3).

Study of lightweight symmetric primitives for low-end processors and nodes with tight energy constraint. This research will include both cryptanalyzing existing proposals and designing new proposals; the work will feed into NIST ongoing standardization activity in the area of Lightweight cryptography) [NISTL].

COED platforms have resulted in new optimization criteria for symmetric cryptography such as the reduction of the number of AND operations per bit or the reduction of the depth of the AND gates. Moreover, it can be advantageous to design symmetric algorithms over fields of characteristic larger than 2. This results in new design principles, in a renewed interest in algebraic attacks and in development of novel cryptanalytic techniques. In addition, these building blocks enable services such as distributed/threshold Oblivious PRFs which have a strong potential for applications.

(RA 4.2.2) Public-key Algorithms

Development and evaluation of new post-quantum primitives based on lattices, isogenies and multivariate quadratic (MQ) systems. This work will analyze existing basic primitives such as digital signatures, encryption and key-establishment and propose new constructions. It will also develop primitives with higher privacy preserving functionalities than that of the current NIST project. These advancements in post-quantum primitives will serve to design higher level post-quantum protocols such as authenticated key-establishment, identification schemes, oblivious transfer or zero-knowledge proofs in Theme 3.

In addition, we will develop threshold constructions of these post-quantum primitives to suit a wide variety of distributed environments and their respective threat models. This will also drive advancement in algorithms for general threshold computation technologies.

Both of the above will be influenced by and feed into NIST standardization activity in the area of post-quantum cryptography [NISTP] and threshold cryptography [NISTL].

The above will also drive advancements on the core algorithmic components of FHE and MPC as well as on dedicated algorithms for specific computations such as private information retrieval (PIR) and private set intersection (PSI). These algorithm-level improvements will take into consideration both lower-level implementation (theme 4) and higher-level protocol (theme 3) requirements in order to provide the most adapted security and performance improvements.

(RA 4.2.3) Proofs and Validation

This research activity focuses on the analysis and exploration of proof techniques for validating the security of cryptographic primitives. It will follow the principles of the provable security paradigm and make use of security models and security reductions to provide formal guarantees of security.

The work will apply across both symmetric (RA4.2.1) and public-key (RA4.2.2) primitives to analyze both existing state-of-the-art proposals and also all new designs constructed within this theme. For each category of primitive, first the best models will be identified, then each new primitive will be analyzed according to its relevant model and different design techniques will be compared based on the security guarantees that they provide. This analysis will drive developments in proof techniques for all categories of cryptographic primitives.

A particular area that will be investigated in symmetric cryptology are modes operation for Pseudo-Random Number Generators (PRNGs): while PRNGs are widely used but they are some of the least studied cryptographic building blocks. In our research, we will analyze the security of widely used PRNGs such as CT-DRBG; hereby we will use alternative designs and proof approaches. A second line of research will be the analysis off forkciphers and alternatives to the Feistel cipher solutions, such as the one of Lai and Massey [LAI91].

D) Expected Outcomes and Road Map

For (RA 4.2.1) we will develop novel cryptanalysis methods to give improved insights in the security of lightweight algorithms (Y1). A next goal will be the design of novel symmetric-key algorithms over fields with characteristic larger than 2 for COED applications (Y2). A longer term effort is the design of algorithms for applications with very tight energy constraints (Y5).

For (RA 4.2.2) we will develop proofs for the classical symmetric modes in the MPC setting (Y1). We will develop and design novel MPC-friendly proofs, schemes and models. The combination of both results will enable us to provide a comprehensive security model and a set of symmetric tools suitable for the MPC setting.

For (RA 4.2.3) we will analyze existing primitives within advanced new security models (such as PRNGs) and new constructions (such as the forkcipher) within established models (Y1). We will refine the requirements for advanced new models and established models and then analyze and compare novel designs, and variations thereof, of primitives within the relevant models (Y2). As a

5.3. Cryptographic Protocols

While cryptographic algorithms form the core protection technology, these algorithms are typically assembled in cryptographic protocols to achieve specific goals. A central goal is the creation of secure channels, which builds on Authenticated Key Establishment (AKE). Widely deployed examples are SSL/TLS, IPsec, the 3GPP Authenticated Key Agreement (AKA) protocols, and the Signal protocol.

Advanced cryptographic protocols can be extremely beneficial for enhancing security and privacy in the cloud and in complex application protocols; there is a need for sophisticated protocols that avoid single points of failure (through distributed cryptography) and that allow to reconcile the mutually conflicting interests of the stakeholders (e.g., privacy of user data, correctness of results). While there has been substantial progress in the past decade, there is still a very large gap between academic work and industrial practice; there are challenges related to performance, the availability of cryptographic libraries with the right features, and the lack of understanding of the potential of this approach.

The areas studied in this theme will be distance bounding protocols to defeat relay attacks, Oblivious Transfer protocols to enhance Multi-Party Computation (MPC) protocols, distributed consensus protocols for blockchain and mix networks to hide metadata in transaction and messaging systems.

The research is divided into five activities: distance bounding, protocols for MPC, protocols for blockchain, protocols for protecting metadata and security analysis for protocols.

A) Industry Needs – Use Cases

As for cryptographic algorithms, there is a growing industry demand for novel cryptographic protocols. For basic protocols such as Authenticated Key Agreement (AKA), the developments are strongly driven by standardization and in particular by IETF and ETSI. For more advanced protocols, such as distance bounding, Multi-Party Computation (MPC), distributed consensus algorithms and protection of metadata, there is more room for dedicated solutions. One can expect that in the next decade standardization will become increasingly important in some of these areas.

There is a need for further analysis tools to support the analysis and development of cryptographic protocols. The basic protocols for the protection of network traffic and data at rest are well understood, but there are still issues with widely protocols that protect the core of the Internet today. Overall, there is strong need for tools and methods to support the security analysis and validation of cryptographic algorithms.

Many environments require secure authentication of users or devices to other devices. It is well known that standard protocols for entity authentication are vulnerable to relay attacks, in which an opponent interjects itself between the two entities; in the past years such relay attacks have been deployed by criminal organizations, for example to defeat Passive Keyless Entry Systems (PKES) for cars. The automotive and financial industries are currently exploring several options; one can expect that this will expand to other sectors that are relying on IoT.

Advanced cloud interactions will increasingly use Multi-Party Computation (MPC) to create value out of data while protecting the legitimate interests of each party; Oblivious Transfer (OT) protocols form a central building block of MPC protocols. While these protocols are not directly visible to the users, they form a central element to improve the performance.

Blockchain technologies show great progress in achieving distributed consensus in an innovative an open way. By combining immutability and transparency, blockchains are a great tool to store transaction logs or any other data that can then later be audited.

Customer story 1: company requires authenticated key management protocol to securely contact and update a large number of sensor nodes in the field.

Customer story 2: company in the automotive sector needs to provide secure passive keyless entry system into a car that is robust against sophisticated relay attacks.

Customer story 3: company wants to deploy a new blockchain protocol that offers increased privacy and robustness in order to optimize a large logistics chain with multiple parties.

B) State of the art: Highlights

Whilst early works quickly arrived at stable notions of security for cryptographic algorithms, the state of security models for more complex protocols has not been so definite. Each area of protocol design is rich with a multitude of models; for example, for Authenticated Key Establishment (AKE) early works date to 1994 [BELLARE94] but there have been many additions to and divergences from the original definitions [CANETTI01B,CREMERS12]. For other areas, an overarching framework that captures protocols with different goals (such as the Universal Composability (UC) framework [CANETTI01A] for MPC) has been developed and it has allowed for flexible design of security notions despite its complexity. There has been extensive analysis of widely used protocols such as TLS (e.g., [MORRISSEY10]). Recently, there has also been an emergence of automated verification tools, such as [PROVERIF] or [EASYCRYPT] which can be used to ease the analysis of new protocols.

Many contactless, radio-based systems, for example electronic payment or Passive Keyless Entry Systems (PKES), are vulnerable to distance-based frauds. Particularly relay attacks are of interest, and have been demonstrated both in a lab setting as well in daily life by criminal organizations. Beth and Desmedt [BETH90] introduced the notion of timed message exchanges to prevent relay attacks. This idea has been generalized later on, first by Brands and Chaum [BRANDS94] and later by Hancke and Kuhn [HANCKE05], resulting in the concept of distance bounding protocols. Distance bounding protocols allow a verifier to both authenticate a prover and evaluate whether the latter is located in his vicinity. To determine an upper bound on the distance between verifier and prover, distance bounding protocols combine physical and cryptographic properties. However, it is a challenging research problem to design and implement these protocols in practice. One of the first attempts was made by Rasmussen and Capkun [RASMUSSEN10], and later improved by Ranganathan et al. [RANGANATHAN12], based on an analog or hybrid design approach. Some proposals rely on Ultra-Wide Band (UWB), for example by Tippenhauer et al. [TIPPENHAUER15] and Singh et al. [SINGH19]. However, this is still a rather unexplored research area, and there is a clear need for practical, low-cost and secure designs and implementations of distance bounding protocols.

The last ten years have seen a remarkable advance in practical secure multi-party computation (MPC) protocols, to the extent that many problems arising in practice can be solved. One of the main cryptographic building block of many efficient MPC protocols, is Oblivious Transfer (OT), originally proposed by Rabin in 1981 [RABIN81]. OT protocols are used in the construction of a range of protocols, in particular, OT is sufficient and necessary for secure multi-party computation, and is often also used in special purpose protocols for tasks such as private set intersection. Thus, the efficiency characteristics of the OT protocol directly affect the efficiency of the resulting secure computation protocol. Unfortunately, OT requires public-key machinery, so even the most efficient oblivious transfer constructions come with a relatively high cost. In 2003, it was discovered how one can “extend OT” starting with a small number of base OTs to create essentially an unlimited number of OTs using only symmetric primitives [ISHAI03]. This passive-secure protocol was later generalized to support active-security in 2015 by [KELLER15]. Other efficient general-purpose secure-multiparty protocols [NIELSEN12,WANG17,CRAMER18] are based on OT. Even though these extension protocols are very efficient from a computational point of view, they still require high communication costs. This impacts the communication efficiency of the resulting secure computation protocols with a large number of parties. Other than efficiency, another problem regarding OT is security: the most efficient OT constructions [PEIKERT08,CHOU15] are not post-quantum secure since they are based on number-theoretic assumptions such as the Decisional Diffie-Hellman problem (DDH). This means that the possible future advent of quantum computer will immediately render insecure not only OT constructions, but also all the protocols and applications based on it.

Blockchain technology is poised to make a major impact on the way organizations deal with data and how transactions are being processed. The innovation driven by blockchain was inspired by the development of Bitcoin [NAKAMOTO08] and other cryptocurrencies as well as the emergence of novel platforms for smart contracts such as Ethereum [WOOD14]. Bitcoin provided an innovative solution to distributed consensus (i.e. the Byzantine generals’ problem) for an open system, that is, a system in which the number of players is a

priori unknown and can change over time [GARAY15]. This solution made use of chained Merkle trees, a technology invented for digital timestamping [HABER90,MASSIAS99], and proofs of work, a technology developed to fight SPAM [BACK02,DWORK92,DWORK03] and also proposed for electronic cash [DAI98]. The main challenges of the cryptocurrencies and transaction systems building on these proofs of work are the scalability [CROMAN16] (in particular, the energy consumption of the Bitcoin proof of work is estimated to be comparable to that of Austria or Greece, while less than 10 transactions per second can be processed) and the stability against malicious adversaries [ZHANG19]. At this moment there is a quickly growing body of research that attempts to address these problems. Moreover, the transparency of public blockchains raises a trilemma among public verifiability, privacy and performance, which researchers are striving to break.

The protection of metadata in communications and transactions is becoming increasingly important. Due to the fast reduction of the cost of public-key cryptography, many applications can now deploy end-to-end security to protect the contents of communications. However, at the same time there is a growing understanding that most of our modern systems expose huge quantities of metadata such as location, communication patterns, social relations, and transaction patterns. Besides being unprotected and available to a variety of entities, metadata is in machine readable formats and therefore very easy to process and analyze at scale [DIFFIE07]. A growing number of players are collecting and exploiting this metadata for their own purposes. This is a concern, as “metadata is data” and in particular metadata can reveal highly sensitive information. For example, the fact that a person interacts with an oncologist, a divorce lawyer, or an investigative journalist might be very revealing even if the contents of their conversation are not available. Moreover, while there is extensive research on technologies to hide metadata, the deployment of these technologies is rather limited [SHIRAZI19]. The only solution that has millions of users is the Tor network [DINGELDINE04]. Tor provides low-latency bidirectional channels suitable for real-time communications such as web traffic; that can hide the identity of clients and offer hidden services. While Tor is a valuable tool, it has become clear that it has serious limitations, which are inherent to its low latency and low bandwidth constraints [JOHSON13, OVERDORF17]. Therefore, there has been an increased interest in mix networks, which were first proposed by Chaum in the 1980s [CHAUM81]. While mix networks impose longer delays, they also allow to offer much stronger security guarantees, including strong anonymity properties towards global adversaries. Furthermore, the increased bandwidth capacity of modern networks allows for the use of dummy traffic techniques in combination with mixing, which further strengthens the privacy properties offered by mix networks [PIOTROWSKA17].

C) Main Areas of Work

The research is divided into five activities: distance bounding, protocols for MPC, protocols for blockchain, protocols for protecting metadata and security analysis for protocols.

(RA 4.3.1) Cryptographic Protocols for Distance Bounding

Relay attacks can be prevented by the use of secure distance bounding protocols. However, one of the main research challenges is to design and implement a distance bounding protocol that offers good trade-offs between security, cost and deployability:

- Security: Two categories of attacks need to be considered: protocol attacks (e.g., a relay attack) and implementation attacks. The latter are attacks that aim to exploit specific implementation details or characteristics of the wireless communication process to carry out a protocol attack. Examples are early-detect and late-commit attacks, CICADA attacks on UWB systems, etc.
- Cost: Ideally, the distance bounding protocol should be implemented on low-cost devices such as a key fob, an IoT device or a medical implant. Another interesting target platform will be a smartphone. In all cases, the cost of integrating distance bounding should be limited.
- Deployability: This is strongly related to the cost. The distance bounding protocol should be compatible with existing communication standards, such as Bluetooth, UWB or VLC (Visible Light Communication). Specifically Bluetooth is of interest, as it is already supported by a large number of devices and systems.

This Research Activity will focus on the secure implementation of distance bounding protocols, considering the trade-off mentioned above. The following research tasks are planned:

- Design and implementation of distance bounding protocols.

- Security evaluation, with particular focus on the prevention of relay attacks. This could include formal verification, as well as practical experiments.
- Building security services on top of distance bounding: This could include authentication and access control, device pairing, secure localization, but also secure network topology services.

