

UiO : **University of Oslo**

Elahe Fazeldehkordi

Security and Privacy Solutions in IoT and Distributed Systems Design

Thesis submitted for the degree of Philosophiae Doctor

Department of Technology Systems and
Department of Informatics

Faculty of Mathematics and Natural Sciences



2021

© **Elahe Fazeldehkordi, 2021**

*Series of dissertations submitted to the
Faculty of Mathematics and Natural Sciences, University of Oslo
No. 2375*

ISSN 1501-7710

All rights reserved. No part of this publication may be reproduced or transmitted, in any form or by any means, without permission.

Cover: Hanne Baadsgaard Utigard.

Print production: Representralen, University of Oslo.

Abstract

New distributed systems in computing technology are appearing on the market day by day. Consequently, new security and privacy challenges arise when designing these systems. These challenges demand a need to look for comprehensive and more precise approaches that can provide higher levels of security, privacy, and trust from the design phase in these systems.

In order to develop systems including open distributed systems, we need techniques and tools for specification, design, and code generation. For the initial ideas of such systems, it is desirable to use graphical notations, so that ideas can easily be understood. This is relevant for the early design phase. Once the ideas are agreed upon, it is then desirable to have a formal basis for the specifications of the desired system, in order to support rigorous reasoning about specifications and designs. This can be done in the late design phase.

Existing documents for security and privacy requirements and functionalities in IoT systems lack some of the security functionalities, and in particular, they do not focus on privacy issues or are presented only in textual form, without defining a framework. Structures of these documents are also complicated. Systems are often made without the help of security and privacy experts. So a comprehensive graphical representation framework is needed and should be easy to follow, even for non-experts.

Formal methods have come into wide use in the design and verification of safety, security, and privacy of systems in the industry, because they are very effective in verifying safety, security, and privacy requirements of systems, requirements for which testing is mostly ineffective. Formal verification techniques can guarantee that a design is free of specific flaws.

Provability of IoT and distributed systems depends on the language used to model the system, its semantics, and the kind of system properties to prove and the techniques used to verify them. One may take a bottom-up approach, starting with low level languages in use, or one

may take a top-down approach, starting with a more abstract language with an expressiveness suitable for IoT and distributed systems. Such a language can be used to model IoT and distributed systems, and if the language is executable, it can be used for simulation and prototyping of IoT and distributed systems, and if the language is imperative with standard mechanisms such as object orientation, IoT and distributed models made in the language can easily be translated to more low level languages. The latter approach also has the advantage that one can define the language and its semantics so that it is amenable to semantical analysis and verification. Additionally, this approach makes program reasoning simple and powerful.

The active object paradigm provides a natural way of modeling distributed systems in general, and IoT systems in particular, because it covers the fundamental aspects of IoT systems such as distribution of concurrent, autonomous units communicating by message passing, where each unit can run on a device with limited processing power and limited storage.

In this thesis, we focus on two aspects of the design phase to develop methodologies to improve security, privacy, and trust levels of IoT and distributed systems. Firstly, we consider the early design phase, where the architecture of a system is being developed, focusing on the setting of IoT systems. Secondly, we consider the late design phase, where models, in particular executable models and prototypes are designed, focusing on IoT systems and also the more general setting of distributed systems. For the modeling related to the late design phase, we limit our work to the active object paradigm. In particular, we find techniques for increasing the security level of the overall IoT systems using a security and privacy functionality framework, and a technique for detecting potential vulnerabilities that can cause distributed denial of service (DDoS) attacks using static analysis technique. In addition, we developed a language-based approach to provide trust, safety, security, and privacy for smart contracts.

Dedication

This thesis is dedicated to my beloved family and my fiancée for their endless support, encouragement, patience, and love in my life and their belief in my dreams.

Acknowledgements

This thesis would have been impossible without the support and mentoring of my main supervisor Olaf Owe. I would like to express my deep gratitude to him for his priceless time and help whenever I needed. I am grateful for his continuous encouragement, inspiration, patient guidance, fruitful discussions, constructive and constant cooperations, and optimism throughout this work, and for all the knowledge I have learned from him. I am also grateful for his great friendship and for teaching me various things of life. I would like to express my special thanks to my co-supervisor Josef Noll for his great help and support, critical suggestions, patient guidance, and active collaboration in various papers. I would also like to thank Martin Steffen for his constructive and useful feedbacks, comments, discussions, interesting lectures, and for all his help whenever I needed. I would like to thank my graduate committee members, Jüri Vain, Svetlana Boudko, and Paal Engelstad for their great feedbacks, help, and encouragement. I also thank Toktam Ramezanifarkhani and Christian Johansen for their feedbacks and guidance during this research.

I would like to thank everybody involved in the IOTSEC project, in particular I thank Habtamu Abie for his great help, support, encouragement, and constructive feedbacks. I am grateful for all the great assistance and facilities provided by the administration staff and the technical support at the Department of Technology Systems and the Department of Informatics, University of Oslo. My thanks also go to all researchers and professors who have crossed my path as friends and colleagues over the years, in particular Jüri Vain.

I am grateful to my family and my fiancée. Words cannot express how grateful I am to them. They have always encouraged and supported me in pursuing my dreams with their love and patience, and they have always been with me despite the distance. Finally, I would like to thank my friend Mona for her company during these years.

This work was supported by the IOTSEC project. Additional funding was provided by the Department of Technology Systems, SCOTT and ConSeRNS projects.

Contents

| | |
|---|-------------|
| Abstract | i |
| Dedication | iii |
| Contents | vii |
| List of Figures | xi |
| List of Tables | xiii |
| | |
| Part I Overview | 1 |
| | |
| 1 Introduction | 3 |
| 1.1 Background and Motivation | 4 |
| 1.2 Research Questions | 12 |
| 1.3 Methodology | 13 |
| 1.4 Structure of the Thesis | 14 |
| | |
| 2 Theoretical Background | 17 |
| 2.1 Object-Oriented Systems | 17 |
| 2.2 Concurrent Object-Oriented Systems | 19 |
| 2.3 Creol/ABS Language | 27 |
| 2.4 Formal Verification and Specification | 29 |
| | |
| 3 Summary of Research Papers and Contributions | 41 |
| 3.1 Paper 1 | 42 |
| 3.2 Paper 2 | 43 |
| 3.3 Paper 3 | 43 |
| 3.4 Paper 4 | 44 |
| 3.5 Additional Publications | 47 |
| | |
| 4 Conclusions and Future Work | 49 |

| | | |
|----------|--|------------|
| 4.1 | <i>Conclusions</i> | 49 |
| 4.2 | <i>Future Work</i> | 54 |
| | <i>Bibliography</i> | 57 |
| | <i>Part II Scientific Contributions</i> | 67 |
| 5 | <i>Paper 1: Security and Privacy Functionalities in IoT</i> | 69 |
| 5.1 | <i>Introduction</i> | 69 |
| 5.2 | <i>IoT-Related Standards and Guidelines</i> | 72 |
| 5.3 | <i>Related Work</i> | 73 |
| 5.4 | <i>Framework Explanation</i> | 75 |
| 5.5 | <i>Security Classification</i> | 79 |
| 5.6 | <i>Pacemaker Case Study</i> | 81 |
| 5.7 | <i>Conclusion</i> | 92 |
| | <i>Bibliography</i> | 95 |
| 6 | <i>Paper 2: A Case Study of Healthcare Products</i> | 99 |
| 6.1 | <i>Introduction</i> | 100 |
| 6.2 | <i>Related Work</i> | 102 |
| 6.3 | <i>Security Classification</i> | 103 |
| 6.4 | <i>Pacemaker Case Study</i> | 104 |
| 6.5 | <i>Discussion</i> | 113 |
| 6.6 | <i>Conclusion</i> | 115 |
| | <i>Bibliography</i> | 117 |
| 7 | <i>Paper 3: A Language-Based Approach to Prevent DDoS Attacks</i> | 121 |
| 7.1 | <i>Introduction</i> | 121 |
| 7.2 | <i>Overview</i> | 124 |
| 7.3 | <i>Related Work</i> | 126 |
| 7.4 | <i>Our Framework for Active Object Systems</i> | 128 |
| 7.5 | <i>Static Analysis to Prevent Attacks</i> | 130 |
| 7.6 | <i>Examples of Possible DoS/DDoS Attacks</i> | 134 |
| 7.7 | <i>Conclusion</i> | 143 |

| | |
|---|------------|
| Bibliography | 145 |
| 8 Paper 4: An Approach to Smart Contracts Supporting Safety and Security | 149 |
| 8.1 Introduction | 150 |
| 8.2 Smart Contracts and Blockchain | 153 |
| 8.3 A High-Level Language for Active Object systems | 156 |
| 8.4 The Proposed Framework | 165 |
| 8.5 Specification of the Auction Example and Improve- ments | 179 |
| 8.6 Verification | 184 |
| 8.7 Evaluation | 188 |
| 8.8 Related Work | 193 |
| 8.9 Conclusion | 196 |
| Bibliography | 199 |

List of Figures

| | | |
|------|--|-----|
| 1.1 | <i>An illustration of the system development life cycle . . .</i> | 5 |
| 5.1 | <i>IoT security and privacy functionality framework . . .</i> | 73 |
| 5.2 | <i>Security mechanisms</i> | 73 |
| 5.3 | <i>Human resource security</i> | 76 |
| 5.4 | <i>Physical and environmental security</i> | 76 |
| 5.5 | <i>Privacy protection</i> | 76 |
| 5.6 | <i>Operations security</i> | 77 |
| 5.7 | <i>Development, maintenance, and audit</i> | 77 |
| 5.8 | <i>IoT security and privacy functionality framework – de-commissioning</i> | 77 |
| 5.9 | <i>Basic inputs for defining a security class</i> | 80 |
| 5.10 | <i>Case study – scenario 1; with a security controller along with the pacemaker inside the body</i> | 87 |
| 5.11 | <i>Case study – scenario 2; with a security controller for the pacemaker out of the body for enhanced security</i> | 87 |
| 6.1 | <i>Case study – scenario 1; with a security controller along with the pacemaker inside the body</i> | 109 |
| 6.2 | <i>Case study – scenario 2; with a security controller for the pacemaker out of the body for enhanced security</i> | 109 |
| 7.1 | <i>Distributed communication (s-obj stands for server object and c-obj for consumer object)</i> | 123 |
| 7.2 | <i>Distributed object communication in DDoS</i> | 124 |
| 7.3 | <i>Control flow graph</i> | 131 |
| 7.4 | <i>Algorithm for detecting flooding by means of calls and comps sets in a given cycle</i> | 132 |
| 7.5 | <i>The interfaces of the units in the subscription example</i> | 135 |
| 7.6 | <i>Classes providing an implementation of the subscription example</i> | 136 |
| 7.7 | <i>DoS attack by a variation of the subscription example</i> | 137 |

| | | |
|------|---|-----|
| 7.8 | <i>The graph and call/comp sets for the original version of the program</i> | 137 |
| 7.9 | <i>The graph and call/comp sets for the modified version of the program</i> | 137 |
| 7.10 | <i>Flooding by unbounded creation of innocent clients targeting the same server</i> | 139 |
| 7.11 | <i>Static detection of flooding using unbounded creation</i> | 139 |
| 8.1 | <i>Interfaces for the auction example</i> | 158 |
| 8.2 | <i>Examples of function definitions</i> | 161 |
| 8.3 | <i>The auction class</i> | 166 |
| 8.4 | <i>The interfaces of futures and history objects</i> | 167 |
| 8.5 | <i>History objects</i> | 168 |
| 8.6 | <i>Predefined types for transactions</i> | 170 |
| 8.7 | <i>A class implementation of futures and history objects</i> | 172 |
| 8.8 | <i>Illustration of the future mechanism</i> | 174 |
| 8.9 | <i>An implementation of history objects where blacklisted bidders are blocked</i> | 175 |
| 8.10 | <i>The privacy-extended history class for auction</i> | 176 |
| 8.11 | <i>A history class implementation with safety check</i> | 178 |
| 8.12 | <i>The auction history interface and class</i> | 180 |
| 8.13 | <i>A trust-extended history class for auction</i> | 183 |

List of Tables

| | | |
|-----|---|-----|
| 5.1 | <i>Security classes</i> | 80 |
| 5.2 | <i>Exposure</i> | 80 |
| 5.3 | <i>Security and privacy challenge comparison of scenario 1 and 2</i> | 89 |
| 5.4 | <i>Security class of the sensor</i> | 91 |
| 5.5 | <i>Security class of the pacemaker security controller</i> . . . | 91 |
| 5.6 | <i>Security class of the mobile phone</i> | 91 |
| 6.1 | <i>Exposure</i> | 103 |
| 6.2 | <i>Security classes</i> | 103 |
| 6.3 | <i>Security and privacy challenge comparison of scenarios 1 and 2</i> | 109 |
| 6.4 | <i>Security class of the sensor</i> | 113 |
| 6.5 | <i>Security class of the pacemaker</i> | 113 |
| 6.6 | <i>Security class of the mobile phone</i> | 113 |

Part I

Overview

Introduction

The modern society is critically depending on information systems, in particular distributed systems and the Internet of Things (IoT). A distributed system is a system that consists of components placed on different networked computers which communicate concurrently by passing messages to each other. Examples of distributed systems include computer networks like the Internet, aircraft control systems, distributed information processing systems like banking systems, mobile phones, etc. The Internet of Things (IoT) consists of different distributed systems and is considered one of the most important areas of future technology, it has received massive attention from a wide variety of industries. IoT is made from a variety of *things* or *objects* – like sensors, mobile phones, television, refrigerators, actuators, etc. – that are able to interact with each other and cooperate with their neighbours through the Internet to reach a common goal using unique addressing schemes.