(RA 4.3.2) Cryptographic Protocols Design for MPC Applications

A first line of research consists of improving the communication costs of OT extensions protocols using code-based assumptions, or relaxed definitions of OT, that can still be used in the design of secure multiparty protocols and applications.

A second line of research can be outlined by noting that OT can be constructed by different assumptions; whether the construction assumes the Random Oracle Model (ROM), or a Common Reference String (CRS), or has no such assumption; as well as the underlying hard problem on which security is based, for example the Computational Diffie-Hellman (CDH) assumption CDH or the Decision Diffie-Hellman (DDH) assumption, or Learning Parity with Noise (LPN) or Learning with Errors (LWE) for lattice-based construction, or McEliece style problems for coding-based constructions. Unfortunately, the most efficient constructions are those based on DDH and CDH. However, any future quantum computer would enable efficient breaking of security of these discrete-logarithm-based variants. We will focus our research on efficient constructions based on lattices, isogenies and coding theory problems.

We will integrate OT-based protocols in our current MPC framework SCALE/MAMBA [SCALE], which is mainly based on lattice-based homomorphic encryption, allowing efficient binary and arithmetic circuit evaluation.

(RA 4.3.3) Cryptographic Protocols for Blockchain

Further analysis of Bitcoin's Nakamoto consensus (NC): As the first consensus protocol in an open system, NC receives more attention from academia than all its successors. However, to date, security analysis of NC only focuses on simplified versions of the protocol in simplified network settings. A series of research questions remain unanswered when more realistic network settings are considered. For example, how are the participants' incentives affected in the presence of a network-level attack? As NC fails to achieve security in an asynchronous model for which it has been designed, is it possible to modify the protocol so that it can achieve stronger security properties in a synchronous model? As we learned that slower block propagation is in favor of big miners, what are the boundary conditions of this phenomenon? With proper modeling, we can answer these questions, which allows us to design more robust protocols that take network conditions into account.

Unified framework for analyzing complex consensus protocols: A considerable number of consensus protocols are proposed following the footstep of NC, all claim to achieve stronger security properties. However, it is difficult to evaluate whether there is actual technological progress as these protocols prove their security under different assumptions. With tools from artificial intelligence, we can quantify the effectiveness of these protocols under the same security assumptions and compare their security and performance, which will lead to additional insights.

Performance evaluation for alternative consensus protocols: Acknowledging the scalability barrier of the underlying consensus protocols, a growing number of blockchains adopt alternative ledger topologies, i.e., direct acyclic graphs (DAG), or sharding—splitting transaction addresses into different zones—for better throughput. However, each of these approaches introduces additional performance tradeoffs. We will simulate these protocols to compute their actual throughput and determine whether these approaches are feasible and if so, when can their performance gain outweighs the cost. By eliminating impractical approaches, we aim to unify the community for further throughput improvements on the practical ones.

Breaking the public-verifiability, privacy and performance trilemma: Most existing smart contract platforms sacrifice privacy and performance for public verifiability. To enable better privacy or better performance, researchers propose to use slow and heavy cryptography or to partially sacrifice public verifiability, respectively. We aim to look into the details of real-world smart contract use cases and design protocols and smart contracts that achieve all three properties simultaneously. This work has the potential to accelerate the adoption of smart contracts.

(RA 4.3.4) Cryptographic Protocols for Mix Networks

Mix network designs that provide secure, scalable and deployable anonymous communication channels remain an open challenge. In this research we will tackle a number of questions that need to be solved to make such networks deployable.

First, we will study the security of mix network protocols considering adversaries who observe large portions of the network and have the ability to corrupt a significant number of mixes. The goals are to devise strategies for mix topology selection that cannot be influenced by adversaries, to detect misbehavior in a way that enables the exclusion of malicious mixes, and to understand how privacy properties degrade when the percentage of corrupted mixes increases. We will study how MPC protocols can be used in mixnet designs to improve resilience towards adversarial compromise.

Second, we will investigate the effects of routing different classes of traffic over a mixnet. Existing work in mixnet design typically considers that only one class of traffic is being routed by the network. It is not known how aggregating diverse types of traffic over the same mix network affects the security guarantees. Yet this is a crucial question in order to design general-purpose mixnets that can support multiple use cases and thus function as a privacy enhanced communication infrastructure.

Finally, we will study the extent to which persistent communication patterns can be recovered by combining observations over an extended duration and devise cover traffic strategies that provide protection for those long-term patterns.

(RA 4.3.5) Security Analysis of Cryptographic Protocols

The research will be conducted along two interacting axes. The first will be the study and design of security models and frameworks. The second will be the application of these models to existing protocols (proposed and deployed) and to those newly designed as part of this research project; this application will also inform further elaboration of appropriate security models.

The core elements of the design axis will be models tailored to specific applications. The construction of such models will be informed by existing models and by newly arising requirements from protocols. This will include research on security notions defined within the UC framework with the refinement of existing definitions and the elaboration of new ones. Alongside these core elements, an overarching framework for defining models of security will be considered, this will allow more direct comparisons of new models. This framework will then be used to expand the scope of the models.

The second axis will construct security proofs for the advanced protocols of this project within the relevant models studied in the first axis. On one hand, this will be used to analyze existing protocols (proposed and deployed), identify construction and proof flaws and, in combination with work from the RAs above, propose mitigation elements. On the other hand, the relevant models and frameworks will be used to construct proofs for all newly developed protocols as part of this theme. This will strongly support every advancement with provable security guarantees.

In addition, the latest parameter estimates for the underlying algorithmic primitives used (as evaluated in Track 2) will allow for the computation of recommended parameters for each new protocol. These concrete parameters will be established to match security levels set by standardization works such as the NIST competition and to follow developments in cryptanalysis.

D) Expected Outcomes and Road Map

For (RA 4.3.1), the first result (Y1) will be the development of the basic building blocks of a distance bounding protocol that is designed particularly for narrow-band systems (e.g., based on Bluetooth). The second result (Y2) is a security evaluation and a proof-of-concept implementation of the full narrow-band distance bounding protocol. The third result (Y5) is the development of distance bounding protocols based on other radio principles, for example UWB.

For (RA 4.3.2), the first result (Y1) will explore different strategies for communication efficient OT-extension and post-quantum-secure OT constructions, based on different post-quantum assumption in the UC-security

model. The second result (Y2) will implement these solutions and integrate them in the SCALE/MAMBA system.

For (RA 4.3.3), the first result (Y1) aims to solve several concrete challenges in order to gain a deeper understanding into the technology: (1) What are the attacker's optimal selfish mining strategies in combination with an eclipse attack? Whether the eclipsed victim has an incentive to defend against this attack? (2) What are the boundary conditions when slower block propagation is in favor of big miners? How will they affect the future of cryptocurrency mining ecosystem? (3) Simulate DAG and sharding protocols to compute their actual throughput and determine whether these approaches are feasible and if so, when can their performance gain outweighs the cost. (4) Design several concrete smart contracts that achieve public verifiability, privacy and performance and the same time. The second result (Y2) will be the design better protocols based on the understandings we gained from the first period: (1) Design a general smart contract framework for breaking the public verifiability, privacy and performance trilemma. (2) Design a modified NC protocol that achieves stronger security properties within a synchronous model. The third result (Y5) aims for a series of complete analyzes of all permissionless consensus protocols.

For (RA 4.3.4) the first result (Y1) will be the development of models that can be used to analyze the security of mixnets towards corrupted nodes and that consider networks with heterogeneous traffic. The second result (Y2) is the development of mixnet designs that offer improved resilience towards corruption and that can effectively blend together traffic with different characteristics (volume and latency constraints). The third result (Y5) will be the development of analytics tools for measuring long-term disclosure, as well as countermeasures to prevent such disclosure.

For (RA 4.3.5), the first result (Y2) will be proofs and analyzes of newly constructed advanced and existing protocols for one or two selected application areas within relevant models. The second result (Y5) will expand the first models and protocols, and design variations thereof, and also focus on additional application areas and their relevant models.

5.4. Secure and Efficient Cryptographic Implementations

The main focus of this research topic is development of efficient cryptographic software and hardware libraries with built-in security against side-channel and fault attacks. This research will be performed in close collaboration with the research on cryptographic algorithms and protocols (see previous two topics), with the technology input of topic 1. The results of this research topic are directly useful in many use cases in Research Track 2 on Security services and Research Track 3 on System and Infrastructure Security.

A) Industry Needs – Use Cases and Technology Outlook

Progress on Moore’s law in classic semiconductor technology and the technology improvement of quantum computers, provide even more computation power to the attacker. The attacker has cloud computing, massively internet connected PCs and even botnets of IoT devices to his disposal. Hence a new generation of algorithms is being proposed (see above in Theme 2) on the one end to support public key infrastructure and homomorphic encryption in the cloud, on the other end to support lightweight cryptography into wireless energy-starved IoT devices. Both, these new generations of post-quantum and lightweight algorithms still need efficient and secure implementations in classic semiconductor technologies. A cloud computing set-up typically requires extremely high throughputs while keeping the cooling requirements and thus power consumption under control. In a typical IoT sensor-actuator set-up e.g., in automotive or health-care, fast reaction time is essential. This requires implementations with low latency while combining this with small footprint and low energy consumption to keep the cost down and to preserve the battery of the IoT device.

Besides efficiency in terms of memory footprint, throughput, latency, power and/or energy, the implementations also have to be secure against physical side-channel and fault attacks. Sophisticated security certification and evaluation methods (FIPS, CC, etc.) have been established to give assurance about the security claims by independent evaluation and testing. The drawback is that certification is time consuming, expensive and sometimes the results are not repeatable. There is additionally an emerging need on one side for further developing provably secure protection methods and automated verification tools and on the other side improving the efficiency and quality of certification by integrating these tools and methods which will

allow assessment of the physical attacks resilience of the implementations without always the immediate necessity to be tested manually in certification labs.

For many applications such as conditional access, media protection and banking, hardware modules are being replaced by software solutions. The key challenge is then to protect cryptographic keys in an adequate way in software. White-box cryptography (WBC) is becoming the solution of choice to create such a protection. Despite all academic white-box implementations being broken, the deployment of WBC in industry keeps increasing.

B) State-of-the-art: Highlights

Multiple implementation aspects for cryptographic algorithms and protocols need to be considered. There are complete new classes of algorithms of which the implementation cost is unknown (RA 4.4.1). Together with the implementation cost also the resistance of the implementation against side-channel and fault attacks is required (RA 4.4.2). Special attention and a different approach is needed when software implementations have no hardware handles or secure storage for protection: in this situation white-box cryptography is required (RA 4.4.3).

(RA 4.4.1) Implementation Challenges of Post-quantum, FHE, Lightweight Crypto on Novel Compute Platforms

The post-quantum secure algorithms mentioned in research topic above (RA 4.2.2), as well as the cryptographic protocols need efficient implementations in existing and future hardware platforms. The computationally very demanding applications, such as fully or somewhat homomorphic encryption applications, will need to run in an efficient way on novel multi-core platforms, consisting of a heterogeneous mix of CPU's, GPU's and hardware co-processors.

There is a long tradition of secure and efficient implementation of public-key algorithms. Examples are extremely lightweight versions of elliptic curve that fit into the power budget of a passive RFID tag [Lee08]. The implementations also need to be resistant to a wide range of attacks [FAN10]. New generations of elliptic curve algorithms with different properties and higher security levels also need implementations [TURAN19].

The first proposals for post-quantum cryptographic algorithms resulted in designs for hardware building blocks for post-quantum computing. An example is the compact Ring Learning With Errors (RLWE) co-processor [SINHAROY14]. These solutions also require co-processor modules resistant to side-channel attacks. Constant time is essential [KARMAKAR18] and masking protects against EM and power attacks [REPARAZ16]. More recently, processors and co-processors to support homomorphic encryption have been introduced. While almost all publications focus on CPU or GPU implementations, our original focus is on hardware acceleration or co-processors with their own instruction set [SINHAROY17, SINHAROY18, SINHAROY19].

The design of lightweight crypto algorithms requires a close cooperation between cryptographers and experts on implementation. Lightweight is a broad term and could cover many different aspects: low area, low memory footprint, low power or low energy or low latency down to single cycle implementations [MAENE15,SIJACIC16].

Besides the algorithms, also the digital platforms change. While in the past, implementations were either focused on software implementations, limited by the instruction set at hand, or hardware implementations, which fixes the implementation, many more hybrid forms of computation appear. One is the appearance of heterogeneous multi-core platforms, with fine-grain and coarse grain computation engines, another is the appearance of much more distributed memory architectures. Moreover, with the integration of FPGAs into mainstream compute platforms, a hardware architecture is no longer 'hard' but can be reconditioned through partial reconfiguration even when deployed. This way, cryptographic agility can be enabled, i.e. the ability of an implementation to be updated after deployment, to thwart newly discovered theoretical or physical vulnerabilities, or to comply with new standards. In order to develop energy-efficient and agile platforms, embedded FPGA (eFPGA) platforms, tailored to cryptographic algorithms, are proposed [MENTENS18]. The features of these new platforms should be taken into account.

(RA 4.4.2) Side-Channel and Fault Attacks

The traditional application of cryptography is the protection of communication lines. It is usually assumed that both sender and receiver have equipment that is protected by physical means against attacks. In modern applications like payment cards, set-top boxes, DRM protection, sensor nodes etc., this assumption is no longer true. The attacker often has physical access to the device that is executing the cryptographic algorithm, and can measure side channels (execution time [KOCHER96], power consumption [KOCHER99], electro-magnetic radiation [GANDOLFI01]) or perform fault attacks (glitching [BAR-EL06], laser injection [SKOROBOGATOV03], electro-magnetic perturbation [QUISQUATER02]). With the advent of the IoT, the interest in embedded cryptographic systems and side-channel/fault attacks on these systems is steadily increasing, both in academia and industry.

Protection against side channel attacks (SCA) is usually done via masking [CHARI99], i.e. by randomizing any sensitive data manipulated during computations. A nice property of masking is that its security can be formally proven using abstract models that capture the leakage behavior of the underlying hardware [ISHAI03]. Yet if such models are inaccurate, the security guarantees do not hold. Examples of lacking features that have been so far studied in the literature include transition leakages in microprocessors (whose security degradation has been investigated in [BALASCH14]) or glitches in combinatorial logic (which has led to the appearance of Threshold Implementations [NIKOVA11]).

Protection against fault injection attacks (FA) is typically done in two ways: (1) checking whether the algorithm was faulted during the execution by using either area or time redundancy (e.g., duplication, concurrent error detection) or (2) using infection, i.e., ensuring that any induced fault results in a garbage output. The problem with duplication is that it does not provide security when faults are duplicated as well. Even with error-detecting codes, a powerful attacker can avoid detection if the injected faults result in valid code words. So far, all infective computations schemes have been broken.