The aim of IoT is to provide an advanced mode of communication among the various systems and devices, and also to facilitate the interaction between humans and the virtual world. With this aim, IoT plays a significant role in the modern society and has applications in almost all fields including healthcare systems, automobile, industrial appliances, sports, homes, entertainments, environmental monitoring, etc. Services provided by IoT applications offer great benefits for human's life, but they can come with a huge price with respect to people's privacy and security protection. IoT devices have already outnumbered the number of people at a computerized workplaces. As a result of this expansion, and as these things are connecting to the Internet, and generate, process, and exchange huge amounts of security-critical and privacy-sensitive information, therefore they are attractive targets of various attacks [21, 29, 41, 47, 54, 55, 57, 61, 62, 63, 75, 76, 77, 83, 86]. The same attack may be used on a large number of (identical) devices.

What makes the IoT different from the traditional Internet is the absence of human role. IoT offers complex systems that can sense the external environment and make decisions with no need of human

intervention. Existing sensors in IoT capture almost every single piece of information from the environment such as image, video, sound, location, proximity, temperature, humidity, acceleration, pressure, and heartbeat. Hence, through these systems a lot of information about human life is going to be collected and processed with some environments like smart wearables, smart healthcare products, or smart homes that are able to sense and manage very sensitive and personal information. This huge amount of uncontrolled collection of sensitive information requires strong data protection and privacy controls.

Despite the advanced abilities provided by IoT and distributed systems in the data communication area, their vulnerability implications from a security and privacy standpoint are still of great concern. Cyberattacks on IoT and distributed systems such as healthcare systems, banking systems, and very recently digital coins are very critical since they can cause significant physical and economic damages, and even threaten human lives. It is therefore vital to make these systems more secure and data privacy-protected. This has made the investment in techniques for preventing security and privacy attacks very important. Appropriate steps should be taken in the initial phase of development and design of these systems.

1.1 Background and Motivation

1.1.1 System Development Life Cycle

System development life cycle (see Figure. 1.1), also referred to as the software/application development life cycle, is whole process of developing and maintaining an information system. system development life cycle comprised of the planning phase, the requirement analysis phase, the system design phase, the implementation phase, and the integration and testing phase.

System design phase defines the architecture, modules, interfaces, and data for a system to satisfy specific requirements. This phase is very important in the sense that, it bridges the gap between the problem domain and the system, supposedly in a manageable way. The focus of this phase is to find a solution of how to make a system that fulfills the desired purpose. This phase finds a solution to how the system can

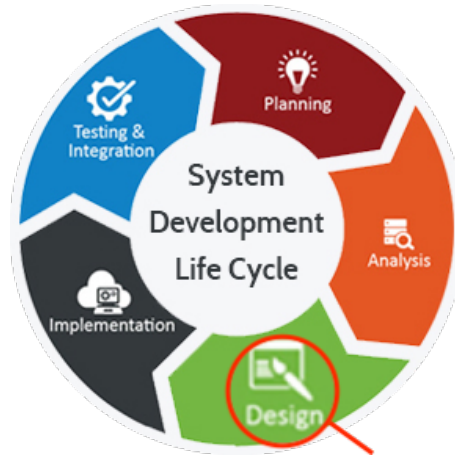


Figure 1.1: An illustration of the system development life cycle [90]

be implemented and operate in such a way that the specification can be satisfied. In the design phase, the complex activity of system development will be divided into several smaller sub-activities, which coordinate with each other to achieve the main objective of the system development.

In order to develop systems including open distributed systems, we need techniques and tools for specification, design, and code generation. For the initial ideas of such systems, it is desirable to use graphical notations, so that ideas can easily be understood. This is relevant for the early design phase. Once the ideas are agreed upon, it is then desirable to have a formal basis for the specifications of the desired system, in order to support rigorous reasoning about specifications and designs. This can be done in the late design phase.

1.1.2 Security and Privacy Taxonomy of IoT Systems

Having a comprehensive overview and *taxonomy* of security and privacy requirements and functionalities in IoT at the design phase is a prerequisite for architecting optimal security solutions, designing, and developing secure and privacy-aware IoT systems that can help both providers and consumers of IoT devices to have a better understanding of the security and privacy aspects. By functionality we mean: “The security and privacy-related features, functions, mechanisms, services, procedures, and architectures implemented within organizational information systems

or the environments in which those systems operate” [79].

General standards, guidelines, and frameworks [1, 2, 25, 32, 33, 68, 79, 81] for establishing, implementing, maintaining, and continually improving an information security management system, and protecting the confidentiality of Controlled Unclassified Information (CUI), exist. Existing IoT-related documents lack some of the security functionalities, and in particular, they do not focus on privacy issues or are presented only in textual form, without defining a framework. Layer structures of these documents are also complicated. Systems are often made without the help of security and privacy experts. So a comprehensive graphical representation framework is needed and should be easy to follow, even for non-experts. Such a representation should extract all the IoT-related security and privacy functionalities and unify them using a common vocabulary. The vocabulary should be categorized, and the guidelines and requirements should be integrated in a uniform style.

1.1.3 Privacy and GDPR

With respect to privacy, the European Union (EU) has passed the *General Data Protection Regulation* (GDPR) [84], the toughest privacy and security law in the world. It came into force on 24th May 2016, and since 25th May 2018 all the organizations and their software are required to be compliant. GDPR replaced the European Data Protection Directive which established minimum data privacy and security standards in 1995 and based on that every member state depends on its implementing law. The aim of GDPR is to protect and empower all EU citizens/residents’ data privacy and to reshape the way organizations approaching data privacy of such people. Even though GDPR is passed by the EU, it forces obligations on organizations anywhere if they collect and process data related to EU citizens/residents, or offer goods or services to these people. Penalties for violating the GDPR are very high and can be 4% of global revenue or €20 million (whichever is greater).

GDPR regulates the following principles:

Accountability; The purpose of limitation principle; Data minimization; Lawfulness, fairness and transparency; Accuracy; Data retention periods; and Data security.

The key rights of Data Subjects, Controllers, and Processors are as follows:

Breach notification, right to access, right to rectification, right to be forgotten, right to restrict processing, right to data portability, right to object, rights in relation to automated decision making and profiling, consent, and data protection by design and by default.

The GDPR requirements relate to almost all modern computer systems and should be taken into consideration in the design phase to support privacy by design.

1.1.4 System Modeling

Modern computer systems comprise of complex hardware and software components, nevertheless ensuring the quality and the correctness of the software part has been often a more significant problem than the one in the underlying hardware. A documentation of this problem providing over 100 “software horror stories” can be found in [30]. The correctness of the software is, in particular, very important from the security and privacy point of view. Any single mistake in the software of a system can be misused by the attackers. We should ensure that the software is reliable and works correctly and does not have any vulnerability, ideally during the design process.

One way to promote correct and reliable programs is through *modeling* and *verification* techniques. *System modeling* develops abstract models of a system. Each model presents a different view of that system [10]. Models help to derive detailed requirements for a system: during the design process they describe the system to engineers, and after implementation document a system’s structure and operation. System models simplify a system and do not represent a complete system. By leaving out details of the system and decreasing the amount of complexity, we will be able to understand the system and analyze specific properties related to the system more easily.

Models can be developed from different perspectives of a system [82]: (1) an external perspective models the context or environment of the system; (2) an interaction perspective models the interactions between a system and its environment, or between the components of the system; (3) a structural perspective models the organization of a system or the

structure of the data processed by the system, and (4) a behavioral perspective models the dynamic behavior of the system and how it responds to events.

1.1.5 Formal Methods

The use of *formal methods* technology in the industry has been an objective for researchers for several decades. Formal methods are mathematically-based approaches to software development that define a formal model of the software. This model can then be formally analyzed to look for errors and inconsistencies. Using mathematical approaches, users will be able to analyze and verify these models at any part of the program life-cycle (requirements engineering, specification, architecture, design, implementation, testing, maintenance, and evolution).

Formal methods are important because they increase reliability of a system. Industrial applications of formal methods and verification technologies have been discussed in an excellent survey by Woodcock et al. [89]. Formal methods have come into wide use in the design and verification of safety-critical systems in the industry, because they are very effective in verifying safety and security requirements of systems, requirements for which *testing* is mostly ineffective. Testing and simulation are techniques that provide a variety of input conditions for a system and ensure that the output is as expected. The inputs can be provided by a designer or randomly generated. These techniques sample the response of the systems to the chosen inputs [74]. They can sample *functional properties* of a system – i.e., what a system is supposed to do, but they are not able to verify *safety and security properties* – i.e., what a system is not supposed to do. For the former properties, given a nominal input it produces the specified output. It is often assumed that the variety of designed-in behaviors are usually small enough to adequately test for compliance. This is only a small fraction of possible valid inputs. The latter properties are predicated on any path that could bring the system into an unsafe or insecure state; therefore, the input state space is vast and almost never accessible by testing. Mathematical analysis and verification of the program are needed to be sure that safety and security requirements are met and to guarantee that a design is free of specific flaws.

Formal verification techniques can guarantee that a design is free of specific flaws. Formal verification is the act of proving or disproving the

correctness of intended algorithms of a system concerning a certain formal specification or property, using formal methods. Formal verification techniques can check the behavior of a design for *all* input vectors.

In distributed systems like smart contracts, there is a need for formal verification in the most popular language, Solidity, that has been used for writing these contracts [4, 39, 94]. Also, in these systems we can achieve a high level of trust and privacy with less expensive (with respect to time, resources, and power consumption) approach using formal methods.

One of the formal verification techniques is program analysis. Program analysis offers static (compile-time) techniques for predicting safe and computable approximations to the set of values or behaviours arising dynamically at runtime when executing a program on a computer [64].

In services like financial systems, static analysis techniques are very important. Even a single unfortunate mistake by a financial service provider could be enough to cause unintended attacks or be misused by attackers and influence customers and this would of course, destroy the customer's trust and confidence and result in financial loss and reputation damage. One should make sure that the software is not harmful for the customers before running it, and this makes the static detection very important in these systems. Most of the existing methods recommend using of Internet traffic monitoring techniques to detect and react to possible threats. But, this kind of runtime protection may slow down or temporarily shut down the websites and their services. Despite the importance of the static detection method as a valuable complement to runtime detection methods, it is underrepresented.

Provability of IoT and distributed systems depends on the language used to model the system, its semantics, and the kind of system properties to prove and the techniques used to verify them. One may take a bottom-up approach, starting with low level languages in use, or one may take a top-down approach, starting with a more abstract language with an expressiveness suitable for IoT and distributed systems. Such a language can be used to model IoT and distributed systems, and if the language is executable, it can be used for simulation and *prototyping* of IoT and distributed systems – i.e., the process of building a model of a system, and if the language is *imperative* – i.e., describes steps that change a program's state and tells *how* a program should operate, with standard mechanisms such as object orientation, IoT and distributed models made in the language can easily be translated to more low level IoT and

distributed languages. The latter approach also has the advantage that one can define the language and its semantics so that it is amenable to semantical analysis and verification. Additionally, this approach makes program *reasoning* (i.e., proof) simple and powerful.

1.1.6 Suitability of Modeling Paradigms for IoT and Distributed Systems

Object orientation is the leading paradigm used in the software industry today and is recommended for open distributed systems [78]. Program code in class-based object-oriented languages is organized in classes, encapsulating variables and methods, and classes can be specialized by subclassing, possibly with a redefinition of methods. Classes can be instantiated at runtime into objects. Many of the most popular programming languages are based on object orientation including Java, C++, C#, Python, etc. The object-oriented paradigm provides flexibility through polymorphism. A single function can shape-shift to adapt to whichever class it is in, we can create one function in the parent class and this one function would work with other functions. Reusable code through inheritance techniques saves time. Using inheritance, one can simply reuse existing code instead of duplicating code. Modularity is provided by the object-oriented paradigm, something which reduces complexity and makes it easier for troubleshooting. Objects may be seen as autonomous units. Object orientation provides good structuring mechanisms and flexibility in applying changes to artifacts produced during the software development process. This allows the structure to remain the same in the final system. By combining object orientation and *concurrency*, we can benefit from the increased responsiveness and throughput of distributed system.

Distributed systems are by nature concurrent, therefore modeling paradigms for distributed systems should support concurrency. Concurrent computing increases program throughput, parallel execution in concurrent programming allows the number of tasks completed in a given time to increase at best proportionally to the number of processors. Concurrent computing also increases responsiveness for input/output, it allows the time that would be spent mostly waiting for input or output operations to complete, to be used for other tasks. Concurrent computing

can also have more appropriate program structure, while some problems and problem domains are better-suited for representation as concurrent tasks or processes.

Standard models of object interaction do not address the specific challenges of distributed computation. Object interaction based on method calls is usually synchronous. However, synchronous communication in a distributed setting, causes undesired and uncontrolled waiting, and possibly deadlocks. Asynchronous message passing provides better control and efficiency, although it does not provide the structure and discipline inherent in method calls.

The *active object paradigm* [13, 91, 92] has evolved from the *actor model* [3] and object orientation. In this paradigm each object is concurrent and autonomous, with its virtual processor, and all interaction is through method invocations and replies, thus supporting two-way interaction rather than one-way interaction as in the case of the actor model. Therefore at most, one process is active on an object at a time, and method executions are decoupled from method invocations using underlying asynchronous message passing. Active object languages have mechanisms for efficient method interaction without active waiting, by means of the future concept or suspension.

The active object paradigm provides a natural way of modeling distributed systems in general, and IoT systems in particular, because it covers the fundamental aspects of IoT systems such as distribution of concurrent, autonomous units communicating by message passing, where each unit can run on a device with limited processing power and limited storage [48].

Each IoT device is modeled as one active object. If a single IoT device has several concurrent activities, such a device could be modeled by a group of active objects, using the concept of *object groups* as suggested and formalized in [12, 22, 50]. Such an object group can be seen as a single object for the environment. The modeling of different kinds of IoT networks, including wireless networks, can be modeled as in [52].

1.2 Research Questions

As discussed earlier, new distributed systems in computing technology are day by day appearing on the market, e.g., IoT systems. Consequently,

new security and privacy challenges arise when designing these systems. These challenges demand a need to look for comprehensive and more precise approaches that can provide higher levels of security, privacy, and trust from the design phase in these systems. For example, new taxonomy frameworks that organize all aspects of security and privacy baselines, guidelines, and recommendations, or the need to formally verify these systems at the software level. Therefore, this thesis is seeking solutions to the following main goal:

Main objective: *To improve security and privacy of IoT and distributed systems during the design phase.*

The main objective is a broad area of research that requires a more defined focus. Therefore, we focus on two aspects of the design phase. First, the early design phase, where the architecture of a system is being developed, considering the setting of IoT systems. Second, the late phase of the design, where models, in particular executable models and prototypes are designed, considering IoT systems and also the more general setting of distributed systems. Having enough information at the architectural design phase and using graphical notations make it easier to understand the system components and their specifications for the modeling phase. These give two subgoals:

Objective 1: *To develop a methodology to improve the overall security and privacy of IoT systems in the early design phase.*

Objective 2: *To develop methodologies to improve the security, privacy, and trust level of distributed systems in the late design phase.*

To achieve the first subgoal, this thesis will address the following specific research questions:

RQ1: What are the security and privacy functionalities for IoT systems relevant at the early design phase?

RQ2: How can we build a security and privacy functionality framework for IoT systems relevant to the early design phase?

RQ3: How can the functionality framework help to improve security and privacy of IoT systems at the early design phase?

To achieve the second subgoal, this thesis will address the following specific research questions:

RQ4: How can we model distributed systems at a high level of abstraction?

RQ5: How can we analyze modeling frameworks of distributed systems to identify and detect possible vulnerabilities in the models in order to improve the security and trust level at the late design phase in these systems?

RQ6: How can we improve the security, privacy and trust levels of distributed systems using a methodology at the late design phase?

RQ7: How can we use formal verification and specification for the behavior of models of distributed systems?

1.3 Methodology

This thesis consists of several publications that describe our research results achieved in the direction of answering the above questions. Besides, the methodology we are looking for should meet the following requirements:

Model-based approach. A model is mainly a representation of some aspects of interest from a real-world system [10]. Models provide an abstraction of a real-world system. Therefore they contain less complexity than real-world systems. By decreasing the amount of complexity, we will be able to understand the interesting aspects better, and analyze specific properties related to them. Therefore it is desirable for the methodology presented in this thesis to be model-based. In addition, it is desirable that the model is executable in order to allow prototyping and verification of specification requirements of the system, and that it defines a program structure that can be a guideline for the implementation phase.

Formality. Formal methods are important since they seek to increase the reliability of systems by adding mathematical rigor to the design process [15]. For the methodology presented in this thesis, it is desirable to have a design notation that offers precision, unambiguity, and abstractions based on a formal syntax and a well-defined semantics.

The paradigm of concurrent and distributed active objects. We consider programming and specification languages suitable for modeling IoT and distributed systems. We define these languages, their semantics, and systems that are proof oriented towards reasoning support. In order to make reasoning as simple and powerful as possible, we focus on languages at a high level of abstraction rather than low level languages. We consider executable and imperative languages supporting object-oriented principles, and designed the languages so as to reduce the complexity of further translation to low level languages for IoT and distributed systems. The active object paradigm provides a natural way of modeling distributed systems in general, and IoT systems in particular, because it covers fundamental aspects of IoT systems, such as distribution of concurrent units communicating by message passing, where each unit can run on a device with limited processing power and limited storage [48]. Furthermore, the active object model can be seen as complementary to graphical modeling languages including the Unified Modeling Language (UML) [53], therefore is a suitable continuation of the early design phase.

For the modeling related to objective 2, we therefore limit our work to the active object paradigm.

1.4 Structure of the Thesis

The rest of this thesis is structured as follows: Chapter 2 introduces Creol/ABS language and discusses its underlying concepts. Chapter 3 gives a short summary of each of the papers collected in this thesis and contributions of this thesis by answering the research questions presented in Chapter 1. Chapter 4 concludes the thesis and suggests some directions for future research. Part II collects the research papers in the thesis.

Theoretical Background

In this chapter we discuss some basic concepts of object-oriented and concurrent object-oriented modeling languages and then position the Creol/ABS modeling languages [48, 51] in this context. For this purpose, Sections 2.1 and 2.2 introduce object-oriented languages, and concurrent object-oriented languages, respectively. The presentation in Sections 2.1 and 2.2 is inspired by [6, 60]. Section 2.3 introduces specifically the Creol/ABS languages, which supports the active object modeling paradigm.

To build a reliable software system, it is important to develop techniques which facilitate verification about the behavior of the program code. For this purpose, we organize the discussion in Section 2.4 by first considering formal logical systems and programming logic, then discussing formal verification approaches including Hoare logic and static program analysis.

2.1 Object-Oriented Systems

Object orientation is claimed to be the most used programming techniques, in particular, for distributed systems [18]. The concept of objects was first introduced by Ole-Johan Dahl and Kristen Nygaard in Simula 67 [26, 28]. The ideas of the language which was system description as well as simulation programming originated by Kristen Nygaard at the Norwegian Computing Center in 1961 [66]. Along with Nygaard, Ole-Johan Dahl produced the initial ideas of object-oriented programming, which still is a leading approach for commercial and industrial applications today.

In the object-oriented programming style, a system is described as a collection of objects. An object is an integrated unit of *data* and *procedures* (called methods) that act on these data. We can think of object as a box that stores some data and is able to act on these data. The data in objects is stored in *variables*. A variable's content changes by executing an assignment statement. Variables in an object may be hidden from the environment, so called private variables. One may declare interfaces

to specify what methods can be called from the outside. By seeing an object solely through interfaces, object's variables will be strictly private and are not accessible to other objects. In class-based object-oriented languages, objects are described by *classes*, defining the common features for each kind of objects. Elements (*instances*) of a class have the same names and types for their variables and execute the same code for their methods. A class serves as a blueprint for building its instances. In the object-oriented model, objects exchange information with each other under the general label of *message passing* – a sender object sends a request to a receiver object to execute a certain method. The sender of the message may provide some parameters to the method and the method can return a result, which is passed to the receiver. Interaction between objects only happens based on the precisely determined message interface. Each object is responsible and is able to maintain its own local data in a consistent state.

At any time during the execution of a program, a new object can be generated, so there can be a large number of objects at some point. Objects are not destroyed explicitly. Although, if they certainly do not influence the correct execution of the program they can be removed through garbage collection.

The features listed below are common among languages considered to be strongly class- and object-oriented [60]:

- *Classes*: object-oriented languages are diverse, but the most popular ones are class-based, that means objects are instances of classes. A class is a description of a set of objects that are defined using a list of behavior (operations), properties (attributes), relationships, and semantics.
- *Encapsulation*: is a mechanism upon which certain features of a class will be inaccessible for user's calls. *Information hiding* sometimes used as a synonym for encapsulation. When writing a class, there are features that are part of the implementation of a class, but not of its interface. Other features of the class possibly available to users, may call the features if they need, but it is not possible for a user to call them directly. The encapsulation mechanism should be flexible in such a way that allows designers of a class to specify a feature available to all users, to no user,

or to some specific users; and an attribute that may need to be available for access only, access and restricted modification, or full modification.

- *Inheritance*: the basic idea is that when we want to define a new class it is often easier to start with variables and methods of existing class(s), then add more features to them in order to reach to the desired new class. In this case, we say that the new class inherits the variables and methods of the superclass(s). In other words, a class can be defined as an extension or restriction of another.
- *Polymorphism and dynamic late binding*: the ability of a program entity to be attached to objects of various possible types, and the guarantee that every feature on an entity will trigger the correct feature, corresponding to the attached runtime object's type.
- *Dynamic object creation instantiation*: at any point of the execution of a program a new object can be generated, so there might be a large number of objects.

Additional features of object-oriented languages are given as follows:

- *Multiple and repeated inheritance*: a class can have more than one parent and can inherit from another class in more than one way.
- *Abstraction*: is a mechanism that aims to representing a simplified model of the program without showing irrelevant details in order to handle complexity and enhance understanding [40, 56, 93]. Abstraction is provided by interfaces.

2.2 Concurrent Object-Oriented Systems

The concept of concurrent programming goes back to 1960s within the concept of operating systems. The idea was to build hardware units called *channels* or *device controllers* that operate independent of a controlling processor. They allow an I/O operation to work concurrently with continued execution of the central processor's program instructions. The program is called concurrent when it contains two or more processes that work together in order to do a task. Each of these processes is a sequential

program execution a sequence of statements one after another. In contrast to sequential programs that have a *single thread of control*, concurrent programs have *multiple threads of control*. The term *multithreaded* means a program has more processes or threads than processors to execute the threads. Concurrent programs are in general more complex than sequential programs. Execution in a concurrent program begins in some initial state and as processes execute independently, they examine and modify the program state.

2.2.1 Fundamental Concepts

Here, we look at five fundamental concepts of concurrent programs: program state, atomic actions, the concept of history, safety properties, and liveness properties [6]. The values of all program variables at a point in time are called the *state* of the program.

A sequence of statements executed in a process. A sequence of one or more *atomic actions* implement a statement. Atomic actions indivisibly examine or alter the program state. Therefore, the execution of a concurrent program results in an interleaving of the sequences of atomic actions that are executed by every process.

A *history* (also called a *trace* of the sequence of states) can be seen as a particular execution of a concurrent program: $s_0 \rightarrow s_1 \rightarrow \dots \rightarrow s_n$, where s_0 is the initial state, $s_1 \dots s_n$ are subsequent states, and atomic actions make the transitions that change the state. In a concurrent program, every execution generates a history. In all but the most trivial programs, there is an enormous number of possible histories, since the next atomic action in each one of the processes could be the next one in a history, so there will be many paths that actions can be interleaved no matter a program always begins in the same initial state.

Synchronization aims to constrain the possible histories of a concurrent program to the desirable ones. There are two kinds of synchronization: *mutual exclusion* that ensures at most one process is executing its *critical section* at the same time, and *condition synchronization* that delays an action until the state holds a Boolean condition. Critical sections cannot be interleaved with actions in other processes when they reference the same variables.

In a program a *property* is an attribute that is true of every possible history of that program, so of all executions of the program. There are

two types of properties: safety and liveness. A *safety property* [6] is when the program never enters a bad state for instance a state that some variables have an undesirable value. A *liveness property* [6] is when the program eventually enters a good state for instance a state that variables have desirable values.

Partial correctness [6] is an example of a safety property. A program is called partially correct if it requires that if an answer is returned it will be correct. *Termination* is an example of a liveness property. If every loop and procedure call of a program terminates, then the program terminates, in this case, the length of every history of the program is finite. A *total correctness* property of a program combines partial correctness and termination. If a program always terminates with a correct answer, the program is called totally correct.

2.2.2 Communication in Concurrent Programming

In a concurrent program, the only non-trivial way that processes can work together in order to solve a problem is to *communicate* with each other in some way [6]. And this communication can only be done if one process writes into *something* that the other process can read. This something might be a shared variable or a shared communication channel that it depends on the hardware in which the program will execute on. Therefore, communication can be done either by writing and reading shared variables or by sending and receiving messages.

Shared variable communication is when one process writes into a variable that will be read by another process. Shared variables are mostly used in concurrent programs that execute on hardware where its processors share memory. In distributed-memory machines, if they are supported by a software implementation of so called distributed shared memory, the shared-variable programming model can also be used.

In a *distributed-memory architecture* where the processors have their own private memory, the communication is mostly through a communication network instead of shared memory. In this case, processes are not able to communicate through sharing variables, instead they communicate by exchanging messages with each other or by invoking remote operations. Concurrent programs that use message passing are named *distributed programs*, since processes can be distributed all over the distributed-memory architecture's processors. A distributed

program can also be executed on a shared memory multiprocessor, like any concurrent program that can be executed on a single, multiplexed processor. In this case, channels are implemented through shared memory rather than a communication network.

There are several mechanisms for distributed programming including asynchronous message passing, synchronous message passing, remote procedure call (RPC), and rendezvous. They are different in the way channels are named and used and the way communication is synchronized.

Message passing is when one process sends a message to a channel, another process can acquire the message through receiving it from the channel. Sending a message can be asynchronous (non-blocking) or synchronous(blocking). In asynchronous message passing, a channel is an unlimited queue of messages that have been sent but not received any answer yet. *Send* is a *non-blocking* primitive and never causes delay, but *receive* is a *blocking* primitive and it might cause delay until a message arrives to be received. In contrast, synchronous message passing causes the sender to *block* until the message is received, that means a process can send only one message on any channel at the same time and until that message receives and the receiver sends back a reply to the sender, the sending process cannot continue its execution and send another message. Synchronous message passing has one advantage that there is a limitation on the size of communication channels, therefore on buffer space. However, it has two disadvantages. First, concurrency is reduced because when two processes communicate, at least one of them will have to block. Another disadvantage of synchronous message passing is that this form of message passing is more prone to *deadlock* – when two or more processes are waiting for conditions that will never happen, therefore none of them could proceed. Therefore the difference between asynchronous and synchronous message passing is having more concurrency, less waiting time and having bounded communication buffer, respectively. Since memory is plentiful and compare with synchronous send, asynchronous message passing is less prone to deadlock, most programmers prefer it. In particular, for distributed environments considering that communication in a network takes time, asynchronous send is more suitable.

Remote procedure call (RPC) and rendezvous are ideally suitable for programming of client/server interactions. In RPC and rendezvous, an operation is a two-way communication channel from the caller to

the process that services the call and then back to the caller instead of sending information through one-way communication channels in message passing. One-way communication channels need a large number of channels for the two-way information flow between clients and servers that they need two explicit message interactions through two different message channels. In addition, every client needs a different reply channel. Also, in both RPC and rendezvous, the execution of call delays the caller until the called operation is executed and any result is returned as in synchronous message passing.