The research direction of combined countermeasures - that is, countermeasures against both SCA and FA - is quite young and experimental. Examples of schemes that combine masking against SCA with redundancy against FA are ParTI [SCHNEIDER16] and Private Circuits II [IHSAI06], but these countermeasures naturally inherit the drawbacks of redundancy. Recently new countermeasures that combine masking against SCA and information-theoretic MAC tags against FA have been proposed (e.g., CAPA [REPARAZ18] and M&M [DEMEYER19]). In CAPA an actively secure multi-party computation protocol was adapted to the context of embedded systems in order to provide security against combined attacks. M&M is a new family of countermeasures that extends any SCA-secure masking scheme with information-theoretic MAC tags against Differential Fault Analysis (DFA) and combines them with an infective computation mechanism.

Formal verification of masked hardware implementations is achievable during several stages of the design flow: at algorithmic, at implementation or at physical level. In [REPARAZ14] a tool for verification at algorithmic level is presented. The tool receives a software implementation of the secured function and performs a leakage assessment over simulated traces. However, glitches are not considered. At implementation level, [BERTONI16] present a tool suitable for hardware implementations in the presence of glitches. Nevertheless, their tool only analyzes combinational logic with a simple power model. Bloem et al. [BLOEM18] also present a formal verification method for hardware implementations in the presence of glitches. Given a hardware implementation, they extract the netlist and model the logic gates using a Fourier representation. To simulate glitches, the model of a gate is extended to compute any Boolean function from its original inputs. Additionally, to analyze higher-orders, a SAT solver is instantiated. In both cases the analysis complexity increases enormously.

(RA 4.4.3) White-Box Cryptography

For many applications such as conditional access, media protection and banking hardware modules are being replaced by software solutions. The key challenge is then to protect cryptographic keys in an adequate way in software. White-box cryptography (WBC) is becoming the solution of choice to create such a protection. In spite of many powerful cryptanalytic results, the number of deployments keeps increasing

WBC started with the seminal white-box implementation of AES [CHOW02], a software implementation of AES designed to hide against attackers with full control over the implementation. This work inspired many other white-box implementations, but all of them have been broken so far [BILLET04, WYSEUR07, LEPOINT14].

Due to the lack of secure solutions in the literature, several companies developed white-box implementations but keep the designs secret to make the attacks harder to apply. Nevertheless, new design-agnostic attacks [BOS16] were proposed, which do not require knowledge of the design employed to secure the implementation. The effectiveness of these attacks was proved in the WhibOx Contest, a competition launched in 2017 where participants submitted white-box implementations of AES and broke other participants' submissions. All 97 submissions were broken.

Apart from key-extraction security, other security notions have been considered in the white-box model [Saxena09]. In particular, some authors have proposed new ciphers admitting huge implementations that cannot be compressed by an attacker without the knowledge of the key [BOGDANOV16]. These implementations make code-lifting harder by forcing attackers to send vast amounts of data through the network, but cannot be deployed in many practical scenarios.

C) Main Areas of Work

(RA 4.4.1) Implementation Challenges of Post-quantum, FHE, Lightweight Crypto on Novel Compute Platforms

We plan to work on building blocks to support post-quantum and homomorphic encryption schemes. We also plan to work on the implementation aspects of lightweight crypto. These implementations need to be efficient as well as resistant to side-channel and fault attacks. For post-quantum crypto we also really need to consider side- channel security and not only from an implementation point of view, but also at the algorithmic level. This includes also subroutines such as side-channel secure error correcting codes. On the one hand, these implementations and side-channel countermeasure need to be optimized for emerging computing platforms. On the other hand, computing platforms such as embedded FPGAs (eFPGAs) need to be tailored to cryptographic algorithms in order obtain energy-efficient and inherently side-channel resistant platforms.

(RA 4.4.2) Side-Channel and Fault Attacks

Countermeasures against SCA and FA build on abstract models that capture characteristics of the underlying hardware. For SCA typical assumptions are that side-channel leakages depend on the operations executed at individual clock cycles, or that a program execution strictly follows the compiled machine code. For FA it is assumed that injected errors can alter data flows (e.g., set/reset one or more bits in intermediate values) or program flows (e.g., skip/replace machine level instructions), up to a certain precision. While such approximations are generally accepted for simple computing devices, they have seldom been challenged for more complex platforms with advanced micro-architectural features. In this research topic, we aim to investigate to which extent components such as out-of-order engines, instruction-level parallelism, or large speculative instruction pipelines can degrade the security of state-of-the-art SCA and FA countermeasures. We plan to apply this insight to propose novel protection strategies that build on refined hardware models.

Formal security definitions and methods to defend software implementations against combined attacks will be developed. We will consider active adversaries who fault a threshold number of circuit wires and combine them with the usual probing adversaries to model combined attacks. Having formalized our adversaries, we will define combined secure implementations ensuring correct output and sensitive variable privacy. We will extend our models by first considering modular implementations and then studying composable security conditions.

Current state-of-the-art verification tools focus solely on verifying implementations protected against side-channel analysis. These tools range between very formal verification to more practical evaluation. More formal verification tools give a stronger conclusion about the security provided, but they are only applicable to small gadgets. Instead, practical evaluations need special equipment and implementations deployed to perform the evaluation. Some work has been done to extend formal verification methodologies to cope with more practical evaluations, nevertheless only a preliminary security assessment can be made. More robust formal tools are needed, capable of handling entire and practical implementations, as well as capable of giving

a thorough security assessment. We will develop robust verification tools for security evaluation of countermeasures against fault attacks and at the next stage we will extend these tools to combined attacks: side-channel and fault attacks. We will work towards defining metrics for combined security and developing procedures for their evaluation using verification tools to replace (at least partially) the attacks-based evaluation procedure which is currently used.

(RA 4.4.3) White-Box Cryptography

Our research will focus on developing the theory behind white-box cryptography, the security analysis of existing constructions and the development of new constructions. The latter will include both white-boxing constructions for existing ciphers as well as the construction of new ciphers that are specifically designed for white-box cryptography. The proposed solutions will also be thoroughly evaluated based on existing and novel attacks. The research programme will also study how white-box cryptography can be securely integrated with applications.

D) Expected Outcomes and Road Map

(RA 4.4.1) Implementation Challenges of Post-quantum, FHE, Lightweight Crypto on Novel Compute Platforms

- Multiplier architectures for lattice based crypto suitable for FPGA (Y1)
- SW routines of basic building blocks of lattice based crypto suited for small embedded micro-controllers (Y1)
- FPGA implementations of lightweight algorithms submitted to the NIST lightweight competition (Y1).
- Physical implementation and design flow for dedicated configurable platforms, tailored to symmetric-key cryptography (Y2).
- Inherent side-channel protection in dedicated configurable platforms (Y5).

(RA 4.4.2) Side-Channel and Fault Attacks

- Fault attack characterization of an ARM superscalar processor (Y1), security degradation of higher-order masking in the presence of advanced micro-architectural features (Y2), refinement of security models and development of suitable countermeasures (Y5).
- The CAPA methodology on different multiparty computation protocols with adapted security models is the first goal (Y1). New countermeasures against combined attacks (SC and Fault) which have improved performance and more realistic adversary model (Y5).
- Development of robust automated verification tools capable of handling entire and practical implementations (Y1). To define metrics for combined security and develop procedures for their evaluation using verification tools (Y5).

(RA 4.4.3) White-Box Cryptography

The first result (Y1) will be a security analysis of published white-box implementations, focusing on the weaknesses that make them vulnerable. The second result (Y2) will be the design of a white-box construction, such as an implementation of an existing cipher or a new dedicated cipher designed for the white-box model. For the long term (Y5) general models and constructions for white-box cryptography will be developed.

5.5. Connections with other Research Tracks

The results of this Research Track are directly useful in many use cases in Research Track 2 on security services and Research Track 3 on system and infrastructure security. Where relevant, close interactions have been planned in order to align the research efforts.

Theme 1 of Research Track 2 (Security Services) deals with identity management; the work on privacy preserving authentication will interact with the research on cryptographic protocols (Theme 3 of this Track). A substantial part of the work of Theme 3 of Track 2 deals with computing on secure data using advanced encryption techniques. The cryptographic algorithms and protocols necessary for this research are developed in Theme 2 and Theme 3 of this Track.

Theme 1 of Research Track 3 (System Security) relies on secure implementations of cryptographic algorithms and protocols, for example to protect confidentiality and integrity of code and data in processors and in part on hardware security; hence there is a strong synergy with the whole track. There is an obvious connection between hardware-based side channels (Theme 4 of this track) and architectural side channels on general-purpose hardware (Theme 1 of Track 2). While they share some techniques to obtain information, they operate at different abstraction levels; hence the required countermeasures are very different.

Theme 2 of Research Track 3 (Network Security) requires efficient cryptographic algorithms (Theme 2 of this track) to achieve high bandwidth or low latency, which are highly relevant to secure communications in IoT and industry 4.0. For cryptographic protocols, the work in Track 3 focuses on the evaluation of the security of cryptographic protocol implementations, while theme 3 of this Track focuses on the analysis of cryptographic protocols.

5.6. References

[ANDREEVA19] E. Andreeva, V. Lallemand, A. Purnal, R. Reyhanitabar, A. Roy, and Damian Vizár. 2019. ForkAE. submission to the NIST Symmetric Lightweight competition.

[ANDREEVA18] E. Andreeva, R. Reyhanitabar, K. Varici, Damian Vizár. 2018. Forking a Blockcipher for Authenticated Encryption of Very Short Messages. In Cryptology ePrint Archive. Report 2018/916.

[BACK02] A. Back, "Hashcash - a denial of service countermeasure,"http://www.hashcash.org/papers/hashcash.pdf, 2002.

[BALASCH11] J. Balasch, B. Gierlichs, I. Verbauwhede. 2011. An In-depth and Black-box Characterization of the Effects of Clock Glitches on 8-bit MCUs. In Fault Diagnosis and Tolerance in Cryptography – FDTC 2011, IEEE, pp. 105-114.

[BARAK06] B. Barak, R. Impagliazzo, A. Wigderson. 2006. Extracting Randomness Using Few Independent Sources. In SIAM J. Computing 36(4), pp. 1095-1118.

[BAR-EL06] H. Bar-El, H. Choukri, D. Naccache, M. Tunstall, C. Whelan. 2006. The Sorcerer's Apprentice Guide to Fault Attacks. In Proceedings of the IEEE 94(2), pp. 370-382.

[BARKER15] E. Barker, J. Kelsey. 2016. NIST SP 800-90a rev. 1 Recommendation for random number generation using deterministic random bit generators. Retrieved September, 3:2016, 2015.

[BEULLENS18] W. Beullens, S.R. Blackburn. 2018. Practical Attacks Against the Walnut Digital Signature Scheme. In Advances in Cryptology – ASIACRYPT 2018, LNCS 11272, Springer-Verlag, pp. 35-61.

[BEIERLE16] C. Beierle, J.Jean, S. Kölbl, G. Leander, A. Moradi, T. Peyrin, Y. Sasaki, P. Sasdrich, S.M. Sim. 2016. The SKINNY Family of Block Ciphers and Its Low-Latency Variant MANTIS. In Advances in Cryptology – CRYPTO (2) 2016, LNCS, Springer-Verlag, pp. 123-153.

[BERNSTEIN15] D.J. Bernstein, D. Hopwood, A. Hülsing, T. Lange, R. Niederhagen, L. Papachristodoulou, M. Schneider, P. Schwabe, Z. Wilcox-O'Hearn. 2015. SPHINCS: Practical Stateless Hash-Based Signatures. In Advances in Cryptology – EUROCRYPT (1) 2015, LNCS, Springer-Verlag, pp. 368-397.

[BELLARE94] M. Bellare, P. Rogaway. 1994. Entity Authentication and Key Distribution. In Advances in Cryptology - CRYPTO 1993, LNCS 773, Springer-Verlag, pp. 232-249.

[BERTONI16] G. Bertoni, M. Martinoli, M.C. Molteni. 2017. A methodology for the characterization of leakages in combinatorial logic. In Journal of Hardware and Systems Security 1(3), pp. 269-281.

[BETH90] Thomas Beth, Yvo Desmedt. 1990. Identification Tokens - or: Solving the Chess Grandmaster Problem. In Advances in Cryptology - CRYPTO '90, LNCS 537, Springer-Verlag, pp. 169-177.

[BILLET04] O. Billet, H. Gilbert, C. Ech-Chatbi. 2004. Cryptanalysis of a White Box AES Implementation. In Selected Areas in Cryptography - SAC 2004, LNCS 3357, Springer-Verlag, pp. 227-240.

[BLOEM18] R. Bloem, H. Gross, R. Iusupov, B. Knighofer, S. Mangard, J. Winter. 2018. Formal Verification of Masked Hardware Implementations in the Presence of Glitches. In Advances in Cryptology - EUROCRYPT 2018, LNCS 10821, Springer-Verlag, pp. 321-353.

[BOGDANOV16] A. Bogdanov, T. Isobe, E. Tischhauser. 2016. Towards Practical Whitebox Cryptography: Optimizing Efficiency and Space Hardness. In Advances in Cryptology – ASIACRYPT 2016, LNCS 10031, Springer-Verlag, pp. 126-158.

[BORGHOFF12] J. Borghoff, A. Canteaut, T. Güneysu, E. Bilge Kavun, M. Knezevic, L.R. Knudsen, G. Leander, V. Nikov, C. Paar, C. Rechberger, P. Rombouts, S.S. Thomsen, T. Yalçin. 2012. PRINCE - A Low-Latency Block Cipher for Pervasive Computing Applications. In Advances in Cryptology, ASIACRYPT 2012, LNCS, Springer-Verlag, pp. 208-225.

[BOS16] J. W. Bos, C. Hubain, W. Michiels, P. Teuwen. 2016. Differential Computation Analysis: Hiding Your White-Box Designs is Not Enough. In Cryptographic Hardware and Embedded Systems – CHES 2016, LNCS 9813, Springer-Verlag, pp. 215-236.

[BRANDS94] S. Brands, D. Chaum. 1994. Distance-Bounding Protocols. In Advances in Cryptology - EUROCRYPT '93, LNCS 765, Springer-Verlag, pp. 344–359.

[CAESAR19] CAESAR: Competition for Authenticated Encryption: Security, Applicability, and Robustness https://competitions.cr.yp.to/caesar.html – last retrieved 1 May, 2019.

[CANETTI01A] R. Canetti. 2001. Universally composable security: a new paradigm for cryptographic protocols. In IEEE Symposium on Foundations on Computer Science, pp. 136-145.

[CANETTI01B] R. Canetti, H. Krawczyk. 2001. Analysis of Key-Exchange Protocols and Their Use for Building Secure Channels. In Advances in Cryptology - EUROCRYPT 2001, LNCS 2045, Springer-Verlag, pp. 435-474.

[CRAMER18] R. Cramer, I. Damgård, D. Escudero, P. Scholl, C. Xing. 2018. SPD Z_{2^k}: Efficient MPC mod 2^k for Dishonest Majority, In Advances in Cryptology – CRYPTO (2) 2018, LNCS, Springer-Verlag, pp. 769-798.