The difference between RPC and rendezvous is in the way of servicing the invocations of operations. RPC declares a procedure for every operation and provides a new process to handle every call, the caller and procedure body may be on different machines that's why it is called *remote procedure call*. In contrast, rendezvous uses an existing process to handle calls (rendezvous is sometimes called *extended rendezvous*).

2.2.3 Actor Model

The term *actor* was defined by Carl Hewitt at MIT [43] in 1973. The idea was to describe the concept of reasoning agents. However, over the years it has been refined into a model of concurrency.

An actor is a fundamental unit of computation that embodies all three essential elements of computation: processing, storage, communication. When an actor receives a message it can create some more actors, it can send messages to actors that it has addresses before and it can designate how to handle the next message that it receives. Actors are isolated from each other and they do not share memory, they have state and the only way to change the state is by receiving a message. Each actor has a unique mail address that is determined at the time of its creation. This address is used to specify the recipient of a message. Every actor has its own mailbox which is similar to message queue. Messages are stored in actors mailboxes until they are processed. Actors after they are created wait for the messages to arrive, actors can communicate with each other only through messages. Messages are sent to actor's mailboxes and processed in First In First Out (FIFO) order. Messages have immutable data structures that can be sent over the network. Conceptually an actor can only process one message at the time. Actors work asynchronously and they do not need to wait for a response from another actor. Actors

have addresses, it is possible for an actor to send a message to another actor by knowing its address, an actor can only communicate with actors whose addresses it has. An actor has addresses of the actors that it has created itself and also it can obtain other addresses from a message it receives. One actor can have many addresses.

The actor model supports maximum concurrency, however it includes excessive process creation overhead. With the actor model, developers do not have to worry about locking and race conditions. Properties of communication in the actor model are as follows: actors do not use a channel or intermediaries so there is nothing like type and semantics issues which exist in case of having a channel or things like buffer or guarantee delivery. Actors use best effort delivery. Actors have at-most-once-delivery. Messages can take arbitrarily long to be delivered. There is no message ordering guarantees. The running state of an actor is monitored and managed by another actor (called a *supervision*). The supervision can perform actions based on the state of the actor.

2.2.4 Active Object Model

The active object model was first proposed by Yonezawa et al. [92] in an object-oriented concurrent programming language called ABCL/1. Their computation model evolved from the actor model [9, 42], and it differs from actor computation model in many respects. Here, we explain the important differences between the two models.

Every object in the active object model of Yonezawa et al. has its own thread of control and it may have its local persistent memory. An object always has one of the three modes: *dormant*, *active*, or *waiting*. An object is initially in dormant mode. If an object receives a message that satisfies one of the specified patterns and constraints, it changes to the active mode. When the sequence of actions that are performed in response to an accepted message have been completed by an active object and there is no more subsequent messages have arrived, an active object changes to the dormant mode again. An active object sometimes needs to stop its current activity and wait for a message with a specified patterns to arrive. In this situation, the active object changes to the waiting mode and when it receives a required message it changes to the active mode again.

In the active object model, two distinct modes of message passing have been considered: *ordinary* and *express*. For each object T , there are two message queues: one for messages sent to T in the ordinary mode and the other for messages sent in the express mode. Messages are enqueued in arrival order.

Ordinary mode message passing [92]:

Suppose a message M sent in the ordinary mode arrives at an object T when the message queue associated with T is empty. If T is in the dormant mode, M is checked as to whether or not it is acceptable according to T 's *script* (a description which satisfies its behavior: what messages it accepts and what actions it performs when it receives such messages). When M is acceptable, T becomes active and starts performing the actions specified for it. When M is not acceptable, it is discarded. If T is in the active mode, M is put at the end of the ordinary message queue associated with T . If T is in the waiting mode, M is checked to see if it satisfies one of the pattern-and-constraint pairs that T accepts in this waiting mode. When M is acceptable, T is reactivated and starts performing the specified actions. When M is not acceptable, it is put at the end of the message queue.

Express mode message passing [92]:

Suppose a message M sent in the express mode arrives at an object T . If T has been previously activated by a message which was also sent to T in the express mode, M is put at the end of the express message queue associated with T . Otherwise, M is checked to see if it satisfies one of the pattern-and-constraint pairs that T accepts. If M is acceptable, T starts performing the actions specified for M even if T has been previously activated by a message sent to T in the ordinary mode. The actions specified for the previous message are suspended until the actions specified for M are completed. If so specified, the suspended actions are aborted. But, in default, they are resumed.

In the active object model, there are three types of message passing: *past*, *now*, and *future*.

Past type message passing:

An active object O sends a message M to an object T and it proceeds its own task without waiting for a response to the requested task from the object T .

This type of message passing significantly increases the concurrency of activities in a system and corresponds to the standard way of message passing in the actor model.

Now type message passing:

An active object O sends a message M to an object T and it waits for a response to the requested task from the object T .

This convention is useful in assigning a returned value to a variable in an assignment.

Future type message passing:

An active object O sends a message M to an object T but it does not need the expected requested result immediately. Therefore, object O does not wait for object T to return the result and proceeds its tasks and computations. Meanwhile or later when object O needs the result, it checks its special *private* object called *future object* that was specified at the time of sending M . If the result is available in the future object, it can use it otherwise it continues its computations again until a result is available.

The future object behaves like a queue and the contents of the queue can be checked or removed by the object O . The concurrency of a system increases by future message passing.

One of the important difference between the active object model and actor model is that in the active object model, an object in the waiting mode can accept a message which is not necessarily at the head of the message queue, however an actor in the actor model can only accepts a message that is at the head of the message queue. Moreover, now type and future type message passing are not supported in the actor model. Hence, when an actor A sends a message to a target actor T and expects a result from T must terminate its current activity before being able to receive the response from actor T , like any other incoming messages arriving to A . In addition, this necessity of the termination of actor A 's current activity

in order to receive actor T 's response leads to an unnatural decomposition of A 's task and requires additional procedures to identify a message that responds to the reply value of now type message passing.

2.3 Creol/ABS Language

The Creol language proposes programming constructs for distributed concurrent objects depending on asynchronous method calls and processor release points [48], building on the OUN language [70] and the language idea of [78]. Concurrent objects are active, meaning that an object encapsulates an execution thread. Method calls are represented by pairs of asynchronous messages. The language is class-based and supports inheritance. Objects have identity meaning that an object's "name" is unique. Communication is between named objects and object names may be exchanged between objects. In Creol object variables (references) are typed by interfaces. The language is extended with mechanisms for static type control in dynamically reconfigurable systems. The language is strongly typed, depending on a nominal type-system that makes sure invoked methods are supported by the callee (if not nil) and formal and actual parameters match [49]. Creol has a modular and compositional semantics meaning that each execution step in the semantic involves only one object. Its operational semantics is defined in rewriting logic [59] and is executable in the tool Maude [23].

ABS is an *abstract behavioral specification* language building on Creol, addressing *executable formal model* specifications for distributed object-oriented systems and has code generators for Java [37], Scala [67], Maude [24], and Erlang [7]. Both Creol and ABS allow a high level specification of a system including its concurrency and synchronization mechanisms, and also local state updates. Therefore, the concurrent control flow of object-oriented systems are captured by Creol/ABS models, while implementation details that are not desirable at the modeling level are abstracted away, like the concrete representation of internal data structures, which are captured by an applicative data type language, the scheduling of method activations, which is simply non-deterministic, and the properties of the communication environment. Like Creol, ABS uses interfaces as types for object variables, while the classes that implement the functionality of these objects are abstracted away

in the type system. ABS has a notion of object groups (cogs), which is similar to the one in JCoBox [80], it generalizes the concurrency model of Creol from single concurrent objects to concurrent object groups sharing the same processor.

In ABS, the communication mechanism of Creol [48] for remote communication has been used, it is based on asynchronous method calls using first-class futures [14]. Another feature of the ABS and Creol communication mechanism is their use of *cooperative scheduling* between asynchronous called methods. However, in ABS this is lifted from the level of objects to the level of concurrent object groups. Cooperative scheduling is a mechanism, allowing passive waiting which means there will be a process queue holding all suspended processes of a given object, the executing process of an object can be suspended so that a required result from another process of the queue of the object can be provided. The type system in ABS is similar to the one in Featherweight Java [46], a core calculus for Java. Featherweight Java is class-based with single inheritance, with subtyping as a reflexive and transitive closure of the subclass relation. Both Creol and ABS distinguishes between classes and types. In Creol asynchronous calls and interfaces are combined with multiple class inheritance, choice operators, and a notion of *cointerface* at the interface level in order to accommodate type-safe calls backs to callers [49]. The cointerface is the interface of the caller, through which the callee can make calls back to the caller. In this case the cointerface will ensure type correctness of these call-backs. Creol introduces a suspension mechanisms, which is also used in ABS. A process (or a method execution) may suspend while waiting for a result to arrive or condition to be satisfied. Suspension releases the processor of the object using a process queue to store suspended processes, allowing the object to continue with other (enabled) processes. Process release (suspension) points allow non-blocking waiting.

The language version that we consider in this thesis is inspired by both Creol and ABS and we refer to it by Creol/ABS. It ignores some of the features in ABS (like cogs) and supports some features that ABS does not support (inheritance, cointerfaces). Our work is mainly using local futures, rather than shared futures.

For the setting of IoT and distributed systems, local futures (in particular method-local ones) are suitable since they can be deleted by each object as soon as they are no longer needed so there is no future to

be garbage collected. However, if the active object paradigm supports *first-class futures (also called shared futures)* – i.e., futures that can be shared with other objects if they need them, it is not clear when none of the objects need the shared future(s) and when they can be deleted. In this case, garbage collection of futures will in general be needed. This is of special importance for the setting of IoT systems where garbage collection, in particular distributed garbage collection, is undesired due to resource limitations. For distributed systems, first-class futures may pose security and privacy issues, since they can be seen as shared variables with read-only access. Local futures are not shared between objects and security and privacy control is easier.

2.4 Formal Verification and Specification

Correctness proofs in programs are important in order to establish program properties. Results of the execution of processes when they are executed in parallel, can depend on the unpredictable order in which the actions from different processes are executed. This will result in a complexity that is very massive to handle informally. Also, program testing is not able to detect all mistakes since the particular interactions in which errors are visible may not happen. Therefore, it is important to have a proof system in order to remove some of this complexity, structure concurrent programs in a simple way, and to verify their correctness.

Below we will first look at verification methods (in Section 2.4.3) including semi-automatic ones, and methods that are fully automatic.

2.4.1 Formal Logical Systems

A formal logical system [27, 36] is used for inferring theorems from axioms based on a set of rules. These rules are defined in the form of: a set of *symbols*, a set of *formulas* built from the symbols, a set of *axioms*, and a set of *inference rules*. Formulas are a sequence of symbols. Axioms are formulas that are assumed to be true and are used as a starting point for further reasoning and arguments. Inference rules determine a way to derive additional true formulas from axioms and other true formulas.

Inference rules are in the form bellow:

$$\frac{H_1 H_2 \dots H_n}{C}$$

where H_i is a *hypothesis*, and C is a conclusion. This means: if all the hypotheses are true, then conclusion is also true. Hypotheses and conclusion are formulas or schematic representations of formulas.

A *proof* in a formal logical system is a finite sequence of line, each line (any line in a formal proof is called *theorem*) can be an axiom or can be derived from previous lines by applying an inference rule.

A formal logical system is a mathematical abstraction (i.e., a collection of symbols and relations between them). In order to be able to know if the statements that the formulas represent are true or not, we need to provide an interpretation of the formulas.

An *interpretation* of a classical logic maps formulas to either true or false. In a logic, if all its axioms and inference rules are sound, the logic is called *sound*. Soundness of axioms and inference rules means all the axioms map to true, and all the inference rules conclusions map to true by the assumption that the hypotheses map to true.

The dual of soundness is *completeness*. In a logic, if each semantically true formula is provable, the logic is called *complete*.

2.4.2 Programming Logic

A *programming logic* is a formal logical system that allows us to state and prove properties of programs. A programming logic aims in facilitating proving properties of program execution. Like any formal logical system, programming logic also contains symbols, formulas, axioms, and inference rules. The symbols of programming logic are predicates, braces, and programming language statements.

One way to formally proof correctness of systems is Hoare logic. Hoare logic is a formal system with a set of logical rules for correctness proofs of programs including concurrent programs. Hoare [45] was the first to develop a formal logic for proving partial correctness properties of sequential programs, this work was inspired by the first proposed technique for proving correctness of programs by Floyd [34]. Later Susan Owicki in her dissertation supervised by David Gries, developed the first

complete logic for proving partial correctness properties of concurrent programs [71, 72, 73] based on Hoare logic for parallel programs [44].

Hoare introduced the concept of a *triple*, interpretation for triples, axioms, and inference rules for sequential statements. Since then logical systems that follow this style are called Hoare logic. Programming logic is an example of a Hoare logic. Formulas in programming logic and Hoare logic are triples of the following form:

$$\{P\} S \{Q\}$$

where the predicates P and Q (often called *assertions*) describe relations between the values of program variables, and S is a statement or statement list. The *interpretation* of a triple has the following meaning:

Whenever execution of a program S starts in a state where assertion P holds, then the assertion Q will hold in a resulting state, assuming the execution of the program terminates.

This interpretation is named *partial correctness*. Predicates P and Q are called the *precondition* and *postcondition* of S , since they characterize conditions that must hold before and after the execution of S .

2.4.3 Formal Verification

The act of proving or disproving the correctness of the algorithms and programs in a system corresponding to certain *specifications* and properties of the system, by applying formal methods of mathematics, is called *formal verification*. *Formal specification* precisely describes system properties using mathematic-based languages.