[CHARI99] S. Chari, C. S. Jutla, J. R. Rao, P. Rohatgi. 1999. Towards Sound Approaches to Counteract Power-Analysis Attacks. In Advances in Cryptology - CRYPTO 1999, LNCS 1666, Springer-Verlag, pp. 398-412.

[CHAUM81] D. Chaum. 1981. Untraceable Electronic Mail, Return Addresses, and Digital Pseudonyms. In Commun. ACM 24(2): pp. 84-88

[CHOW02] S. Chow, P. Eisen, H. Johnson, P.C. Van Oorschot. 2003. White-Box Cryptography and an AES Implementation. In Selected Areas in Cryptography - SAC 2002, LNCS 2595, Springer-Verlag, pp. 250-270.

[CHUANG18:2] K. Chuang, E. Bury, R. Degraeve, B. Kaczer, D. Linten, I. Verbauwhede. 2018. A Physically Unclonable Function with 0% BER Using Soft Oxide Breakdown in 40nm CMOS. In Asian Solid-State Circuits Conference IEEE, 4 pages.

[CHOU15] T. Chou, C. Orlandi. 2015. The simplest protocol for oblivious transfer. In Progress in Cryptology - LATINCRYPT 2015 - 4th International Conference on Cryptology and Information Security in Latin America, LNCS, Springer-Verlag, pp. 40–58.

[CREMERS12] C. Cremers, M. Feltz. 2012. Beyond eCK: Perfect Forward Secrecy under Actor Compromise and Ephemeral-Key Reveal. In Computer Security - ESORICS 2012, LNCS 7459, Springer-Verlag, pp. 734-751.

[CROMAN16] K. Croman, C. Decker, I. Eyal, A.E. Gencer, A. Juels, A.E. Kosba, A. Miller, P. Saxena, E. Shi, E. Gün Sirer, D. Song, R. Wattenhofer. 2016. On Scaling Decentralized Blockchains - (A Position Paper). In Financial Cryptography Workshops 2016, LNCS, Springer-Verlag, pp. 106-125.

[CRYSTALS17] CRYSTALS-KYBER and CRYSTALS-Dilithium Teams, NIST PQCrypto Submission, 2017. Available at https://pq-crystals.org/ – last retrieved Wed 24 April, 2019.

[DAEMEN02] J. Daemen, V. Rijmen. 2002. The Design of Rijndael: AES – The Advanced Encryption Standard. Springer. ISBN 3-540-42580-2

[DAI98] W. Dai, "b-money," http://www.weidai.com/bmoney.txt, 1998.

[DEMEYER19] L. De Meyer, V. Arribas, S. Nikova, V. Nikov, V. Rijmen. 2019. M&M: Masks and Macs against Physical Attacks. In IACR Transactions of Cryptographic Hardware and Embedded Systems 2019(1), pp. 25-50.

[DIFFIE07] W. Diffie, S. Landau. 2007. Privacy on the Line: The Politics of Wiretapping and Encryption, Updated and Expanded Edition. The MIT Press.

[DIFFIE76] W. Diffie, M. Hellman. 1976. New Directions in Cryptography. In IEEE Transactions on Information Theory 22(6), pp. 644-654.

[DINGELDINE04] R. Dingledine, N. Mathewson, P. Syverson. 2004. Tor: the second-generation onion router. In Proceedings of the 13th conference on USENIX Security Symposium - Volume 13 (SSYM'04), Vol. 13. USENIX Association, Berkeley, CA, USA, pp. 303-320

[DOERNER18] J. Doerner, Y. Kondi, E. Lee, A. Shelat. 2018. Secure Two-party Threshold ECDSA from ECDSA Assumptions. In 2018 Symposium on Security and Privacy, IEEE Computer Society Press, pp. 980-997.

[DOBRAUNIG18] C. Dobraunig, M. Eichlseder, L. Grassi, V. Lallemand, G. Leander, E. List, F. Mendel, C. Rechberger. 2018. Rasta: A Cipher with Low ANDdepth and Few ANDs per Bit. In Advances in Cryptology – CRYPTO (1) 2018, LNCS, Springer-Verlag, pp. 662-692.

[DODIS04] Y. Dodis, A. Elbaz, R. Oliveira, R. Raz. 2004. Improved Randomness Extraction from Two Independent Sources. In APPROX-RANDOM 2004, pp. 334-344.

[DODIS13] Y. Dodis, D. Pointcheval, S. Ruhault, D. Vergniaud, D. Wichs. 2013. Security analysis of pseudo-random number generators with input:/dev/random is not robust. In ACM Conference on Computer and Communications Security (CCS 2013), pp. 647-658.

[DWORK03] C. Dwork, A. Goldberg, M. Naor. 2003. On memory-bound functions for fighting spam. In Advances in Cryptology – CRYPTO 2003, LNCS, Springer-Verlag, pp. 426-444.

[DWORK92] C. Dwork, M. Naor. 1992. Pricing via Processing or Combatting Junk Mail. In Advances in Cryptology – CRYPTO 1992, LNCS, Springer-Verlag, pp. 139-147.

[EASYCRYPT] EasyCrypt team. https://www.easycrypt.info/trac/ - last retrieved Fri 26 April, 2019.

[ETSI17] 3GPP TS 35.201 Specification of the 3GPP confidentiality and integrity algorithms, Version 3.2.0, 2002 (updated to v14.0.0 in 2017).

[ETSI19] https://www.etsi.org/security-algorithms-and-codes/etsi-algorithms?jjj=1557692656400 – last retrieved Wed 24 April, 2019.

[FRODOKEM17] FrodoKEM Team. NIST PQCrypto Submission, 2017. Available at https://frodokem.org/ - last retrieved Wed 24 April, 2019.

[GANDOLFI01] K. Gandolfi, C. Mourtel, F. Olivier. 2001. Electromagnetic Analysis: Concrete Results, In Cryptographic Hardware and Embedded Systems - CHES 2001, LNCS 2162, Springer-Verlag, pp. 251-261.

[GARAY15] J. Garay, A. Kiayias, N. Leonardos. 2015. The iBtcoin backbone protocol: Analysis and applications, In Advances in Cryptology – EUROCRYPT (2), LNCS, Springer-Verlag, pp. 281-310.

[GENNARO18] R. Gennaro, S. Goldfeder. 2018. Fast Multiparty Threshold ECDSA with Fast Trustless Setup. In ACM Conference on Computer and Communications Security (CCS 2018), ACM Press pp. 1179-1194.

[GRASSI16] L. Grassi, C. Rechberger, D. Rotaru, P. Scholl, N.P. Smart. 2016. MPC-Friendly Symmetric Key Primitives. In ACM Conference on Computer and Communications Security, pp. 430-443.

[HABER90] S. Haber, W.S. Stornetta. 1990. How to Time-Stamp a Digital Document. In Advances in Cryptology – CRYPTO 1990, LNCS, Springer-Verlag, pp. 437-455.

[HANCKE05] G.P. Hancke, M.G. Kuhn. 2005. An RFID Distance Bounding Protocol. In Proceedings of the 1st International Conference on Security and Privacy for Emerging Areas in Communications Networks (SECURECOMM '05), IEEE Computer Society, pp. 67–73.

[IETF18] Internet Engineering Task Force. The Transport Layer Security (TLS) Protocol Version 1.3, RFC 8446, 2018. Available at https://tools.ietf.org/html/rfc8446 – Last retrieved Wed 24 April 2019.

[ISHAI03] Y. Ishai, A. Sahai, D.A. Wagner. 2003. Private Circuits: Securing Hardware against Probing Attacks. In Advances in Cryptology - CRYPTO 2003, LNCS 2729, Springer-Verlag, pp. 463-481.

[ISHAI06] Y. Ishai, M. Prabhakaran, A. Sahai, D.A. Wagner. 2006. Private Circuits II: Keeping Secrets in Tamperable Circuits. In Advances in Cryptology - EUROCRYPT 2006, LNCS 4004, Springer-Verlag, pp. 308-327.

[ISHAI13] Y. Ishai, J. Kilian, K. Nissim, E. Petrank. 2003. Extending oblivious transfers efficiently. In Advances in Cryptology – CRYPTO 2003. LNCS 2729, Springer-Verlag, pp. 145–161.

[JOHNSON13] A. Johnson, C. Wacek, R. Jansen, M. Sherr, P. Syverson. 2013. Users get routed: traffic correlation on tor by realistic adversaries. In Proceedings of the 2013 ACM SIGSAC Conference on Computer & Communications Security (CCS '13). ACM, New York, NY, USA, pp. 337-348.

[KEANE10] J. Keane, X. Wang, D. Persaud, C.H. Kim. 2010. An All-In-One Silicon Odometer for Separately Monitoring HCI, BTI, and TDDB. In IEEE Journal of Solid-State Circuits, vol. 45, no. 4, pp. 817-829.

[KILLMANN11] W. Killmann, W. Schindler. 2011. A Proposal for: Functionality classes for random number generators. BSI, Bonn.

[KLLP16] M. Kaplan, G. Leurent, A. Leverrier, M. Naya-Plasencia. 2016. Breaking Symmetric Cryptosystems Using Quantum Period Finding. In Advances in Cryptology – CRYPTO (2) 2016, LNCS, Springer-Verlag, pp. 207-237.

[KRAITSBERG19] M. Kraitsberg, Y. Lindell, V. Osheter, N.P. Smart, Y. Talibi Alaoui. 2019. Adding Distributed Decryption and Key Generation to a Ring-LWE Based CCA Encryption Scheme. To appear at the 24th Australasian Conference on Information Security and Privacy, LNCS, Springer-Verlag.

[KOCHER96] P.C. Kocher. 1996. Timing Attacks on Implementations of Diffie-Hellman, RSA, DSS, and Other Systems. In Advances in Cryptology - CRYPTO 1996, LNCS 1109, Springer-Verlag, pp. 104-113.

[KOCHER99] P.C. Kocher, J. Jaffe, B. Jun, 1999. Differential Power Analysis. In Advances in Cryptology - CRYPTO 1999, LNCS 1666, Springer-Verlag, pp. 388-397.

[KELLER15] M. Keller, E. Orsini, P. Scholl. 2015. Actively Secure OT Extension with Optimal Overhead. In Advances in Cryptology – CRYPTO 2015, Springer-Verlag, pp. 724-741.

[KELLER16] M.Keller, E. Orsini, P. Scholl. 2016. MASCOT: Faster Malicious Arithmetic Secure Computation with Oblivious Transfer. In ACM Conference on Computer and Communications Security 2016, pp. 830-842.

[LINDELL18] Y.Lindell, A. Nof. 2018. Fast Secure Multiparty ECDSA with Practical Distributed Key Generation and Applications to Cryptocurrency Custody. In ACM Conference on Computer and Communications Security, ACM Press, pp. 1837-1854.

[LAI91] X. Lai, J.L. Massey. 1991. A proposal for a new block encryption standard. In Advances in Cryptology EUROCRYPT'90, Aarhus, Denmark, LNCS 473, Springer-Verlag, pp. 389–404.

[LUOV17] LUOV Team. NIST PQCrypto Submission, 2017. Available at https://www.esat.kuleuven.be/cosic/pqcrypto/luov/ - last retrieved Wed 24 April, 2019.

[MASSIAS99] H. Massias, X.S. Avila, J.-J. Quisquater. 1999. Design of a secure timestamping service with minimal trust requirements. In 20th Symposium on Information Theory in the Benelux.

[MATHEW16] S.K. Mathew, D. Johnston, S. Satpathy, V.B. Suresh, P. Newman, M. Anders, H. Kaul, A. Agarwal, S. Hsu, G.K. Chen, R. Krishnamurthy. 2016. µRNG: A 300–950 mV, 323 Gbps/W All-Digital Full-Entropy True Random Number Generator in 14 nm FinFET CMOS. in IEEE Journal of Solid-State Circuits, vol. 51, no. 7, pp. 1695-1704.

[MORRISSEY10] P. Morrissey, N.P. Smart, B.Warinschi. 2010. The TLS Handshake Protocol: A Modular Analysis. J. Cryptology 23(2), pp. 187-223.

[MOUHA14] N. Mouha, B. Mennink, A. Van Herrewege, D. Watanabe, B. Preneel, I. Verbauwhede: Chaskey. 2014. An Efficient MAC Algorithm for 32-bit Micro-controllers. In Selected Areas in Cryptography 2014, LNCS, Springer-Verlag, pp. 306-323.

[NAKAMOTO08] S. Nakamoto. 2008. Bitcoin: A peer-to-peer electronic cash system. Available at http://www.bitcoin.org/bitcoin.pdf

[NIKOVA11] S. Nikova, V. Rijmen, M. Schläffer. 2011. Secure Hardware Implementation of Nonlinear Functions in the Presence of Glitches. In Journal of Cryptology 24(2), pp. 292-321.

[NIST16] National Institute of Standards and Technology. On-line communication, 2016. Available at https://csrc.nist.gov/News/2016/Public-Key-Post-Quantum-Cryptographic-Algorithms – last retrieved Wed 24 April, 2019. Project homepage available at https://csrc.nist.gov/projects/post-quantum-cryptography – last retrieved Wed 24 April, 2019.

[NIST17] NIST Interagency Report 8114, Report on Lightweight Cryptography, https://nvlpubs.nist.gov/nistpubs/ir/2017/NIST.IR.8114.pdf – last retrieved Wed 24 April, 2019.

[NIST19] National Institute of Standards and Technology. On-line publication, 2019. Available at https://www.nist.gov/publications/threshold-schemes-cryptographic-primitives – last retrieved Wed 24 April, 2019.

[NIST19L] NIST, https://csrc.nist.gov/projects/lightweight-cryptography – last retrieved 1 May, 2019.

[NIST19P] NIST, https://csrc.nist.gov/Projects/Post-Quantum-Cryptography– last retrieved 1 May, 2019.

[NIST19T] NIST, https://csrc.nist.gov/Projects/Threshold-Cryptography – last retrieved 1 May, 2019.

[NIST94] NIST. Federal Information Processing Standards Publication 186 (FIPS 186), 1994. Latest version available at https://csrc.nist.gov/publications/detail/fips/186/4/final – last retrieved Wed 24 April, 2019.

[NIELSEN12] J.B. Nielsen, P.S. Nordholt, C Orlandi, S.S. Burra. 2012. A new approach to practical active-secure two-party computation. In Advances in Cryptology–CRYPTO 2012, LNCS, Springer-Verlag, pp. 681-700.

[OVERDORF17] R. Overdorf, M. Juarez, G. Acar, R. Greenstadt, C. Diaz. 2017. How Unique is Your onion?: An Analysis of the Fingerprintability of Tor Onion Services. In Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security (CCS '17). ACM, New York, NY, USA, pp. 2021-2036.