There are many approaches to formally verify programs. The specification formalism that represent a specification, provides a good discrimination criterion between different approaches. Different formalisms are classified to three groups [38]: logic, automata and language theory, and hybrid formalisms. In the first one, specifications are expressed as statements in a specific kind of logic. In the second one, a specification is represented in the form of a language (or an automaton). And, the last one uses a hybrid approach meaning that a specification first expressed using a logic statement then it is converted to an equivalent language/automaton

representation. Below we give a hint of approaches that use each of these formalisms. Due to extensive literature available in each area of formal verification, we cannot possibly do justice to the fine details of every approach and we ask the reader to refer to the references for more details.

2.4.3.1 Logic

Logic in general is the study of principles of reasoning and a subject of study in multiple disciplines like philosophy, mathematics, and computer science. We only consider *formal logic*, in which a *formal language* and a *formal system* represent logical truths and logical methods of reasoning. A formal language is a set of *well-formed-formulas* specified syntactically by a set of symbols and a set of rules for the formation of formulas. A formal system as explained earlier in this section consists of a formal language and a deductive apparatus. A deductive apparatus, also called a deductive system consists of a set of axioms and a set of inference rules that can be used to derive theorems of the system.

First-order predicate logic. An example of deductive system is first order predicate logic also known as first-order logic consists of a countable sets of symbols for constants, functions, and predicates, variables, and a set of standard Boolean connectives ($\wedge, \vee, \Rightarrow, \equiv$) and quantifiers (\exists, \forall). This distinguishes first-order logic from propositional logic, in which quantifiers or relations are not used. In Floyd-Hoare [34, 45] assertional methods also typically first-order predicate assertions are used. Most of the approaches that have used first-order logic actually draw upon the Hoare-style verification techniques for software.

Boyer-Moore computational logic. This logic is a restricted form of first-order logic. Boyer-Moore [16] introduced a quantifier-free first-order logic with equality. In this logic, terms composed of variables and function applications. Constants are demonstrated as functions without arguments. Logical connectives, like *not*, *and*, *or* are defined in terms of primitive logical constants *true*, *false*, and the primitive connective *if*. Equality is in the form of a function *equal* and axioms that characterize it. Through the *Shell Principle*, a user can introduce inductively constructed object type (characterized by a set of axioms). According to the shell principle, all type definitions should be accompanied by specification of a recognizer function, a constructor function, and an accessor function

for that object type. A user can also introduce axioms defining new functions. These axioms have to satisfy the *Principle of Definition*, to avoid inconsistencies caused by their addition. Using this principle, one can make sure that all new functions are defined either non-recursively in case of pre-defined functions, or there exists a well-founded ordering on some measure of the arguments that reduces with every recursive call in case of recursive definitions. In this logic, standard rules of inference used for propositional logic and the principle of induction, which is based on the same well-founded ordering used by the definition axioms, have been used for reasoning.

The Boyer-Moore theorem prover equips an automated facility for generating proofs in the described logic above. The process of proof generation though, is not totally automatic, and the user might need to assist the theorem prover for setting up intermediate lemmas and helpful definitions. The basic theory used in the system (logic plus theorem prover) is including: axiomatize natural numbers, negative integers, lists and character strings, also commands are provided for adding shells, defining new functions and proving theorems. This system is an effective tool that has a number of applications in different areas [17], due to the strong mathematical foundation and heuristics built into the system.

Higher-order logic. This logic is a form of predicate logic that has additional quantifiers and sometimes stronger semantics. Higher-order logic is more expressive than first-order logic because of the ability to quantify over predicate symbols, however its model-theoretic properties are less well-behaved than in the first-order logic.

Modal logic. This is the logic of necessities and possibilities. Modal logic is considered the development of logic of increasing ability to express change [58]. Propositional logic is the logic of absolute truths, that means in a domain, propositions are either true or false. In predicate logic, the notion of truth has been extended by making it relative, that means the truth of a predicate may depend on the related actual arguments (variables), as well as the domain (since the arguments can change over the elements in the domain of discourse). Modal logic extends this notion further by providing additional variability, that means the meaning of a predicate (or a function) symbol may also vary depending on what point it is in (this point is often called a “world”). Therefore, modal formulas

are interpreted according to a state in a universe (the universe consists of a set of states), a domain of discourse that appropriate logic symbols are interpreted by every state, as well as an accessibility relation between states. There are two basic modal operators: the necessity operator that is represented by $\Box P$ and the possibility operator that is represented by $\Diamond P$. The meaning of modality operators is as follows:

- $\Box P$: a property is true in state s if the property P is true in all states accessible from s
- $\Diamond P$: a property is true in state s if the property P is true in at least one state accessible from s

Temporal logic. This logic is a specialized form of modal logic that places additional restrictions on the accessibility relation to represent passage of time. In temporal logic, one can talk about the past, present, and future. In addition to basic modalities $\Box P$ and $\Diamond P$, there are two other operators in temporal logic: \bigcirc and U . In the temporal framework, these four operators are referred to as *Always (Henceforth, G)*, *Sometimes (Eventually, F)*, *Next-time (X)*, and *Until (U)*, respectively. The meaning of these operators is as follows:

- $\Box P$: is true in state s , if P is true in all future states accessible from s
- $\Diamond P$: is true in state s , if P is true in some future state accessible from s
- $\bigcirc P$: is true in state s , if P is true in the next state accessible from s
- $P U Q$: is true in state s , if either Q is true in s , or it is true in some future state of s , and until then P is true at every intermediate state

There are different kinds of temporal logic. If the truth of a formula is determined with respect to an interval between states, it is called *Interval Temporal Logic* [5]. If the truth of a formula is determined with respect to a state, depending on the difference in viewing the notion of time, there are other categorization. If time is categorized as a single linear sequence of events, it is called *Linear Time (Temporal) Logic (LTL)* [58], but if a branching view of time is considered, that means in any instant of

time there are a branching set of possibilities into the future, it is called *Branching Time (Temporal) Logic (BTTL)* [31].

Extended temporal logic (ETL). The original work on ETL was first suggested by Wolper [87]. He showed that in propositional version of LTL (called PTL), a property like $even(p)$ is not expressible. Therefore, he proposed to add new operators corresponding to right-linear grammars to temporal logic. To interpret grammars as operators, he established a correspondence between words generated by a grammar and sequences. The basic idea is that if there is a word w (finite or infinite) over a finite alphabet σ and an assignment of formulas to every letter of σ , a sequence satisfies the word w for the given assignment if for all i , the formula associated with the letter appearing in the i th position of the word w is true in the i th state of the sequence. Later, the original work on ETL was generalized by Wolper, Vardi, and Sistla [88].

Mu-Calculus. This is another extension framework of the temporal logic context for extending expressiveness that in addition to the usual predicate logic operators, provides operators for denoting fixed points of predicate transformers (functions from relations to relations) [20].

2.4.3.2 Automata and Language Theory

Machine equivalence. An approach that both the implementation and the specification are represented as finite-state machines. An equivalence between the corresponding machines should be approved in order to show that each behavior of the implementation satisfies the specification.

Language containment. In language containment approach a containment between the languages representing the implementation ($l(Impl)$) and the specification ($l(Spec)$) has been considered instead of an equivalence in the previous approach.

Trace conformation. In trace theory a system's behavior is modeled as a set of traces that each of them is a sequence of events.

2.4.3.3 Hybrid Formalisms

LTL and finite-state automata. This verification techniques were used by Fujita, Tanaka, and Moto-oka [35]. In this technique, a finite-

state automata (FSA) description of an implementation is given, then using traditional automata techniques can verify that it satisfies a LTL specification.

Temporal logics and Büchi automata. This techniques presented by Vardi and Wolper [85]. The basic idea of their technique is as follows: a PTL formula is interpreted according to a computation comprising of an infinite sequence of states. Every state can be entirely described by a finite set of propositions, therefore a computation can be expressed as an infinite word over the alphabet comprising of truth assignments to these propositions. A constraint on the computations can be directly translated to a constraint in the form of these infinite words. In this way, by given any PTL formula, a finite-state automaton on infinite words can be built that accepts the exact set of sequences that satisfy the formula. Temporal logic formulas here can be seen as *finite-state acceptors* of infinite words, which is in the form of a Büchi automaton (this is named after Büchi who first studied them [19]).

Approaches like *automated theorem proving*, *model checking*, and *static analysis* are automatic techniques. In automated theorem proving, the relationship between a specification of a system and an implementation is considered as a theorem in logic, that should be proved within the context of a proof calculus. By giving a description of the system, a set of logical axioms, and a set of inference rules, a system produces a formal proof from the scratch. By its very nature, theorem proving is a *deductive process* – reasoning from one or more statements to reach a logical conclusion. This raises theoretical and practical concerns regarding managing its complexity. The theorem proving approach is automatic to some degree like the form of rewriting rules, specialized tautology-checkers, etc. But most of the automated theorem provers available today are actually semi-automated and in order to guide the proof searching process they need some kind of human interaction. Theorem proving systems are very general in their applications. Logic allows representation of virtually all mathematical knowledge in the form of domain-specific theories. The ability of theorem proving that can define appropriate theories, and reason about them by a common set of inference rules, makes a unifying framework that can perform all kinds of verification tasks. This very generality makes theorem proving very complex.

Model checking [8] is an algorithmic method that determines whether

a finite-state model of a system satisfies a correctness specification (the specification or property of a system is in the form of a logical formula). This is done by means of an exhaustive search of all possible states that a system could enter during its execution. Attention in model checking is focused on a single model, and it does not require to encode incidental knowledge of the whole world, therefore the complexity of this task is more manageable compare with theorem proving. Mostly, model checking provides a clear algorithm that can be completely automatic. Typically, these algorithms also have counter-example mechanism that is useful for debugging purposes, and is important from a practical point of view. However, since these systems are not general like theorem provers, they can only work for the kind of logic and model that it is designed for. Static analysis technique will be explained later in this chapter.

Other approaches. Deductive verification, type and effect systems, confluent rewriting systems (such as SOS-style rules for defining operation semantics), relational calculus, refinement calculus that is a formal system (inspired from Hoare logic) that promoting program refinement (program refinement is the verifiable transformation of an abstract (high level) formal specification into a concrete (low level) executable program).

In this thesis, we have benefited from formal verification using a Hoare-style logic for partial correctness, and static analysis technique in order to prove our approaches.

2.4.4 *Static Program Analysis*

The availability of powerful tools is crucial for the design, implementation, and maintenance of large programs. Advanced programming environments provide many such tools like syntax-directed editors, optimizing compilers, and debuggers as well we tools for transforming programs and estimating their performance. Program analysis has a significant role in many of these tools. Although program analysis techniques are extensively used in compiler optimization, we focus on their use in program verification.

Program analyses [64, 65] are static compile-time techniques that give useful information for predicting safe and computable approximations about the dynamic behaviour of programs that will arise at runtime when

executing a program on a computer. The analyses are *static*, it means that to find the results of the analysis, the programs are *not* run at all. The analyses are *safe*, it means that the result of the analysis defines all possible runs of the program. A static analysis is said to be syntactic if it is only concerned with the grammatical structure of a program, and semantic if it involves the meaning of grammatically correct programs. Thus semantic analysis is, in general, more effective in recognizing potential problems in a program.

Traditionally, program analyses have been developed for *optimizing compilers*. They are used at different levels in the compiler including optimizations apply to the source program, optimizations apply to the various intermediate representations used inside the compiler, and optimizations that exploit the architecture of the target machine and directly improve the target code. The improvements and the associated transformations caused using these analyses can have dramatic reductions on the running time. A main application is allowing compilers to generate codes to avoid *redundant* computations for instance by reusing available results or by moving loop invariant computations out of loops, etc. One of the more recent applications is the validation of software in order to decrease the probability of malicious or unintended behaviour. Usually in these applications, information from different parts of the program should be combined.

Approaches to program analysis have many features in common. One common feature among the approaches is that to be able to remain computable they can only provide *approximate answers*. In general, program behaviour is *undecidable*. All non-trivial, and semantical properties of programs are undecidable. Thus, program analysis *efficiently* compute approximate but *sound* guarantees – guarantees that are true. Therefore, what is expected in the program analysis is producing a possibly *larger set* of possibilities than what will ever happen during the program execution. The challenge here is to not produce the safe approximation too often, otherwise the analysis might be useless. Even though program analysis does not give precise information, it gives useful information for example we can know that a value will be positive, or a variable will fit within 1 byte of storage, a variable will always hold a value of its declared type, etc.

Another common feature is that program analyses should be *semantic based*, meaning that the information that we get from the analysis can be

proved to be correct based on *semantics* of the programming language – i.e., mathematical definitions of how a program behaves.

Some of the main approaches to program analysis are: data flow analysis, constraint based analysis, abstract interpretation, and type and effect systems. In *data flow analysis* which is a traditional form of program analysis, it is common to think of a program as a graph in which nodes are elementary blocks and edges show how control might pass from one elementary block to another.

There are two kinds of analysis *intraprocedural analysis* and *interprocedural analysis*. Intraprocedural analysis only deals with simple languages without functions and procedures. In contrast in interprocedural analysis functions and procedures are also taken into account. Interprocedural analysis becomes complicated for instance when we want to ensure that calls and returns match one another.

The purpose of *control flow analysis* that is often expressed as *constraint based analysis* is to specify information concerning what “elementary blocks” may lead to what other “elementary blocks”. Control flow analysis is usually expressed as a constraint based analysis. *Constraint based analysis* is in fact one way to specify the control flow analysis using a collection of constraints. Constraints are used to determine the control flow between functions. In this analysis, the labeling scheme has been used. All *subexpressions* will be labeled. The control flow analysis is specified using a pair of functions $(\hat{C}, \hat{\rho})$, where $\hat{C}(\ell)$ contains the values that the subexpression labeled ℓ may evaluate to and $\hat{\rho}(x)$ contains the values that the variable x may be bound to.