[PIOTROWSKA17] A.M. Piotrowska, J. Hayes, T. Elahi, S. Meiser, G. Danezis. 2017. The loopix anonymity system. In Proceedings of the 26th USENIX Conference on Security Symposium (SEC'17), USENIX Association, Berkeley, CA, USA, pp. 1199-1216.

[PROVERIF] ProVerif team. https://prosecco.gforge.inria.fr/personal/bblanche/proverif/ - last retreived Fri 26 April, 2019.

[PEIKERT08] Chris Peikert, Vinod Vaikuntanathan, Brent Waters. 2018. A Framework for Efficient and Composable Oblivious Transfer. In Advances in Cryptology – CRYPTO 2018, LNCS, Springer-Verlag, pp. 554-571.

[QUISQUATER02] J.-J. Quisquater, D. Samyde. 2002. Eddy Current for Magnetic Analysis with Active Sensor. In E-Smart 2002, pp. 185-194.

[RABIN81] M.O. Rabin, M.O. 1981. How to exchange secrets with oblivious transfer. Harvard University Technical report 81.

[RANGANATHAN12] A. Ranganathan, N.O. Tippenhauer, B. Skoric, D. Singelée, S. Capkun. 2012. Design and Implementation of a Terrorist Fraud Resilient Distance Bounding System. In Proceedings of ESORICS 2012, LNCS 7459, Springer-Verlag, pp. 415-432.

[RASMUSSEN10] K. Rasmussen, S. Capkun. 2010. Realization of RF Distance Bounding. In Proceedings of the 19th USENIX Security Symposium, pp. 389-402.

[RECHBERGER18] C. Rechberger, H. Soleimany, T. Tiessen. 2018. Cryptanalysis of Low-Data Instances of Full LowMCv2. In IACR Trans. Symmetric Cryptol. 2018(3): pp. 163-181.

[REPARAZ14] O. Reparaz. 2016. Detecting flawed masking schemes with leakage detection tests. in Fast Software Encryption - FSE 2016, LNCS 9783, Springer-Verlag, pp. 204–222.

[REPARAZ18] O. Reparaz, L. De Meyer, B. Bilgin, V. Arribas, S. Nikova, V. Nikov, N.P. Smart. 2018. CAPA: The Spirit of Beaver Against Physical Attacks. In Advances in Cryptology - CRYPTO 2018, LNCS 10991, Springer-Verlag, pp. 121-151.

[RIVEST78] R. L. Rivest, A. Shamir, L. Adleman. 1978. A method for obtaining digital signatures and public-key cryptosystems. In Communications of the ACM 21(2), pp. 120-126.

[ROTARU17] D. Rotaru, N.P. Smart, M. Stam. 2017. Modes of Operation Suitable for Computing on Encrypted Data. In IACR Trans. Symmetric Cryptol. 2017(3), pp. 294-324.

[ROZIC18] V. Rozic, I. Verbauwhede. 2018. Hardware-Efficient Post-processing Architectures for True Random Number Generators. In IEEE Transactions on Circuits and Systems II: Express Briefs 1(1), pp. 1-1.

[SABER17] SABER Team. NIST PQCrypto Submission, 2017. Available at https://www.esat.kuleuven.be/cosic/pqcrypto/saber/ - last retrieved Wed 24 April, 2019.

[SCHNEIDER16] T. Schneider, A. Moradi, T. Güneysu. 2016. ParTI: Towards Combined Hardware Countermeasures against Side-Channel and Fault-Injection Attacks. In Advances in Cryptology – CRYPTO 2006, LNCS 9815, Springer-Verlag, pp. 302-332.

[SHIRAZI19] F. Shirazi, M. Simeonovski, M.R. Asghar, M. Backes, C. Diaz. 2018. A Survey on Routing in Anonymous Communication Protocols. In ACM Comput. Surv. 51, 3, Article 51 (June 2018), 39 pages.

[SHUMOW07] D. Shumow, N. Ferguson. 2007. On the possibility of a back door in the NIST SP800-90 dual EC PRNG. In CRYPTO rump session.

[SIKE17] SIKE Team. NIST PQCrypto Submission, 2017. Available at https://sike.org/ - last retrieved Wed 24 April, 2019.

[SINGH19] M. Singh, P. Leu, S. Capkun. 2019 (to appear). UWB with Pulse Reordering: Securing Ranging against Relay and Physical Layer Attacks. In Proceedings of the Network and Distributed System Security Symposium (NDSS).

[SKOROBOGATOV03] S. Skorobogatov. 2003. Optical Fault Induction Attacks. In Cryptographic Hardware and Embedded Systems - CHES 2002, LNCS 2523, Springer-Verlag, pp. 2-12.

[TIPPENHAUER15] N.O. Tippenhauer, H. Luecken, M. Kuhn, S. Capkun. 2015. UWB rapid-bit-exchange system for distance bounding. In Proceedings of the 8th ACM Conference on Security & Privacy in Wireless and Mobile Networks (WiSec '15), pp. 2-12.

[TURAN18] M.S. Turan, E. Barker, J. Kelsey, K. McKay, M. Baish, M. Boyle. 2018. Recommendation for the Entropy Sources Used for Random Bit Generation. NIST Special Publication 800-90B.

[WANG14] X. Wang, J. Keane, T.T. Kim, P. Jain, Q. Tang, C.H. Kim. 2014. Silicon Odometers: Compact In Situ Aging Sensors for Robust System Design. in IEEE Micro, vol. 34, no. 6, pp. 74-85.

[WOOD14] G. Wood. 2014. Ethereum: A secure decentralised generalised transaction ledger. Ethereum project yellow paper 151 (2014), pp. 1-32.

[WOODAGE19] J. Woodage, D. Shumow. 2019. An Analysis of the NIST SP800-90A Standard. In EUROCRYPT 2019. See also IACR Cryptology ePrint Archive 2018: 349.

[WANG17] X. Wang, S. Ranellucci, J. Katz. 2017. Global-scale secure multiparty computation. In ACM Concnference on Computer and Communications Security, pp. 39–56.

[WU18] M.-Y. Wu, T.-H. Yang, L.C. Chen, C.-C. Lin, H.-C. Hu, F.-Y. Su, C.-M. Wang, J. P.-H. Huang, H.-M. Chen, C. C.-H. Lu, E. C.-S. Yang, R. S.-J. Shen. 2018. A PUF scheme using competing oxide rupture with bit error rate approaching zero. In 2018 IEEE International Solid - State Circuits Conference - (ISSCC), pp. 130–132.

[ZHANG19] R. Zhang, B. Preneel. 2019. Lay Down the Common Metrics: Evaluating Proof-of-Work Consensus Protocols' Security. In IEEE Security & Privacy 2019, to appear.

6. Prototypes, Validation and Inroads to Industry Implementation

The socio-economic importance of the cybersecurity technologies investigated in the proposed research programme is crucial. It is expected that targeted technologies, methods and solutions to cybersecurity problems are generically applicable to most (if not all) industry sectors. This chapter further addresses the relevance of and demand for cybersecurity solutions. The first section sketches three representative business domains to illustrate how the proposed programme addresses the broad demand for cybersecurity. The aim is to show the potential value of research results in a realistic business context. These illustrations do not intend to limit the application scope of security technologies to a specific industry sector.

Section 6.1 illustrates market relevance and highlights three example industry sectors with an outspoken demand for cybersecurity. Section 6.2 highlights three technology settings that are relevant for many industries: IoT platforms, cloud platforms and data sharing platforms. Section 6.3 illustrates how the Consortium will combine multiple cybersecurity research results to enhance the cybersecurity posture of the technology platforms mentioned above. Finally, Section 6.4 summarizes the approach to transfer cybersecurity solutions to industry implementations.

6.1. Illustration of Market Relevance:Three Strategic Industry Sectors for Cybersecurity

This section further illustrates the socio-economic relevance of cybersecurity research by highlighting three industry sectors and illustrating their demand for cybersecurity: healthcare, financial services and manufacturing (including industry4.0). We refer to market studies from various parties, without explicitly using the most recent reports.

This analysis also builds upon insights that were formulated by regional domain experts[5] in healthcare (e.g., from Barco, Televic, UZ Leuven), finance services (e.g., from Atos Worldline, Bancontact Payconiq Company, Euroclear, SWIFT), and manufacturing (e.g., from Atlas Copco, Siemens, Dematic). These insights were confirmed by recent market studies, for instance from KPMG, McKinsey, PwC and Wipro.

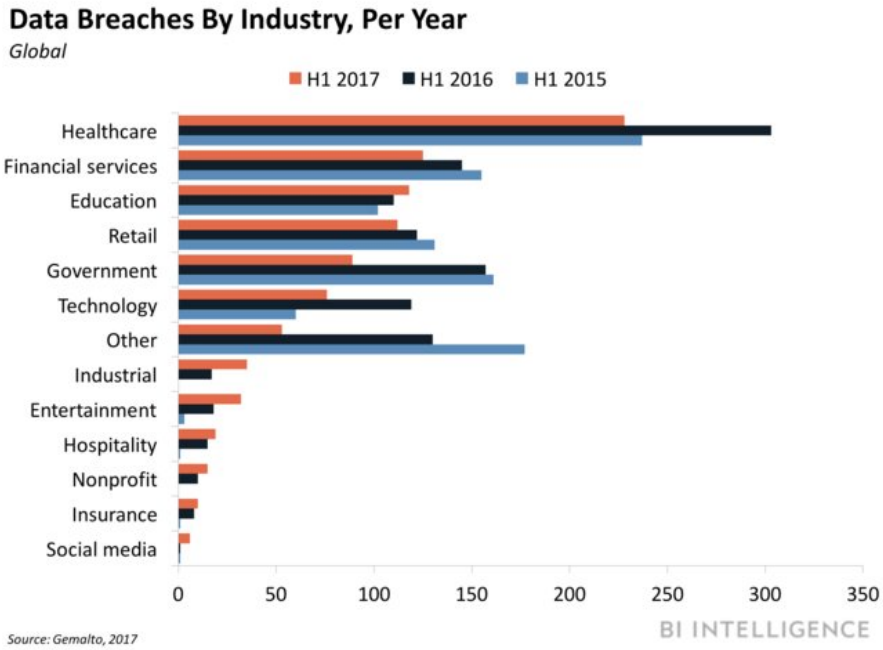

Figure 6-1: BI Intelligence Market Survey Identifies Healthcare and Financial Services as the Most Vulnerable Market Sectors for Cyberattacks.

[5] Most of these companies participated in the industry feedback session organized by the Cybersecurity Consortium on April 2, 2019. The aim of this full-day event was to engage industry in refining the focus and scope of the Cybersecurity Research Programme and in setting priorities that support their operations and strategic product/service roadmaps.

Healthcare

According to a 2017 market study by Wipro[6], 41% of all security breaches targeted the healthcare industry, followed by banking & financial services with 18%. This is consistent with insights from Business Insider[7] (see Figure 6.1) and Verizon. Verizon's 2018 Data Breach Investigations Report[8] shows that the healthcare sector experiences five times the number of breaches experienced by any other industry.

Forbes reports that in 2017, a typical healthcare organization experienced on average 32,000 intrusion attacks per day[9]. As a partial explanation and according to KPMG, medical data is worth at least 10 times as much as financial data on the dark web black market[10].

The impact of data breaches is massive. Over the last five years, Forbes[11] identified a strong increase of attacks in the healthcare industry, with the largest breaches impacting as many as 80 million people. In July 2018, it was revealed that data of 150,000 NHS patients in the UK was shared over a three-year period following a major breach. In the US, the 2015 cyber-attack on Anthem saw hackers steal 78.8 million patient records, claiming highly sensitive personal data. In 2018, hackers breached the Singapore government's health database with a targeted cyber-attack, accessing the data of 1.5 million patients.

Financial services

In their 2018 Top Financial Services Issues Report[12] PwC confirms: "cyberattacks against financial services and other sectors have grown in number, size, and sophistication. Fraud incidents, both online and offline, increased by more than 130% during the past year, resulting in significant monetary and reputational losses for financial institutions." The World Economic Forum released a white paper[13] in 2018 in which they stress that the "financial services system faces challenges, both internal and external, in managing innovation-driven cyber-risk. Internally, challenges around technology and expertise; externally, challenges around coordination with regulators and across the industry." Microsoft Asia, supported by a Frost & Sullivan study commissioned by them, reports[14] that "despite financial services being a highly regulated industry, more than half (56%) of

---

[6] Wipro. Wipro's State of Cybersecurity Report 2018, https://www.wipro.com/content/dam/nexus/en/service-lines/applications/latest-thinking/state-of-cybersecurity-report-2018.pdf

[7] Business Insider, BI Intelligence Report: The strategies companies are using to protect their customers - and themselves - in the age of massive breaches, August 30, 2018, https://www.businessinsider.com.au/intelligence-data-breaches-australia-2018-8, Visited April 18, 2019.

[8] Verizon, Data Breach Investigations Report 2018 (Executive Summary), https://enterprise.verizon.com/resources/reports/dbir, Visited April 9, 2019.

[9] Gary Alterson, Forbes Insights, "Confronting One Of Healthcare's Biggest Challenges: Cyber Risk", February 2019, https://www.forbes.com/sites/insights-intelai/2019/02/11/confronting-one-of-healthcares-biggest-challenges-cyber-risk/#67eb68427b83, visited April 9, 2019.

[10] KPMG, The healthy approach to cybersecurity, 2017, https://advisory.kpmg.us/content/dam/advisory/en/pdfs/cyber-report-healthcare.pdf, Visited April 9, 2019.

[11] Kate O'Flaherty, Forbes Insights, "Why Cyber-Criminals Are Attacking Healthcare -- And How To Stop Them", October 2018, https://www.forbes.com/sites/kateoflahertyuk/2018/10/05/why-cyber-criminals-are-attacking-healthcare-and-how-to-stop-them/#555e22b57f69, visited April 9, 2019.

[12] PWC, Top financial services issues of 2018, https://www.pwc.com/us/en/financial-services/research-institute/assets/pwc-fsi-top-issues-2018.pdf, Visited April 18, 2019.

[13] Oliver Wyman, Innovation-Driven Cyber-Risk to Customer Data in Financial Services, http://www3.weforum.org/docs/WEF_Cyber_Risk_to_Customer_Data.pdf, White paper World Economic Forum, March 2018, Visited April 18, 2019.

[14] Microsoft Asia. Fear of cyberattacks slows down the progress of digital transformation in financial services companies in Asia Pacific, https://news.microsoft.com/apac/2018/11/15/fear-of-cyberattacks-slows-down-the-progress-of-digital-transformation-in-financial-services-companies-in-asia-pacific/, Visited April 10, 2019.

the organizations surveyed have either experienced a security incident (27%) or are not sure if they have had a security incident as they have not checked (29%)."