Abstract Interpretation is a framework for computing over-approximations of a behavioural property of the program based on a fixpoint computation over a lattice representing it. Abstract Interpretation cannot reason about the exact program behavior because of undecidability, rather a conservative over-approximation will be obtained; however, this can be enough to prove program correctness.

A simple *type and effect systems* can be described as follows: a statement S maps a state to a state (if it terminates), therefore can be considered to have type $\Sigma \rightarrow \Sigma$ where Σ denotes the type of states; this is written as the judgement:

$$S : \Sigma \rightarrow \Sigma$$

A type and effect system can often be viewed as a combination of an *Effect System* and an *Annotated Type System*. In an Effect System,

judgements are usually in the form: $S : \Sigma \xrightarrow{\varphi} \Sigma$ where φ is the *effect* and tell us something about what will happen when S is executed, for example this can be which error might happen, or which file might be changed. In an Annotated Type System, judgements are usually in the form: $S : \Sigma_1 \rightarrow \Sigma_2$ where the Σ_i describes certain *properties of states*, for instance it can tell that a variable is positive or certain invariant is maintained.

In this thesis we use control flow analysis.

Summary of Research Papers and Contributions

This chapter presents research contributions of the thesis, summarizing four research papers. The chapter then returns to the research questions formulated in Section 1.2 and discusses them in connection with the contributions of this thesis. The full content of the papers appears in Part II, as in their original publications, but reformatted to fit the structure of this thesis.

We first recall the research questions documented in Section 1.2:

RQ1: What are the security and privacy functionalities for IoT systems relevant at the early design phase?

RQ2: How can we build a security and privacy functionality framework for IoT systems relevant to the early design phase?

RQ3: How can the functionality framework help to improve security and privacy of IoT systems at the early design phase?

RQ4: How can we model distributed systems at a high level of abstraction?

RQ5: How can we analyze modeling frameworks of distributed systems to identify and detect possible vulnerabilities in the models in order to improve the security and trust level at the late design phase in these systems?

RQ6: How can we improve the security, privacy and trust levels of distributed systems using a methodology at the late design phase?

RQ7: How can we use formal verification and specification for the behavior of models of distributed systems?

The following sections give a summary of each paper and explain how the paper answers above research question(s).

3.1 Paper 1

Title: Security and Privacy Functionalities in IoT
Authors: Elahe Fazeldehkordi, Olaf Owe, and Josef Noll
Publication: Presented at 17th International Conference on Privacy, Security and Trust (PST), 26–28 August 2019 (pp. 1–12). Published in IEEE.
DOI: <https://doi.org/10.1109/PST47121.2019.8949054>

Summary and contributions: This paper presents a new taxonomy framework that organizes all aspects of security and privacy baselines, guidelines, and recommendations by summarizing the most current standards and documents related to security and privacy functionalities in the setting of IoT. The paper then explores how the framework can facilitate the process of improving IoT security and privacy, in combination with a security classification method and demonstrates the application of the framework on a case study.

For the framework, this paper investigates the most important IoT-related security baselines and guidelines developed by ENISA [32, 33], OWASP [68], the Industrial Internet Consortium [81], the Cloud Security Alliance [25], and the Broadband Internet Technical Advisory Group [11], etc., as well as security and privacy guidelines from ISO [1, 2] and NIST [79], which could be relevant for securing IoT systems. The paper extracted the parts of these documents that deal with IoT and security/privacy, and unified them using a common vocabulary, and then categorized and integrated the resulting guidelines and requirements in a uniform style, and embedded them in a graphical representation by means of a tool based on diagrams called xmind.

By presenting the security and privacy functionality framework and explaining how we have developed it, we address RQ1 and RQ2. Also, this paper addresses RQ3 by combining the framework with a security classification method and demonstrating how the framework can help to improve the security and privacy of a system using a case study.

3.2 Paper 2

- Title:** Security and Privacy in IoT Systems: A Case Study of Healthcare Products
- Authors:** Elahe Fazeldehkordi, Olaf Owe, and Josef Noll
- Publication:** Presented at 13th International Symposium on Medical Information and Communication Technology (ISMICT), 8–10 May, 2019 (pp. 1–8). Published in IEEE. DOI: <https://doi.org/10.1109/ISMICT.2019.8743971>

Summary and contributions: This paper delves deeper and demonstrates the application of the framework presented in Paper 1 by a case study of healthcare products, in combination with a recent security classification method. To give an understanding of how the framework can help to improve security and privacy in practice, this paper gives a comprehensive discussion on the details of two scenarios in the case study, compares security and privacy challenges of the two scenarios, and then shows how to soften these challenges using security classifications and functionalities of the framework in Paper 1. This paper addresses RQ3.

3.3 Paper 3

- Title:** A Language-Based Approach to Prevent DDoS Attacks in Distributed Financial Agent Systems
- Authors:** Elahe Fazeldehkordi, Olaf Owe, and Toktam Ramezani-farkhani
- Publication:** Presented at Security for Financial Critical Infrastructures and Services (FINSEC), part of the 24th European Symposium on Research in Computer Security (ESORICS) Conference, 23–27 September 2019. Published in Lecture Notes in Computer Science, vol 11981 (pp. 258–277). Springer. DOI: https://doi.org/10.1007/978-3-030-42051-2_18

Summary and contributions: This paper presents an approach as an additional layer of defense in distributed agent systems to combat Denial

of Service(DoS) and Distributed DoS(DDoS) attacks. A high level concurrent object-oriented modeling framework for distributed systems, based on actor model with support of asynchronous and synchronous method interaction and futures has been considered. The language supports efficient interaction by features such as asynchronous and non-blocking method calls and first-class futures.

The approach in this paper uses static analysis to identify and prevent potential vulnerabilities caused by asynchronous communication that can directly or indirectly cause call-based flooding of agents. More precisely, a general algorithm has been adapted for detecting call flooding [69] to the setting of security analysis and for detection of distributed denial of service attacks by adding support for one-to-many and many-to-one attacks. The algorithm detects call cycles that might overflow the incoming queues of one or more communicating agents. Each cycle may involve any number of agents, possible involving the attacked agent(s).

The paper also illustrates the approach on examples of distributed systems in the financial sector, including versions of a one-to-many attack and a many-to-one attack and shows how static detection can solve the situation in each case.

By showing examples modeling distributed systems in the financial sector using object-oriented modeling frameworks based on the active object paradigm, this paper addresses RQ4. Also, by applying static analysis on these examples in order to detect and prevent potential vulnerabilities caused by asynchronous communication that can directly or indirectly cause call-based flooding of agents resulting DoS and DDoS attacks, this paper achieves aspects of RQ5.

3.4 Paper 4

- Title:** A Language-Based Approach to Smart Contracts Supporting Safety and Security
- Authors:** Elahe Fazeldehkordi and Olaf Owe
- Publication:** Submitted to the Journal of Logic and Algebraic Programming (JLAMP), April 2020, 51 pages. (under revision)

Summary and contributions: This paper presents a new approach to define lightweight smart contracts with associated history objects, supporting similar trust, immutability, and transparency of smart contracts based on blockchain but at the software level, without use of blockchain. The approach can apply to a wide range of contracts, not only financial ones, and opens up for lightweight smart contracts without the resource and energy costs of blockchain. However, it can also be combined with the blockchain technology in order to improve the overall trust level on insecure platforms. Furthermore, this approach offers better privacy control and comes with a theory for formal specification and verification.

The framework integrates trust at the language level through the notion of history objects. A history object can be used to provide safety, security, and privacy, as well as runtime checking. For each contract, a history object will record all related transactions and generated future values. Because of these recorded transactions, a history object can be seen as a ledger, but local to a given contract. Contract partners may interact with the history objects through predefined interfaces. The history objects are specially protected objects by predefined interfaces and provide read-only access from the programmer side, and increment-only access by the underlying system. They are separated from the contract provider, and can be used by contract users to check the validity of the transactions made, in a way similar to blockchain. The approach protects against tampering and fraud, each history object is immutable and corruption-proof, and each user can observe the contract behavior through the history object to ensure validity.

Moreover, the paper gives a theory for formal specification and verification of smart contracts. In particular, the approach supports class-wise verification, which is essential in open distributed systems where the contracts interact in an unknown environment. The verification is based on sequential reasoning augmented with effects on the transaction history. These advantages have been achieved by defining a version of a high level language based on the active object paradigm and interface abstraction. The language supports multiple inheritance and allows subclasses and redefinitions without behavioral restrictions from the superclasses.

Furthermore, the paper shows how formal specifications can be checked automatically by the history objects at runtime, thereby protecting users of a faulty contract provider, as well as protecting the contract provider from illegal users. In addition, it shows how to achieve security

and privacy. The paper illustrates the approach on a typical smart contract example, namely an auction system. The contract implementation has been verified and various improvements with respect to added safety, security, and privacy have been shown.

The approach may be combined with the notion of dynamic and concurrent object groups [25]. This allows a contract provider to appear as a single object to the outside while internally consisting of a number of cooperating concurrent objects.

A history objects can provide runtime checking of specified behavioral properties of the contracts. The framework allows runtime roll-backs, since the transaction history gives sufficient information to rerun a contract provider and stop at the last state before the execution of method that results in error. Reasoning about simple error recovery has been considered here.

This paper addresses RQ4 by showing examples of modeling smart contracts using a high level language based on active object paradigm. Also, this paper achieves aspects of RQ6 by demonstrating that how this modeling framework of smart contracts can support the main advantages of smart contracts based on blockchain, including trust, immutability, and transparency at the programming level and without use of blockchain. Besides, this modeling framework offers better privacy control using the notion of history objects. So it would be less expensive to implement compared with smart contracts built on blockchain that are costly with respect to time, resources, and power consumption. Furthermore, the paper explains that the approach can be combined with the blockchain technology in order to improve the overall trust level on insecure platforms. Moreover, RQ7 has been achieved in this paper by developing an imperative language for contracts, an executable functional language for writing smart contract specifications, and a theory for class-wise verification, with support of privacy, security, and model checking of contract specifications.

3.5 Additional Publications

This section lists other publications to which this thesis has contributed during her PhD research, but which are indirectly included as part of this

thesis. These papers are related to the goal of this thesis or correspond to shorter and/or preliminary versions of the work reported in this thesis.

- **Title:** Static Detection of Distributed Denial of Service Attacks in Active Object Systems

Authors: Elahe Fazeldehkordi, Olaf Owe, and Toktam Ramezani-farkhani

Publication: Technical Report of the Western Norway University (pp. 1–3). (presented at the PhD Symposium at the 15th International Conference on integrated Formal Methods (iFM), December 2–6, 2019)

- **Title:** Futures, Histories, and Smart Contracts

Authors: Elahe Fazeldehkordi and Olaf Owe

Publication: Proceedings of 31st Nordic Workshop on Programming Theory NWPT 2019 (p. 25–28),

DOI: <https://digikogu.taltech.ee/et/Download/536ca083-81f0-4f6c-9e59-c6b90b35125a>.

(presented at the 31st Nordic Workshop on Programming Theory – NWPT’19, November 13–15, 2019) (Invited paper for JLAMP journal)

- **Title:** A Framework for Flexible Program Evolution and Verification of Distributed Systems

Authors: Olaf Owe, Elahe Fazeldehkordi, and Jia-Chun Lin

Publication: In International Conference on Model-Driven Engineering and Software Development (pp. 320–349), February 20–22, 2019, Springer.

- **Title:** A Flexible Framework for Program Evolution and Verification

Authors: Olaf Owe, Jia-Chun Lin, and Elahe Fazeldehkordi

Publication: In MODELSWARD (pp. 177–189), February 20–22, 2019.

- **Title:** A Language-Based Approach to Prevent DDoS Attacks in Distributed Object Systems

Authors: Toktam Ramezanifarkhani, Elahe Fazeldehkordi, and Olaf Owe

Publication: Proceedings of the 29th Nordic Workshop on Programming Theory (p. 19–21),

DOI: <https://research.it.abo.fi/nwpt17/proceedings/NWPT2017proceedings.pdf>

(presented at the 29th Nordic Workshop on Programming Theory - NWPT'17, November 1–3, 2017)

- **Title:** Hoare-Style Reasoning from Multiple Contracts

Authors: Olaf Owe, Toktam Ramezanifarkhani, and Elahe Fazeldehkordi

Publication: In the International Conference on integrated Formal Methods (pp. 263–278), September 18–22, 2017, Springer.

- **Title:** Security Functionality of IoT Devices

Authors: Elahe Fazeldehkordi, Olaf Owe, and Toktam Ramezanifarkhani

Publication: Proceedings of the PhD Symposium at iFM'17 of Formal Methods: Algorithms, Tools, and Applications (pp. 44–47),

DOI: <https://www.duo.uio.no/handle/10852/57814>

(presented at PhD Symposium at the 13th International Conference on integrated Formal Methods (iFM), September 18–22, 2017)

Conclusions and Future Work

4.1 *Conclusions*

The main goal of this thesis is to improve security and privacy of IoT and distributed systems during the design phase. From this goal, two subgoals are identified that resulted in seven research questions presented in Section 1.2. This thesis has contributed to presenting four research papers which address the research questions and the related subgoals.

The focus of this thesis is on two aspects of the design phase. First, the early design phase, where the architecture of a system is being developed, considering the setting of IoT systems. Second, the late design phase, where models, in particular executable models and prototypes are designed, considering IoT systems and also the more general setting of distributed systems. Having sufficient information at the architectural design phase and using graphical notations make it easier to understand the system components and their specifications for the modeling phase.

The methodology proposed in this thesis meets the requirements stated in Section 1.3. We now revisit these requirements and their rationale:

Model-based approach. The methodology proposed in this thesis uses a modeling language which is executable in order to allow prototyping and verification of specification requirements of the system, so that it defines a program structure that can be a guideline for the implementation phase.