According to the 2019 IBM X-Force Cybersecurity Intelligence Index[15], the impact of financial malware is immense (e.g., banking malware families such as TrickBot, Gozi, Ramnit, or IcedID). According to the same report, criminals are increasingly leveraging coin-mining malware over ransomware, installing miners on victim endpoints and enslaving them, thus slowly generating coins for the attacker. Spreading to every part of the globe, financially motivated threat actors in Eastern Europe and North Korea have taken special notice of the profitability of coin-mining malware since consumers in these regions have adopted the use of cryptocurrency as a regular payment method for everyday transactions.

Manufacturing

The manufacturing sector represents a strategic sector for cybersecurity that is experiencing a drastic transition towards digitization (also referred to as Industry 4.0). This transition is enabled by emerging technologies such as the Internet-of-Things (IoT) and the Cyber-Physical Systems[16] (CPS) connected to it.

The business potential associated with this transition is substantial. Bosch[17], for instance, estimates the overall consolidated market potential over € 500B by 2022. General Electric[18] envisioned in 2012 that over the next 20 years "the Industrial Internet could add a sizable $ 10-15 trillion to global GDP – the size of today's U.S. economy – over the same horizon." Cisco[19] estimates that a potential of $14.4 trillion in value is at stake arising from the combination of increased revenues and reduced costs in the period from 2013 to 2022.

In addition to this massive business potential, the ongoing digitization in the manufacturing sector implies considerable cybersecurity risks. The enabling IoT/CPS infrastructure creates a new and extended attack surface by using wireless sensor networks, for instance, to monitor and control in-factory machines and robots. According to the 2019 IBM X-Force Cybersecurity Intelligence Index[20], the number of IoT vulnerabilities recorded in 2018 increased with 5,400 percent over the number reported just five years earlier. Think about the Mirai botnet in 2016 (which caused internet-wide disruption) and its successors Aidra, Wifatch and Gafgyt, or newcomers such as the BCMUPnP_, Hunter52 and Torii53 botnets, which have amassed access to hundreds of thousands of devices to spread their Distributed Denial of Service (DDoS) attack malware, coin-mining malware and spam. A multinational expert survey on IoT security by McKinsey[21] indicates that of the 400 IoT-involved experts surveyed, 75% say that IoT security is either important or very important, 70% expect that its relevance will increase, but only 16 percent say their company is well prepared for the challenge.

---

[15] IBM X-Force Threat Intelligence Index 2019. https://www.ibm.com/downloads/cas/ZGB3ERYD. Visited April 10, 2019

[16] The Internet-of-Things (IoT) is the network of physical objects that contain embedded technology to communicate and sense or interact with their internal states or the external environment. Cyber-Physical Systems (CPS) are embedded intelligent ICT systems that provide the interface between ICT processes and the physical world to make products smarter, more interconnected, interdependent, collaborative, and autonomous. If we refer to IoT in this document, we always refer to the combination of both IoT (the distributed platform) and CPS (the individual systems connected to it).

[17] http://blog.bosch-si.com/wp-content/uploads/20140403_Infographic_Key-Markets_72dpi_992x709px_02.png

[18] P. Evans, M. Annunziata, Industrial Internet: Pushing the Boundaries of Minds and Machines, General Electric, Nov 2012.

[19] J. Bradley, J. Barbier, and D. Handler, Embracing the Internet of Everything to capture your share of $14.4 trillion, Cisco white paper, 2013.

[20] IBM X-Force Threat Intelligence Index 2019. https://www.ibm.com/downloads/cas/ZGB3ERYD. Visited April 10, 2019.

[21] The Internet of Things: How to capture the value of IoT, McKinsey, https://www.mckinsey.com/featured-insights/internet-of-things/our-insights/the-internet-of-things-how-to-capture-the-value-of-iot, Visited April 9, 2019.

6.2. Essential Types of Technology Platforms

This section identifies three technology perspectives, representing technology platforms that cover the end-to-end architecture of a contemporary distributed application. This inventory includes IoT platforms (enabling the ubiquitous presence of ICT systems and services), cloud and web platforms (enabling the outsourced deployment and delivery of many critical services), and data sharing platforms (enabling virtually unlimited information sharing and the building up of intelligence in advanced services).

- **IoT/CPS** system technologies, i.e. the physical device and low-level system software that create the front-end for an end user or that integrate with the environment (e.g., a sensor, a camera, a smartphone, an industrial control system, or a smart card).
- **Cloud and web** technologies, i.e. the back-end server infrastructure and middleware on the Web/Internet that enable the computation, visualization and storage of services and applications (e.g., an online IaaS, PaaS or SaaS service to host an online service, for instance for online content management, document sharing, or electronic payments).
- **Data** sharing technologies, i.e., to share the business events and data samples that are collected and/or provided in digital applications (e.g., medical records, bank transactions, images or video streams, or any sensor data collected from a robot or machine).

The overall idea is to build prototypes of cybersecurity solutions that are relevant to any (or a combination) of the platform/technology types listed above. These relatively isolated security techniques will be combined in larger platforms and lead the way towards prototypes and try-outs of more sophisticated cybersecurity solutions.

Such an approach will allow to zoom in on specific security challenges that were also raised[22] by system integrators and technology providers specialized in IoT/CPS, in cloud solutions and/or data management (e.g., Atlas Copco, Cisco, Cegeka, Dimension Data, GuardSquare, Hexagon, IBM, Microsoft, Nallian, Newtec, Nokia, NXP, OneSpan, Siemens, Proximus, Telenet, Zion Security).

6.3. Leveraging on Research Results: Illustration with Three Business Cases

As presented in detail in the previous sections, the cybersecurity research programme presents four synergetic Research Tracks that cover the whole hardware/software stack, from the root of trust hardware (Research Track 4), over the system and network infrastructure (Research Track 3) and security services (Research Track 2), to the application software level (Research Track 1).

In order to illustrate the coherence of these research activities, we describe three example cases and sketch how they can benefit from combining cybersecurity research results. Each of the cases is dominated by one of the types of platforms described above in Section 6.2 (i.e. IoT, cloud or data sharing) and relates to one of the high-potential sectors identified in Section 6.1 (i.e. manufacturing, healthcare and financial services).

A) Industrial Control Systems – Leveraging on IoT

Problem context for CPS/IoT infrastructure providers. The production process of, amongst others, medical, pharma, chemical and food products is strictly regulated and controlled by third-party quality assurance procedures. Infrastructure providers must enable external audits that assess the quality and correctness of the production process, which requires full traceability and genuineness of the production data. Contemporary industrial control systems rely on production data produced by Cyber-Physical Systems (CPS) on the factory floor (e.g., machines, robots, sensors or cameras) and collected via a (wireless) network, the Internet-of-Things (IoT).

[22] The Consortium has a strong tradition in industry collaboration for more than 20 years. Some members have established a solid network of industry relations in Flanders, Europe and beyond. Specific feedback on the Cybersecurity Programme has been collected recently on April 2, 2019.

This CPS/IoT infrastructure is highly business-critical for two reasons. First, the production data provides valuable insights in the performance and quality of the production process and must not leak to a competitor or malicious third party. Secondly, the CPS/IoT infrastructure must be available at all times to control the operations on the factory floor and avoid massive costs due to unexpected downtime.

No surprise that industrial control systems have been target of sophisticated malware attacks (e.g., the infamous Stuxnet worm). Malware that incorporates and exploits the physical entanglement of systems (e.g., GPS spoofing) poses severe novel threats.

Demands for cybersecurity. Achieving traceability (and thus guaranteeing the integrity of collected monitoring data) requires strict protection of the monitoring data, as well as the cyber-physical systems and the network that connect them. The collected data must be secured, for instance by providing robust encryption algorithms and services for identity and authentication management. The cyber-physical systems that produce the data must be secured, for instance, by strictly isolating independent software modules running on a sensor, camera or PLC controller, or by proving that the code that is running on a CPS is still genuine and not infected by malware. The integrity of the wireless IoT network must be guaranteed, for instance by securely commissioning network devices and intelligently scanning the network for malicious intruders.

Existing security solutions are typically not directly applicable to CPS/IoT due to additional requirements regarding, for instance, resource-constraints, real-time operation and safety. Combining these requirements is far from trivial, especially when processing, memory, and energy resources are (extremely) limited. Think, for instance, on achieving safe and real-time protection on a battery-driven IETF Class-1 device with an 8-bit microcontroller running at 10MHz with 128KB flash and 16KB RAM.

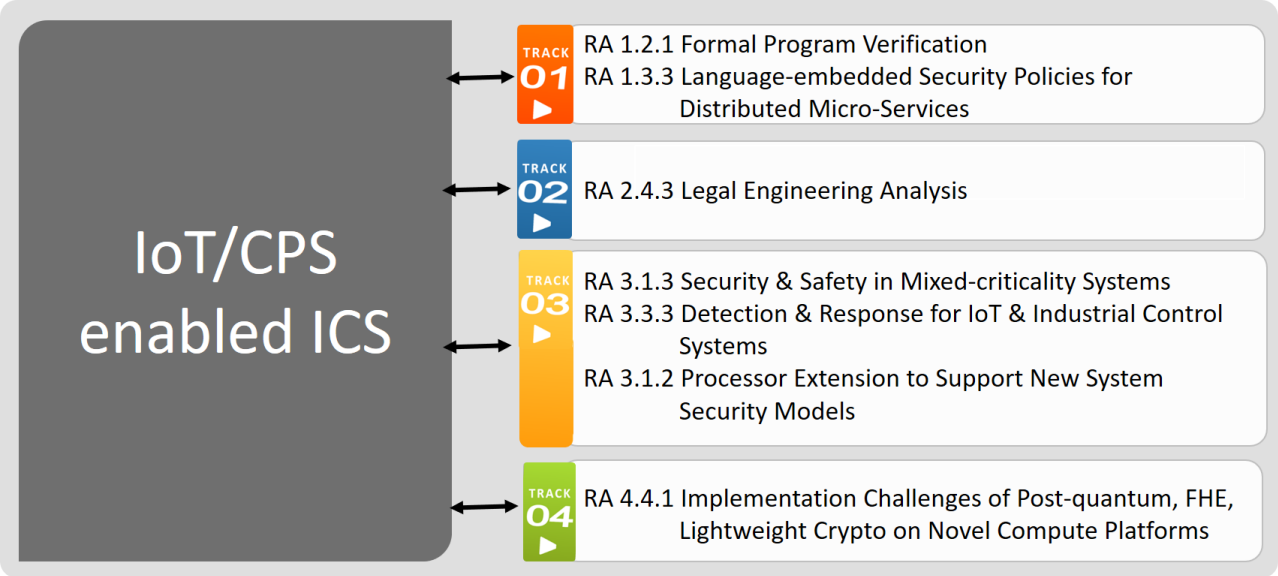

Figure 6-2: Relevant Cybersecurity Solutions for IoT/CPS-Enabled Industrial Control Systems.

Cybersecurity research addressing the demand. The cybersecurity Programme will investigate, design and develop mechanisms for securing legacy safety-critical CPS, such as industrial control systems, medical devices or autonomous vehicles, while complying with, amongst others, safety, cost and real-time requirements. These security mechanisms include hard- and software primitives, language and programming support with corresponding tools, as well as deployment and configuration middleware, with corresponding policy and configuration languages and tools.

As illustrated in Figure 6.2, developers of networked embedded IoT/CPS applications will benefit from development, verification, and language support in Track 1, such as *(RA 1.3.3) Language-embedded Security Policies for Distributed Micro-services*, and *(RA 1.2.1) Formal Program Verification*. Track 2 will investigate *RA 2.4.3 Legal Engineering Analysis* to embed legal values into the design of IoT/CPS systems (e.g., GDPR's data-protection-by-design and security-by-design). At systems and network infrastructure level, Track 3 will investigate *RA 3.1.3 Security & Safety in Mixed-criticality Systems*, and design solutions for *(RA 3.3.3) Detection*

& Response for IoT & Industrial Control Systems, and *(RA 3.1.2) Processor Extension to Support New System Security Models*. All these security solutions build upon the root of trust that is established by investigating, amongst others and *(RA 4.4.1) Implementation Challenges of Novel Algorithms & Platforms* for lightweight authentication and secure channels, and the complex trade-offs between time, on-chip area, power and energy.

B) Remote Health Monitoring – Using Data Sharing

Problem context for data centric application providers. Imagine a remote health monitoring application that observes a person's health condition by collecting data from, amongst others heart rate, blood pressure, sugar level, and oxygen saturation sensors on the body. The Philips Jovia Coach, for example, is a smartphone app that combines IoT, data and cloud technologies with human coaching to support people at-risk of type-2 diabetes. Data analytics tools provide new opportunities to guide users towards a healthier lifestyle.

These health-monitoring applications collect process and store personal and often highly sensitive data. All (meta)data is added to the person's health record, which is stored in a hospital data center and (partly) accessible for multiple medical professionals who may be associated with different care institutions. Medical data sharing is an example of a data-centric application area where challenging security, performance and transparency tradeoffs come into play and where design-level security approaches are essential, as any later changes require an expensive and strenuous redesign.

Demands for cybersecurity. As formulated, for example, in a 2014 study[23] by the SANS Institute, "the trend of pushing sensitive data outside an organization's *protected* environment via cloud computing, mobile access and the Internet of [Care] Things demands that security solutions are pushed closer to the actual data sources". Achieving privacy-aware sharing of medical data requires strict access control as well as transparency on the usage of data ("Which data is used by which applications?"). This demand for control and transparency is not only driven by legal obligations (e.g., the GDPR, which forces application providers to design and enforce strict cybersecurity support throughout the lifetime of the application and which defines substantial penalties when state-of-the-art protection measures are not in place).