Formality. The methodology proposed in this thesis uses a language formalized by means of a formal syntax and an operational semantics given in the structural operational semantics (SOS) style that supports rigorous, abstractions, and formal analysis.

The paradigm of concurrent and distributed active objects. The methodology proposed in this thesis builds on the Creol/ABS language, this methodology uses a concurrent and distributed

object-oriented framework. Object orientation is the leading paradigm used in the software industry today and is suitable for open distributed systems [78]. The active object paradigm provides a natural way of modeling distributed systems in general, and IoT systems in particular, because it covers fundamental aspects of IoT systems, such as distribution of concurrent units communicating by message passing, where each unit can run on a device with limited processing power and limited storage [48]. In order to make reasoning as simple and powerful as possible, we focus on languages at a high level of abstraction rather than low level languages. We consider executable and imperative languages supporting object-oriented principles, and have designed the languages to reduce the complexity of further translation to low level languages for IoT and distributed systems.

In the following, we outline how the methodologies presented in this thesis address security and privacy solutions during the design phase by answering the research questions.

RQ1: What are the security and privacy functionalities for IoT systems relevant at the early design phase?

Paper 1 collects security and privacy functionalities in the setting of IoT by summarizing the most current related documents.

RQ2: How can we build a security and privacy functionality framework for IoT systems relevant to the early design phase?

Paper 1 presents a new taxonomy framework that organizes all aspects of security and privacy baselines, guidelines, and recommendations. The functionalities are separated into two major parts, the *life cycle* aspects of a system and the *management* aspects of security and privacy. The life cycle in this paper relates to the different phases in the life of a system, while the management of security and privacy is the ability to put supporting functionality elements in the system.

RQ3: How can the functionality framework help to improve security and privacy of IoT systems at the early design phase?

To give an understanding of how the framework can help to improve the security and privacy of IoT products and to facilitate developing and designing secure and privacy-aware IoT systems, Paper 1 combines the taxonomy framework with a security classification method and demonstrates the usefulness by a case study. Paper 2 delves deeper and demonstrates the application of the framework described in Paper 1 using a case study of healthcare products, in combination with a recent security classification method. Paper 2 gives a comprehensive discussion on the details of two scenarios in the case study, compares security and privacy challenges of the two scenarios, and then shows how to soften these challenges using the security classifications and functionalities of the framework. The analysis results in Paper 1 and 2 demonstrate that using the security and privacy functionality framework, the security of the overall system has improved.

RQ4: How can we model distributed systems at a high level of abstraction?

Paper 3 considers a high level object-oriented modeling framework for distributed systems based on active objects. This setting is appealing in that it naturally supports the distribution of autonomous concurrent units, and efficient interaction, avoiding active waiting and low level synchronization primitives such as explicit signaling and lock operations. It is therefore useful as a framework for modeling and analysis of distributed systems. The language in this paper supports efficient interaction by features such as asynchronous and non-blocking method calls and first-class futures, which are popular features applied in many distributed systems today. However, the DDoS detection method works better for local futures, since then the association between call statements and get statements can be done with less over-approximation, and the analysis of this is modular, in contrast to the setting of first-class futures where get statements may associate to call statements in other objects, including objects not known or not part of the considered subsystem. The DDoS detection method is therefore well suited for IoT systems with local futures.

Paper 4 considers a high level object-oriented modeling language that allows reasoning support. The language gives fewer runtime errors and less need for roll-backs, which simplifies reasoning. In particular, the

approach in paper 4 supports class-wise verification. This is essential for scalability and open-ended program development, which are highly relevant factors for contracts. The language is based on the active object paradigm. This paradigm offers modular semantics, which is essential when we turn to specification and verification issues. Clients, contract, and history objects in this paper are then described by concurrent and distributed objects. The language builds on the principle of interface abstraction, i.e., remote field access is illegal, and an object can only be accessed through an interface. Each object has one or more interfaces, and the only possible way of object interaction is through the methods defined in the corresponding interfaces. The language combines first-class futures, which is often used in active object languages, and a restricted version of cooperative scheduling. This novel combination gives flexible method interaction, scheduling control, simplified reproducibility of executions, as well as simplified verification. In particular, the restricted cooperative scheduling implies that each object is deterministic, relative to its inputs (i.e., invocation messages), something which is essential for enabling roll-backs. Furthermore, the futures (related to a contract) are now collected in the history objects, for which we have suggested mechanisms for security and privacy control. These futures will not be garbage collected since they are used in the calculations of roll-backs and in the functionality of the history object interfaces.

RQ5: How can we analyze modeling frameworks of distributed systems to identify and detect possible vulnerabilities in the models in order to improve the security and trust level at the late design phase in these systems?

Paper 3 proposes an approach consisting of static analysis to identify and prevent potential vulnerabilities caused by asynchronous communication including call-based DoS or DDoS attacks, possibly involving a large number of distributed actors. The paper illustrates the approach on examples of distributed systems in the financial sector, including versions of a one-to-many attack and a many-to-one attack.

RQ6: How can we improve the security, privacy and trust levels of distributed systems using a methodology at the late design phase?

Paper 4 proposes a new approach to define smart contracts, offering trust at the software level. This approach can apply to a wide range of contracts, not only financial ones, and opens up for lightweight smart contracts without the resource and energy costs of blockchain. The framework in paper 4 integrates trust at the language level through the notion of history objects. For each contract, a history object will record all related transactions. The history objects are specially protected objects, with read-only access at the programming language level. Contract partners may interact with the history objects through predefined interfaces. The paper shows that a history object can be used to provide safety, security, and privacy, as well as runtime checking. A history object can provide runtime checking of specified behavioral properties of the contracts. The paper demonstrates the approach on a smart contract example, namely an auction system.

RQ7: How can we use formal verification and specification for the behavior of models of distributed systems?

Paper 4 presents a framework integrating the notion of history objects giving rise to lightweight smart contracts, by developing an imperative language for contracts, an executable functional language for writing smart contract specifications by means of invariants referring to the transaction history of a contract, and a theory for class-wise verification, with support of privacy, security, and model checking of contract specifications.

4.1.1 Limitations

Our focus has been on the design phase of the system development life cycle, which means that planning, requirement analysis, implementation, and testing phases have not been considered. We have only considered two parts of the design phase, namely the early design phase, where the architecture of a system is being developed, and the late design phase – by means of executable modeling of systems using the active object concurrency model. However, the modeling paradigm considered allows high level implementation, prototyping, testing, as well as program analysis, at the abstraction level of the model.

We have considered only some approaches and some security problems. In the early design phase, we have only considered simple

forms of graphical architecture models in the setting of IoT systems. And at the late design phase we have only looked at the active object setting and have only considered DDoS attack and smart contracts. Security and privacy issues specifically for IoT models have not been considered. However, our analysis framework for DDoS attacks is specifically useful for IoT model where first-class future are avoided.

In the next section, we will look at the shortcomings of each research paper one by one and give some suggestions for possible extensions and future research directions.

4.2 Future Work

We organize the discussion in this section according to the research papers of the thesis and the shortcomings of each paper, this makes it easier to identify possible research directions, also makes it more clear to the reader how the suggested work relates to the papers in this thesis.

Security and Privacy Functionalities in IoT: The framework in this paper is a preliminary classification of privacy and security functionalities based on the available recommendations and standards for IoT systems. This should imply that all aspects are covered, but there is no guarantee for that. This can be seen as a limitation of the work. Secondly, the application of the methodology is ultimately depending on the judgements of software engineers or security experts, and is therefore not 100% precise. If their judgement is wrong, for instance if they choose the wrong connectivity or protection level, it will in general give a wrong estimate. One possible extension would be to allow the framework to be complemented with even more elements and look in more detail at GDPR for privacy functionalities.

Security and Privacy in IoT Systems: A Case Study of Healthcare Products: This paper only shows the application of our framework on one case study of healthcare products. The applications of our framework can be extended to other IoT domains and case studies with several kinds of IoT devices and sensors involved can be considered. This would be a valuable step toward validating the framework presented in this thesis.

A Language-Based Approach to Prevent DDoS Attacks in Distributed Financial Agent Systems: In this paper we used static analysis. Our static detection will overapproximate the set of (potentially) un-

completed calls. In future work, we suggest to complement the static checking with dynamic runtime checking since static detection methods give a degree of over-estimation. This could give a more precise combined detection strategy.

An Approach to Smart Contracts Supporting Safety and Security:

In the current state, our framework can be used for modeling, prototyping, analysis, verification, and model checking of smart contracts. In future work, further mechanisms for error and exception handling can be added and efficient implementation of history objects and roll-backs can be considered.

Bibliography

- [1] 27000, I. *Information technology — Security techniques — Information security management systems — Overview and vocabulary (fourth edition)*. ISO/IEC 2016. 2016.
- [2] 27001, I. *INTERNATIONAL STANDARD ISO/IEC 27001 Information technology—Security techniques —Information security management systems —Requirements*. ISO/IEC 2013. 2013.
- [3] Agha, G. *Actors: A Model of Concurrent Computation in Distributed Systems*. Cambridge, MA, USA: MIT Press, 1986.
- [4] Ahrendt, W., Pace, G. J., and Schneider, G. “Smart contracts: a killer application for deductive source code verification”. In: *Principled Software Development*. Springer, 2018, pp. 1–18.
- [5] Allen, J. F. and Ferguson, G. “Actions and events in interval temporal logic”. In: *Journal of Logic and Computation* vol. 4, no. 5 (1994), pp. 531–579.
- [6] Andrews, G. R. *Foundations of multithreaded, parallel, and distributed programming*. Vol. 11. Addison-Wesley Reading, 2000.
- [7] Armstrong, J. et al. “Concurrent programming in ERLANG”. In: (1993).
- [8] Baier, C. and Katoen, J.-P. *Principles of model checking*. MIT press, 2008.
- [9] Baker, H. and Hewitt, C. “Laws for communicating parallel processes”. In: *MIT Artificial Intelligence Laboratory* (1977).
- [10] Beydeda, S., Book, M., Gruhn, V., et al. *Model-driven software development*. Vol. 15. Heidelberg: Springer, 2005.
- [11] BITAG. *Internet of Things (IoT) Security and Privacy Recommendations*. [https://www.bitag.org/documents/BITAG_Report_-_Internet_of_Things_\(IoT\)_Security_and_Privacy_Recommendations.pdf](https://www.bitag.org/documents/BITAG_Report_-_Internet_of_Things_(IoT)_Security_and_Privacy_Recommendations.pdf). 2016.

- [12] Bjørk, J. et al. “A type-safe model of adaptive object groups”. In: *arXiv preprint arXiv:1208.4630* (2012).
- [13] Boer, F. D. et al. “A Survey of Active Object Languages”. In: *ACM Comput. Surv.* vol. 50, no. 5 (Oct. 2017), pp. 1–39.
- [14] Boer, F. S. de, Clarke, D., and Johnsen, E. B. “A Complete Guide to the Future”. In: *Programming Languages and Systems*. Ed. by De Nicola, R. Berlin, Heidelberg: Springer Berlin Heidelberg, 2007, pp. 316–330.
- [15] Bowman, H. and Derrick, J. “Issues in Formal Methods (chapter 3)”. In: *Formal Methods for Distributed Processing, A Survey of Object-oriented Approaches*. Ed. by Bowman, H. and Derrick, J. Cambridge, UK: Cambridge University Press, Sept. 2001, pp. 18–35.
- [16] Boyer, R. S. and Moore, J. S. “A Computational Logic Handbook, volume 23 of”. In: *Perspectives in computing* (1988).
- [17] Boyer, R. S. and Moore, J. S. *Proof-Checking, Theorem Proving, and Program Verification*. Tech. rep. Texas Univ at Austin Inst for Computing Science and Computer Applications, 1983.
- [18] Broy, M. “Distributed Concurrent Object-Oriented Software”. In: *From Object-Oriented to Formal Methods: Essays in Memory of Ole-Johan Dahl*. Ed. by Owe, O., Kroghdahl, S., and Lyche, T. Berlin, Heidelberg: Springer Berlin Heidelberg, 2004, pp. 83–95.
- [19] Büchi, J. R. “On a Decision Method in Restricted Second Order Arithmetic”. In: *The Collected Works of J. Richard Büchi*. Ed. by Mac Lane, S. and Siefkes, D. New York, NY: Springer New York, 1990, pp. 425–435.
- [20] Burch, J. et al. “Symbolic model checking: 1020 States and beyond”. In: *Information and Computation* vol. 98, no. 2 (1992), pp. 142–170.
- [21] Byres, E. and Lowe, J. “The myths and facts behind cyber security risks for industrial control systems”. In: *Proceedings of the VDE Kongress*. Vol. 116. 2004, pp. 213–218.

-
- [22] Clarke, D., Johnsen, E. B., and Owe, O. “Concurrent Objects à la Carte”. In: *Concurrency, Compositionality, and Correctness: Essays in Honor of Willem-Paul de Roever*. Ed. by Dams, D., Hannemann, U., and Steffen, M. Berlin, Heidelberg: Springer Berlin Heidelberg, 2010, pp. 185–206.
- [23] Clavel, M. et al. “Maude: specification and programming in rewriting logic”. In: *Theoretical Computer Science* vol. 285, no. 2 (2002). *Rewriting Logic and its Applications*, pp. 187–243.
- [24] Clavel, M. et al. *All About Maude – A High-Performance Logical Framework: How to Specify, Program, and Verify Systems in Rewriting Logic*. Vol. 4350. Springer, 2007.
- [25] *Future-proofing the Connected World: 13 Steps to Developing Secure IoT Products*. <https://downloads.cloudsecurityalliance.org/assets/research/internet-of-things/future-proofing-the-connected-world.pdf>. 2016.
- [26] Dahl, O.-J. “The Roots of Object Orientation: The Simula Language”. In: *Software Pioneers: Contributions to Software Engineering*. Ed. by Broy, M. and Denert, E. Berlin, Heidelberg: Springer Berlin Heidelberg, 2002, pp. 78–90.
- [27] Dahl, O.-J. *Verifiable programming*. Prentice Hall PTR, 1992.
- [28] Dahl, O.-J., Myhrhaug, B., and Nygaard, K. *Common base language*. Norsk Regnesentral, 1970.
- [29] Deogirikar, J. and Vidhate, A. “Security attacks in IoT: A survey”. In: *2017 International Conference on I-SMAC (IoT in Social, Mobile, Analytics and Cloud)(I-SMAC)*. IEEE. 2017, pp. 32–37.
- [30] Dershowitz, N. *Software horror stories*. URL: <https://www.cs.tau.ac.il/~nachumd/horror.html>.
- [31] Emerson, E. A. and Halpern, J. Y. ““Sometimes” and “not never” revisited: on branching versus linear time temporal logic”. In: *Journal of the ACM (JACM)* vol. 33, no. 1 (1986), pp. 151–178.
- [32] ENISA. *Baseline Security Recommendations for IoT in the context of Critical Information Infrastructures*. European Union Agency for Network and Inf. Sec. 2017.