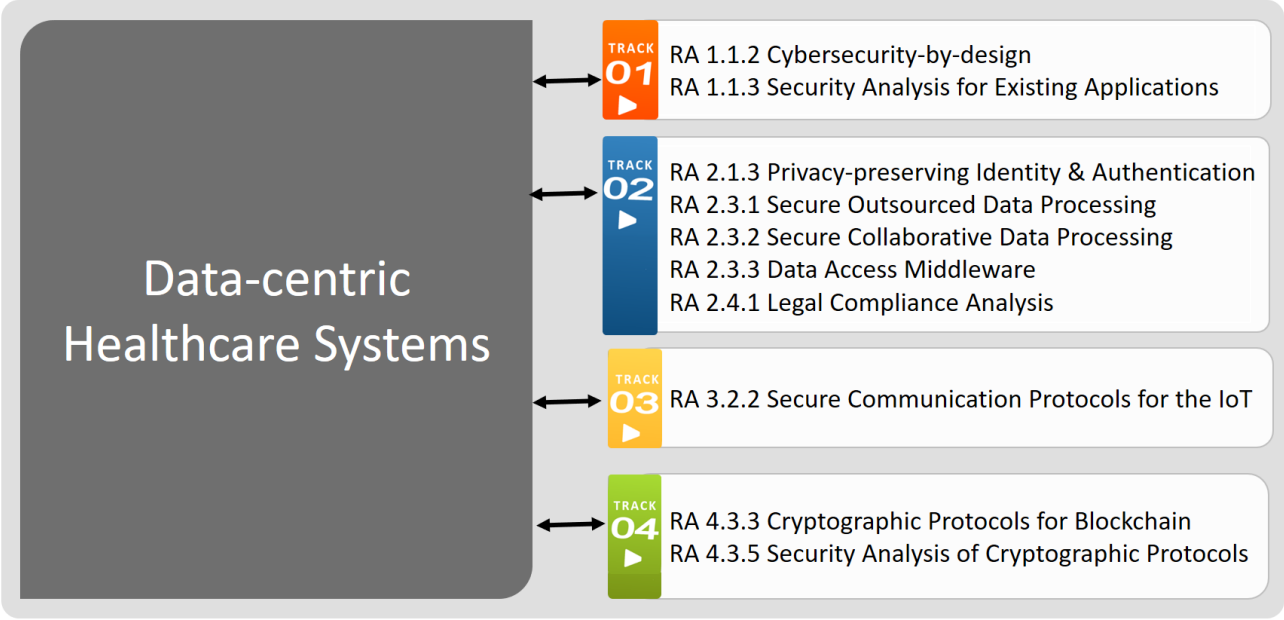

Figure 6-3: Relevant security solutions for data-centric e-healthcare systems.

Various types of data platforms demand for tailored cybersecurity support. Many enterprise solutions build on traditional database technologies (typically SQL-based), while an emerging series of new applications is being developed and deployed in a NoSQL context (be it column-based, based key value-pairs, document-

[23] Barbara Filkins, A SANS Survey – New threats drive improved practices: state of cybersecurity in health care organizations, December 2014, https://www.qualys.com/docs/sans-threats-drive-improved-practices-state-of-cybersecurity-health-care-organizations.pdf, (visited April 9, 2019).

based technologies, to name a few). In other words, while data centric applications emerge, one also observes a growing diversity in underpinning storage technologies. In addition, this wide variety of storage platforms and technologies is deployed within enterprises, by cloud providers, or in a combination (multi-cloud and hybrid cloud being the most prominent settings). Needless to state that the challenge of data protection becomes more complex as the specific architectures evolve. Furthermore, recent evolutions stress the importance and relative appreciation of peer-to-peer, fully decentralized architectures in specific applications. This is typically the case for distributed ledgers and blockchain.

Cybersecurity research to address industry demand. The Cybersecurity Research Program will address the cybersecurity challenges for data-centric applications starting from the early phases of the software engineering process, for instance via next-generation *(RA 1.1.2) Cybersecurity-by-Design Approaches* and *(RA 1.1.3) Security Analysis for Existing Applications* (see also Figure 6.3). These analysis and design approaches will be complemented with essential security services for data-centric application development, such as T2/Privacy-preserving Identity & Authentication, *(RA 2.3.3) Data Access Middleware*, *(RA 2.3.1) Secure Outsourced Data Processing* and *(RA 2.3.2) Secure Collaborative Data Processing*. These services will support application; platform and infrastructure providers to comply with (inter)national cybersecurity and -crime laws and regulations. The *(RA 2.4.1) Legal Compliance Analysis* activities will take into account laws that apply horizontally (i.e. across all sectors, such as GDPR) as well as sector-specific rules (e.g., the Health Insurance Portability and Accountability Act (HIPAA) security and privacy rules). At network and hardware level, The Cybersecurity Research Strategic Programme will provide solutions for *(RA 3.2.2) Secure Communication Protocols for the IoT*, *(RA 4.3.3) Cryptographic Protocols for Blockchain*, and *(RA 4.3.5) Security Analysis of Cryptographic Protocols*.

C) Mobile Payments – Supported by Cloud Platforms

Problem context for cloud/web platform providers. Imagine a payment card financial transaction whereby a cardholder leverages a digital container accessed by a mobile device; a smartphone, for instance, stores wallet applications, payment credentials and loyalty cards that can be used to make proximity and remote mobile payments. Tokenized payment credentials are either stored securely in the mobile phone (if NFC) or in the cloud. Wallet transactions may be completed using, for instance, near field communication (NFC) “Pay wallets (e.g., Apple Pay, Samsung Pay, and Android Pay) or cloud-based card-on-file wallets (e.g., PayPal, Pay, and Amazon). These transactions may also support biometrics, pin and signature for consumer authentication.

In addition, online payment services can support short-term B2B loans to bridge the waiting time between issuing a bill to a customer and the eventual payment by the customer. Using the unpaid bill as a pledge (or security), such online loan services help to improve the cash flow and liquidity of the company, which is crucial for business continuity and operations.

The focus of cloud, web and mobile services is on outsourcing towards third parties; trust plays a key role here. Blockchains offer novel trade-offs between distributed and centralized trust, and result in novel combinations of increased transparency with resilience and security but often at a very high cost in terms of resources.

Demands for cybersecurity. Strict security policies must be enforced, for instance to control that a customer has not yet paid the bill to the company (in this case it cannot be used as a pledge or security), the customer has not used the unpaid bill as a pledge for another loan at another bank (i.e. the well-known “double spending” problem), or to prevent a bank from excluding or blocking other banks to offer loans by formulating false claims (e.g., that the bill is already paid, that the bill is already used for another loan).

This requires that transaction state information must be maintained persistently, while being able to trace changes back to the multiple interacting parties who will change this state. For example, a simple money transfer from party A to party B results in a withdrawal from the balance of party A and a deposit on the account of party B. A next transaction by party A is dependent on the new balance of the account of party A, as the first transaction might have reduced the balance below a certain threshold. In case of the running example described before, if Bank A provides a loan to a company C based on a bill B, then company C cannot request a loan at another bank. If Bank A’ registers a bill B’ as paid, then company C’ can also not request a loan at any bank based on that bill B’. The interacting parties in this running example are clearly the banks, and the shared state is a register with paid bills and provided loans based on specific bills.

Such shared state could be maintained by a centralized authority (e.g., the national bank) or a decentralized blockchain. One of the main challenges for the latter is to make blockchain technology scalable and reconcile it with confidentiality requirements.

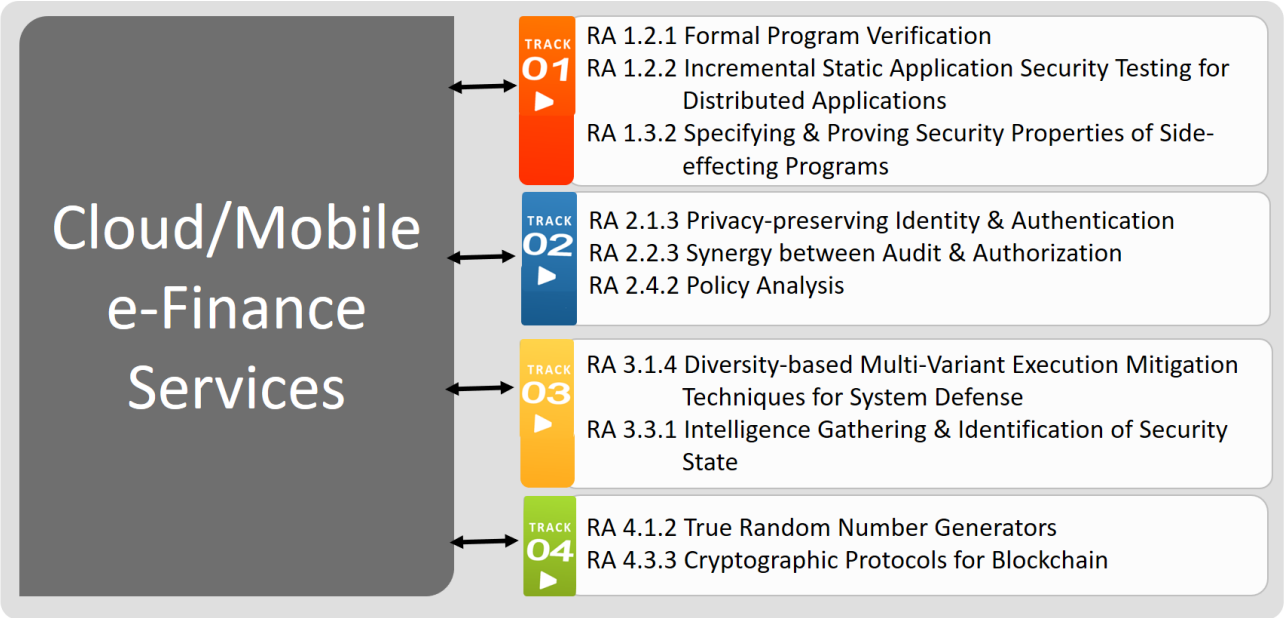

Figure 6-4: Relevant Cybersecurity Solutions for Cloud/Mobile Based e-Finance Services.

Synergetic cybersecurity research to address industry demand. The Cybersecurity Research Programme will diversify and focus on various types of cloud, web and mobile deployment contexts (see Figure 5.4). For example, client-side system security mechanisms such as *(RA 3.1.4) Diversity-based Multi-Variant Execution Mitigation Techniques for System Defense* and *(RA 3.3.1) Intelligence Gathering & Identification of Security State* will be tackled for a diverse range of platforms in web and mobile.

For cloud application and platform developers, dedicated efforts are planned to investigate security services for *(RA 2.1.3) Privacy-preserving Identity & Authentication*, and *(RA 2.2.3) Synergy between Audit & Authorization* in support of isolation and virtualization, containerization and orchestration, multi-cloud and hybrid cloud settings, storage and business continuity. In general, the Consortium will actively engage in *(RA 2.4.2) Policy Analysis* activities to identify regulatory hurdles, assessing implications of recent and upcoming policy initiatives, and influence ongoing policy discussions in relation to cybersecurity and cybercrime.

Special focus will be put on enabling technologies that increase trust in the cloud operator based on technologies such as secure processors (e.g., in line with technologies such as Intel SGX), *(RA 4.3.3) Cryptographic Protocols for Blockchain* and *(RA 4.1.2) True Random Number Generators*.

In support of SaaS-based architectures, storage centric-cloud applications, architectures for workflow and business processes, additional strategic research themes include the Secure Software Development Life-Cycle (SDLC) for cloud-based solutions (process) and especially *(RA 1.2.1) Formal Program Verification*, *(RA 1.2.2) Incremental Static Application Security Testing for Distributed Applications*, and *(RA 1.3.2) Specifying & Proving Security Properties of Side-effecting Programs*.

D) Summary

Multiple research results, emerging from each of the Tracks of the Programme, can contribute to deliver and enhance the cybersecurity posture of a modern ICT platform (IoT-based, cloud-based etc.). While the Strategic Research Programme aims to deliver top class results that each have a stand-alone scientific and technical value, the Consortium also aims for combinations of results that can pave the way for robustly securing key business applications. This section has sketched opportunities by looking into three examples, yet the Programme will dynamically target such application cases by proactively defining and analyzing opportunities in dialogue with industry.

6.4. Delivery & Validation of Technology Assets that Emerge from Research

This section briefly sketches the approach to maintain a synergetic collaboration between the academia and industry (by the cybersecurity Consortium and beyond) and to transfer emerging know how and technology to enable practical use by industry and by other actors in society.

A) From the Perspective of the Strategic Research Programme

The Consortium will pro-actively identify mature, high-potential research results from the four Research Tracks and prioritize these results to provide some strong and robust prototypes that can be validated and demonstrated. These will be strong candidates (future assets) for further use in industrial applications.

This strategy to deliver and validate concrete technology assets will drive three types of activities:

- *[Validation of mature research prototypes]* The Strategic research program will engage with industry from the early Proof-of-Concept phase onwards in order to validate the value and feasibility of specific results in the context of applications.
- *[Consolidation of results into solutions: platforms and tools that enable cybersecurity]* The Strategic Research Programme will integrate related security building blocks (resulting from the Programme or available as state-of-the-art assets) into solutions (platforms, tools, tool chains etc.) for industry.
- *[Instantiation of cybersecurity solutions into industrial pilot applications]* High-potential research results, typically platforms and tools as sketched above will be demonstrated in a realistic business context.

B) From the Perspective of the Flemish Cybersecurity Programme as a Whole

The Strategic Research Programme will deliver prototypes, integrated solutions (tools and platforms that may combine multiple results) and demonstrations that show the application of the former. An important goal of the Cybersecurity programme is to enable industry implementation of these cybersecurity solutions.

Such implementations will emerge from various types of activities, supported by the strategic research Consortium.

These activities include

- *Industrial pilot studies*. Pilot studies evaluate the potential and benefits of resulting hardware and software technologies and keep industrial partners abreast of relevant emerging cybersecurity technologies. (Some of these studies fit with the Flemish COOCK program.)
- *Cybersecurity technology assessments*. A technology assessment will study and position the added value of the cybersecurity technologies, for instance by means of a technology gap analysis, the assessment of third-party products, or a survey of emerging technologies in a specific domain or context
- *Follow-up applied research in collaboration with industry*. The Consortium will actively engage in opportunities for follow-up applied research projects, which are ideal to analyze domain-specific requirements and evaluate promising research results with one or more industrial partners. These projects target industrial proof-of-concept prototypes that integrate cybersecurity technologies with hardware and software used and provided by the industrial partner(s). (Some of these research activities can operate as ICON projects, or other government supported projects for industry.)

The overall Cybersecurity programme creates various opportunities for industry implementation. Funding opportunities are available through VLAIO: e.g., O&O projects, ICON and TETRA projects, and COOCK. These instruments are specific for Flanders.

The European level offers additional possibilities for academia and industry (e.g., Horizon Europe, ECSEL, ITEA3, PENTA, CelticPlus, EIT Innovation projects). The upcoming Horizon Europe program, for example, will include a strong focus on next-generation Internet and Cybersecurity. The ECSEL Joint Undertaking added in

their latest multi-annual strategy plan[24] a specific chapter on safety, security and reliability, and confirms that "Safety, security and reliability are fundamental components of any innovation in the digital economy." In their Vision 2030[25] report, the two leading industry associations in the domains software-intensive systems & services and embedded & Cyber-Physical Systems (ITEA and the ARTEMIS Industry Association) confirm the need for next-generation cybersecurity solutions and request a doubling of the investment in software innovation to keep Europe on par with the rest of the world in sustaining the benefits of Digital Technology for European economy and society.

[24] ECSEL Multi-Annual Strategic Plan 2018, http://ec.europa.eu/research/participants/data/ref/h2020/other/wp/jtis/h2020-wp18-ecsel_en.pdf

[25] ITEA ARTEMIS-IA High-level Vision VISION 2030: 'opportunities for Europe', https://itea3.org/news/itea-artemis-ia-high-level-vision-2030-opportunities-for-europe-officially-launched.html

7. Conclusion

7.1. Summary of the Programme

The proposed Research Programme strengthens the existing, core competences, while reaching out to industry (1) to ensure applicability of knowledge and technology, and (2) to operate with an up-to-date prioritization of topics in cybersecurity. At the same time, collaboration and synchronization with other leading labs in Europe will ensure that Flanders further invests in its research strengths in cybersecurity, while creating synergy and collaboration with other leading centers in Europe – thus avoiding duplication of efforts.

This document presents at a high-level, the Strategic Programme for Cybersecurity Research in Flanders. This programme has to deliver impactful solutions to real-world challenges, while starting from and building upon academic excellence. The execution of the Programme has to strengthen existing core competences in cybersecurity research, while delivering building blocks and solutions that will benefit cybersecurity in industry.

The latter is reflected in Figure 7-1. The core of the Research Programme covers four Tracks, as explained before.

- Track 1 addresses **Application and Software Security** and aims to support all stakeholders that analyze, develop and deploy new application software, while using an evolving set of technologies in the context of secure software development.
- Track 2 includes **Strategic Security Services**, such as authentication, authorization and services for data protection. The overall idea is that many security specific building blocks (reusable components or services that are typically offered as security middleware) will not be built from scratch in new applications, and should be evolving with new demands and expectations – typically reaching beyond the state-of-practice in industry offerings.
- Track 3 covers **System and Infrastructure Security**. Here one expects stable, secured technology that is packaged as a black box in an operating system or in network layers. Software and service developers rely on the robustness of these lower layers – yet we all know that additional research is essential to meet the promise.
- Track 4 covers the **basic Technology Building Blocks for Security**: secure hardware, cryptography and secure cryptographic implementations.

The Consortium gathers research groups with a strong and proven base and history, residing in KU Leuven, UGent, VUB and the strategic research center imec. The programme will strengthen existing teams, groups and activities and make research progress as fast as possible. The Research Tracks have been identified and confirmed in dialogue with industry and the relevance of each of the proposed research themes has been confirmed by many industry stakeholders. The Consortium will develop prototype platforms that apply research results in the context of important technology platforms such as IoT environments, cloud platform and platforms for data sharing.

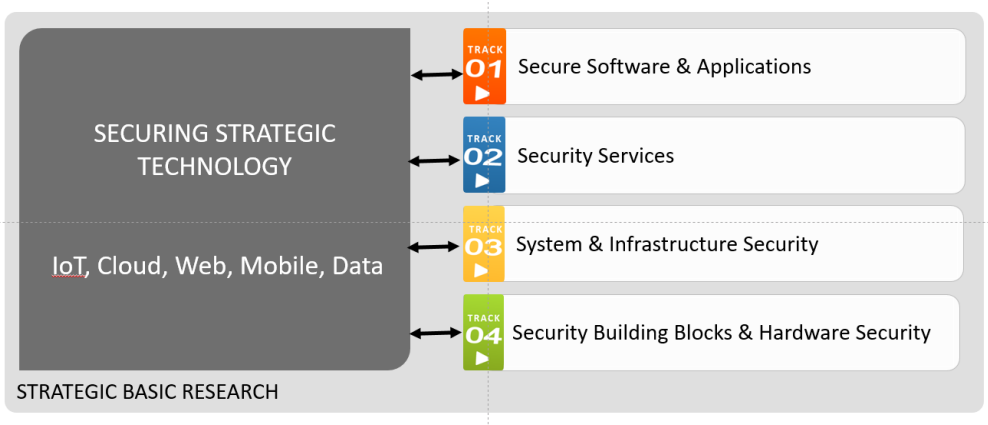

Figure 7-1: A Strategic Program with Four Research Tracks

7.2. Refinement and Operational Plan of the Research Program

The Research Programme will start on September 1, 2019. The ambition of the Consortium is to aggressively hire research staff and grow the planned research activities at maximal speed. Part of the recruitment challenge will be shared amongst the partners by jointly developing a communication plan. This plan is part of the Outreach Programme of the Flemish Cybersecurity Initiative. Formally, the Outreach Programme will service many objectives – and not recruitment as such – but strong communication and outreach obviously will contribute the visibility of the research program and the recruitment opportunities of all research partners in the Consortium .

The Programme is expected to run at full force by the end of 2020. In the meantime, the timing of starting specific research activities will depend on the recruitment process. The potential progress and intensity of the research activities will be monitored and scheduled as a function of the teams that have effectively been assembled. This will demand for dynamic management in year 1 and 2 of the programme (estimating 4 to 6 quarters).

Each of the Research Tracks is coordinated by one principal faculty member:

- Track 1 will be coordinated by Prof. Bart Jacobs (DistriNet). The Track as a whole will be delivered in close collaboration between DistriNet and VUB.
- Track 2 will be coordinated by Prof. Frederik Vercauteren (COSIC). The Track as a whole will be delivered in close collaboration between COSIC, CiTiP and DistriNet (all KU Leuven).
- Track 3 will be coordinated by Prof. Frank Piessens (DistriNet). The Track as a whole will be delivered in close collaboration between COSIC, DistriNet and UGent.
- Track 4 will be coordinated by Prof. Ingrid Verbauwhede (COSIC). The Track as a whole will be delivered in close collaboration between COSIC and imec.

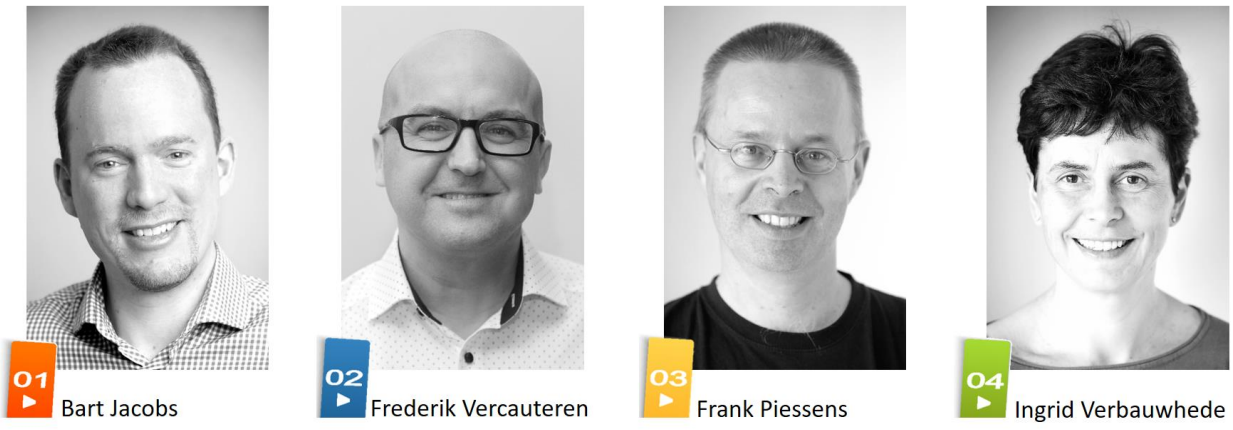

Figure 7-2: One Principal Faculty Member per Research Track

The sub-consortia per Research Track are relatively compact and all partners have a successful history of collaboration. This factor minimizes risk. From a management perspective, we plan for quarterly progress meetings per Research Track, and for a quarterly meeting to manage the overall programme at the Consortium level. The latter forum will also monitor the interaction with industry and the development of prototypes and platforms that facilitate technology validation and transfer of knowledge, as introduced and articulated in Chapter 6. It should be noticed that the development of such platforms will be managed at the Consortium level by Prof. Wouter Joosen and Prof. Bart Preneel, with the support of the three valorization managers that are currently present in the Consortium .

Coordination at the Consortium level will also organize interaction with and feedback from industry. Even though this is largely a continuous and informal process, the Consortium will organize a plenary event twice a year, once in spring and once in the fall.

The Consortium will implement a relatively lightweight management approach; it must be stressed that strong results are expected from collaboration within small teams of focused researchers, and/or collaborations between a small number of these teams. Hence new concepts, techniques and results should be created, refined and evaluated without the governance or strong presence from a central body. In that sense, the cybersecurity research programme embeds a lot of decentralization and bottom-up research dynamics. The Consortium expects a high level of productivity and creativity from this model, yielding a substantial improvement in key indicators such as top publications, PhD dissertations, research prototypes, etc.

7.3. Evaluation and Kick-off

The Cybersecurity Research Programme will be evaluated annually. The process will occur in close collaboration with an International Scientific Advisory Board (ISAB). This board gathers for the first time in July 2019, and will meet at least annually.

Research Tracks, Themes and Activities will be reported upon and results will be summarized. The details of this process will be articulated at the first meeting of the ISAB (July 4-5, 2019). Subject to a positive evaluation of the initial Programme description (presented in this document), the strategic research activities can kick off on September 1, 2019.

Soon thereafter, by the end of the first quarter, the Consortium will organize an industry briefing and collect additional feedback.

8. Annexes

8.1. Consortium

This Annex describes the Consortium in detail by sketching each of the research groups involved (Section 2), including the Faculty Members and Permanent Research Staff who work on Cybersecurity (Section 3).

8.2. Budget Headlines

The administration has required to provision for an allocation of the full budget as of year 1. It is obvious that this budget will not be fully consumed in the first and second year of the programme, as aggressive recruitment will be ongoing for quite a while. In practice, the research groups plan to build up the capacity that will require the entire budget by Q4 2021. The total budget of the basic research program amounts to 8M€ on an annual basis.

The presented budget allocation is proportional to the critical mass of the research groups involved (the number of faculty members and the number of permanent staff in academic research in cybersecurity). In this respect, the consortium description in Annex I clarifies the group-level proportion of the budget.

In addition, it should be noticed that the levels of Track 3 (2.385 M€) and Track 4 (2.252 M€) are substantially larger compared to Track 1 (1.325 M€) and Track 2 (1,457 m€).

	COSIC	**DistriNet**	**VUB**	**CiTiP**	**Ugent**	**imec**	**TOTAL (€ K)**
TRACK 1	-	795,0	530,0	-	-	-	**1.325,0**
TRACK 2	530,0	530,0	-	397,5	-	-	**1.457,5**
TRACK 3	530,0	1.457,5	-	-	397,5	-	**2.385,0**
TRACK 4	1.722,5	-	-	-	-	530,0	**2.252,5**
Total Budget Basic Research (€ K)	**2.782,5**	**2.782,5**	**530,0**	**397,5**	**397,5**	**530,0**	**7.420,0**
Programme Coordination	75,0	75,0	-	-	-	-	**150,0**
TRACK 1 Management	-	45,0	30,0	-	-	-	**75,0**
TRACK 2 Management	30,0	30,0	-	22,5	-	-	**82,5**
TRACK 3 Management	30,0	82,5	-	-	22,5	-	**135,0**
TRACK 4 Management	97,5	-	-	-	-	30,0	**127,5**
Total Budget Management & Coordination (€ K)	**232,5**	**232,5**	**30,0**	**22,5**	**22,5**	**30,0**	**570,0**
Grand Total (€ K)	**3.015,0**	**3.015,0**	**560,0**	**420,0**	**420,0**	**560,0**	**7.990,0**

8.3. Overview Research Activities

Research Track 1: Application and Software Security
Theme 1. Secure SDLC – Secure Software Development Life Cycle
(RA 1.1.1) Cybersecurity Requirements
(RA 1.1.2) Cybersecurity-by-Design Solutions
(RA 1.1.3) Security Analysis for Existing Applications
Theme 2. Program Verification
(RA 1.2.1) Formal Program Verification
(RA 1.2.2) Incremental Static Application Security Testing (SAST) for Distributed Applications
(RA 1.2.3) Efficient Runtime Application Security Protection (RASP) for Distributed Applications
Theme 3. Secure Programming Languages and Secure Compilation
(RA 1.3.1) Mechanically-verified Security Proofs for Capability Machine Programs
(RA 1.3.2) Specifying and Proving Security Properties of Side-Effecting Programs
(RA 1.3.3) Language-embedded Security Policies for Distributed Micro-services.

Research Track 2: Strategic Security Services
Theme 1. Identity Management and Authentication
(RA 2.1.1) Identity
(RA 2.1.2) Frictionless Authentication: Collaborative and Continuous
(RA 2.1.3) Privacy-preserving Identity and Authentication
Theme 2. Authorization and Audit
(RA 2.2.1) Enhancing Authorization Capabilities
(RA 2.2.2) Intelligent Audit
(RA 2.2.3) Synergy between Audit and Authorization
Theme 3. Advanced Encryption Techniques and Data Access Middleware
(RA 2.3.1) Secure Outsourced Data Processing
(RA 2.3.2) Secure Collaborative Data Processing
(RA 2.3.3) Data Access Middleware
Theme 4. Policy and Regulation

(RA 2.4.1) Legal Compliance Analysis
(RA 2.4.2) Policy Analysis
(RA 2.4.3) Legal Engineering Analysis

Research Track 3: System and Infrastructure Security
Theme 1. System Security
(RA 3.1.1) Protection Against Software-Controlled Side-Channel Attacks (on general purpose hardware)
(RA 3.1.2) Processor Extension to Support New System Security Models
(RA 3.1.3) Security and Safety In Mixed Criticality Systems
(RA 3.1.4) Diversity-based Multi-Variant Execution Mitigation Techniques for System Defense
Theme 2. Network Security
(RA 3.2.1) Study of Critical Internet Components and Protocols
(RA 3.2.2) Secure Communication Protocols for the IoT
(RA 3.2.3) Analysis of Protocol Implementations
Theme 3. Security Monitoring and Management
(RA 3.3.1) Intelligence Gathering and Identification of Security State
(RA 3.3.2) Methods and Tools for Secure Deployment
(RA 3.3.3) Detection and Response for IoT and Industrial Control Systems

Research Track 4: Technology Building Blocks: Secure Hardware, Cryptography and Secure Implementations
Theme 1. Secure Hardware: Roots of Trust Anchored Into Technology Foundations
(RA 4.1.1) Developing PUFs
(RA 4.1.2) True Random Number Generators
(RA 4.1.3) Technology Solutions to Secure Circular Economy
Theme 2. Cryptographic algorithms
(RA 4.2.1) Symmetric-key Algoritms
(RA 4.2.2) Public-key Algorithms
(RA 4.2.3) Proofs and Validation
Theme 3. Cryptographic Protocols

(RA 4.3.1) Cryptographic Protocols for Distance Bounding
(RA 4.3.2) Cryptographic Protocols Design for MPC Applications
(RA 4.3.3) Cryptographic Protocols for Blockchain
(RA 4.3.4) Cryptographic Protocols for Mix Networks
(RA 4.3.5) Security Analysis of Cryptographic Protocols
Theme 4. Secure and Efficient Cryptographic Implementations
(RA 4.4.1) Implementation Challenges of Post-quantum, FHE, Lightweight Crypto on Novel Compute Platforms
(RA 4.4.2) Side-Channel and Fault Attacks
(RA 4.4.3) White-Box Cryptography