- [33] ENISA. *Guidelines for SMEs on the security of personal data processing*. <https://www.enisa.europa.eu/publications/guidelines-for-smes-on-the-security-of-personal-data-processing>. 2017.
- [34] Floyd67, R. and Floyd, R. “Assigning meaning to programs”. In: *Mathematical Aspects of Computer Science* vol. 19 (), pp. 19–31.
- [35] Fujita, M., Tanaka, H., and Moto-Oka, T. “Logic design assistance with temporal logic”. PhD thesis. University of Tokyo, 1984.
- [36] Gallier, J. H. *Logic for computer science: foundations of automatic theorem proving*. Courier Dover Publications, 2015.
- [37] Gosling, J. et al. *The Java language specification*. Addison-Wesley Professional, 2000.
- [38] Gupta, A. “Formal hardware verification methods: A survey”. In: *Formal Methods in System Design* vol. 1, no. 2-3 (1992), pp. 151–238.
- [39] Hajdu, Á. and Jovanović, D. “SMT-Friendly Formalization of the Solidity Memory Model”. In: *arXiv preprint arXiv:2001.03256* (2020).
- [40] Henderson-Sellers, B. *A book of object-oriented knowledge: an introduction to object-oriented software engineering*. Prentice-Hall, Inc., 1996.
- [41] Hernandez, G. et al. “Smart nest thermostat: A smart spy in your home”. In: *Black Hat USA*, no. 2015 (2014).
- [42] Hewitt, C. “Viewing control structures as patterns of passing messages”. In: *Artificial intelligence* vol. 8, no. 3 (1977), pp. 323–364.
- [43] Hewitt, C., Bishop, P., and Steiger, R. “Session 8 formalisms for artificial intelligence a universal modular actor formalism for artificial intelligence”. In: *Advance Papers of the Conference*. Vol. 3. Stanford Research Institute. 1973, p. 235.
- [44] Hoare, C. A. R. “Towards a Theory of Parallel Programming”. In: *The Origin of Concurrent Programming: From Semaphores to Remote Procedure Calls*. Ed. by Hansen, P. B. New York, NY: Springer New York, 2002, pp. 231–244.

-
- [45] Hoare, C. A. R. “An axiomatic basis for computer programming”. In: *Communications of the ACM* vol. 12, no. 10 (1969), pp. 576–580.
- [46] Igarashi, A., Pierce, B. C., and Wadler, P. “Featherweight Java: a minimal core calculus for Java and GJ”. In: *ACM Transactions on Programming Languages and Systems (TOPLAS)* vol. 23, no. 3 (2001), pp. 396–450.
- [47] Illera, A. G. and Vidal, J. V. “Lights off! The darkness of the smart meters”. In: *BlackHat Europe* (2014).
- [48] Johnsen, E. B. and Owe, O. “An asynchronous communication model for distributed concurrent objects”. In: *Software & Systems Modeling* vol. 6, no. 1 (2007), pp. 39–58.
- [49] Johnsen, E. B., Owe, O., and Yu, I. C. “Creol: A type-safe object-oriented model for distributed concurrent systems”. In: *Theoretical Computer Science* vol. 365, no. 1 (2006). *Formal Methods for Components and Objects*, pp. 23–66.
- [50] Johnsen, E. B. et al. “A formal model of service-oriented dynamic object groups”. In: *Science of Computer Programming* vol. 115–116 (2016). *Special Section on Foundations of Coordination Languages and Software (FOCLASA 2012) Special Section on Foundations of Coordination Languages and Software (FOCLASA 2013)*, pp. 3–22.
- [51] Johnsen, E. B. et al. “ABS: A Core Language for Abstract Behavioral Specification”. In: *Formal Methods for Components and Objects*. Ed. by Aichernig, B. K., Boer, F. S. de, and Bonsangue, M. M. Berlin, Heidelberg: Springer Berlin Heidelberg, 2012, pp. 142–164.
- [52] Johnsen, E. B. et al. “An Object-Oriented Component Model for Heterogeneous Nets”. In: *Formal Methods for Components and Objects*. Ed. by Boer, F. S. de et al. Berlin, Heidelberg: Springer Berlin Heidelberg, 2008, pp. 257–279.
- [53] Johnsen, E. B. et al. “Specification of distributed systems with a combination of graphical and formal languages”. In: *Proceedings Eighth Asia-Pacific Software Engineering Conference*. IEEE, 2001, pp. 105–108.

- [54] Kabay, M. “Attacks on Power Systems: Hackers, Malware”. In: *Norwich University* (2010).
- [55] Koscher, K. et al. “Experimental security analysis of a modern automobile”. In: *2010 IEEE Symposium on Security and Privacy*. IEEE. 2010, pp. 447–462.
- [56] Ledgard, H. F. *The little book of object-oriented programming*. Prentice-Hall, Inc., 1995.
- [57] Levy, E. “Crossover: online pests plaguing the off line world”. In: *IEEE Security & Privacy* vol. 1, no. 6 (2003), pp. 71–73.
- [58] Manna, Z. and Pnueli, A. *Verification of Concurrent Programs: The Temporal Framework; by Zohar Manna and Amir Pnueli*. Department of Computer Science, Stanford University, 1981.
- [59] Meseguer, J. “Conditional rewriting logic as a unified model of concurrency”. In: *Theoretical computer science* vol. 96, no. 1 (1992), pp. 73–155.
- [60] Meyer, B. *Object-oriented software construction*. Vol. 2. Prentice hall Englewood Cliffs, 1997.
- [61] Miller, B. and Rowe, D. “A Survey SCADA of and Critical Infrastructure Incident”. In: *Proceedings of the 1st Annual Conference on Research in Information Technology*. RIIT '12. Calgary, Alberta, Canada: Association for Computing Machinery, 2012, pp. 51–56.
- [62] Miller, C. and Valasek, C. “A survey of remote automotive attack surfaces”. In: *black hat USA* vol. 2014 (2014), p. 94.
- [63] Nawir, M. et al. “Internet of Things (IoT): Taxonomy of security attacks”. In: *2016 3rd International Conference on Electronic Design (ICED)*. IEEE. 2016, pp. 321–326.
- [64] Nielson, F., Nielson, H. R., and Hankin, C. *Principles of program analysis*. Springer, 2015.
- [65] Nielson, H. R. and Nielson, F. *Semantics with applications: an appetizer*. Springer Science & Business Media, 2007.
- [66] Nygaard, K. and Dahl, O.-J. “The development of the SIMULA languages”. In: *History of programming languages*. 1978, pp. 439–480.

-
- [67] Odersky, M., Spoon, L., and Venners, B. *Programming in scala*. Artima Inc, 2008.
- [68] OWASP. *IoT Security Guidance*.
- [69] Owe, O. and McDowell, C. “On detecting over-eager concurrency in asynchronously communicating concurrent object systems”. In: *Journal of Logical and Algebraic Methods in Programming* vol. 90 (2017), pp. 158–175.
- [70] Owe, O. and Ryl, I. “On combining object orientation, openness and reliability”. In: *Norwegian Informatics Conference*. 1999.
- [71] Owicki, S. “Axiomatic proof techniques for parallel programs. Dept. of Computer Science, Cornell University”. PhD thesis. 1975.
- [72] Owicki, S. and Gries, D. “An axiomatic proof technique for parallel programs I”. In: *Acta Informatica* vol. 6, no. 4 (1976), pp. 319–340.
- [73] Owicki, S. and Gries, D. “Verifying properties of parallel programs: An axiomatic approach”. In: *Communications of the ACM* vol. 19, no. 5 (1976), pp. 279–285.
- [74] Pal, D. and Vain, J. “Model Based Test Framework for Communications-Critical Internet of Things Systems”. In: *Databases and Information Systems X: Selected Papers from the Thirteenth International Baltic Conference, DB&IS 2018*. Vol. 315. IOS Press. 2019, p. 79.
- [75] Pfleeger, C. P. *Security in computing*. Pearson Education India, 2009.
- [76] Pollet, J. and Cummins, J. “Electricity for free? the dirty underbelly of scada and smart meters”. In: *Proceedings of Black Hat USA* vol. 2010 (2010).
- [77] Poulsen, K. “Slammer worm crashed Ohio nuke plant network”. In: <http://www.securityfocus.com/news/6767> (2003).
- [78] Raymond, K. “Reference Model of Open Distributed Processing (RM-ODP): Introduction”. In: *Open Distributed Processing: Experiences with distributed environments. Proceedings of the third IFIP TC 6/WG 6.1 international conference on open distributed processing, 1994*. Ed. by Raymond, K. and Armstrong, L. Boston, MA: Springer US, 1995, pp. 3–14.

- [79] Ross, R. S. et al. *Protecting Controlled Unclassified Information in Nonfederal Information Systems and Organizations [including updates as of 02-20-2018]*. Tech. rep. National Institute of Standards and Technology, 2018.
- [80] Schäfer, J. and Poetzsch-Heffter, A. “JCoBox: Generalizing Active Objects to Concurrent Components”. In: *ECOOP 2010 – Object-Oriented Programming*. Ed. by D’Hondt, T. Berlin, Heidelberg: Springer Berlin Heidelberg, 2010, pp. 275–299.
- [81] Schrecker, S. and al., et. *Industrial Internet of Things Volume G4: Security Framework*. Industrial Internet Consortium. https://www.iiconsortium.org/pdf/IIC_PUB_G4_V1.00_PB.pdf. 2016.
- [82] Sommerville, I. “Software engineering 9th Edition”. In: *ISBN-10* vol. 137035152 (2011), p. 18.
- [83] Stallings, W. et al. *Computer security: principles and practice*. Pearson Education Upper Saddle River, NJ, USA, 2012.
- [84] Union, E. *The General Data Protection Regulation (GDPR)*. URL: https://ec.europa.eu/info/law/law-topic/data-protection_en (visited on 05/24/2016).
- [85] Vardi, M. Y. and Wolper, P. “An automata-theoretic approach to automatic program verification”. In: *Proceedings of the First Symposium on Logic in Computer Science*. IEEE Computer Society. 1986, pp. 322–331.
- [86] Vijayan, J. “Stuxnet renews power grid security concerns”. In: *Computerworld* vol. 26 (2010).
- [87] Wolper, P. “Temporal logic can be more expressive”. In: *Information and Control* vol. 56, no. 1 (1983), pp. 72–99.
- [88] Wolper, P., Vardi, M. Y., and Sistla, A. P. “Reasoning about infinite computation paths”. In: *24th Annual Symposium on Foundations of Computer Science (sfcs 1983)*. IEEE. 1983, pp. 185–194.
- [89] Woodcock, J. et al. “Formal methods: Practice and experience”. In: *ACM computing surveys (CSUR)* vol. 41, no. 4 (2009), pp. 1–36.
- [90] Yadav, N. *Know More About the System Development Life Cycle*. URL: <https://darkbears.com/blog/know-more-about-the-system-development-life-cycle/> (visited on 03/29/2019).

- [91] Yonezawa, A. *ABCL: An Object-Oriented Concurrent System*. Cambridge, MA, USA: MIT Press, 1990.
- [92] Yonezawa, A., Briot, J.-P., and Shibayama, E. “Object-oriented concurrent programming in ABCL/1”. In: *ACM SIGPLAN Notices* vol. 21, no. 11 (1986), pp. 258–268.
- [93] Yourdon, E. et al. *Mainstream objects: an analysis and design approach for business*. Yourdon Press, 1995.
- [94] Zakrzewski, J. “Towards Verification of Ethereum Smart Contracts: A Formalization of Core of Solidity”. In: *Verified Software. Theories, Tools, and Experiments*. Ed. by Piskac, R. and Rümmer, P. Cham: Springer International Publishing, 2018, pp. 229–247.

Part II

Scientific Contributions

Paper 1: **Security and Privacy Functionalities in IoT**

Authors: Elahe Fazeldehkordi, Olaf Owe, Josef Noll

Publication: Presented at 17th International Conference on Privacy, Security and Trust (PST), 26-28 August 2019 (pp. 1-12). Published in IEEE. DOI: <https://doi.org/10.1109/PST47121.2019.8949054>

Abstract

Internet of Things (IoT) offers a variety of technologies for connecting different kinds of heterogeneous devices. Security and privacy are becoming the main issue for IoT systems and their developers. Nevertheless, most works on IoT security and privacy requirements look at these requirements from a high-level view. Hence, the essential aspects of security and privacy functionalities will be disregarded, causing wrong design decisions. To combat this problem, this paper summarizes the most current documents related to security and privacy functionalities in the setting of IoT and provides a new taxonomy framework that organizes all aspects of security and privacy baselines, guidelines, and recommendations. To give an understanding of how the framework can help to improve security and privacy of IoT products, we combine it with a security classification method and demonstrate the usefulness by a case study of health products. Our framework can serve as a cornerstone towards the development of appropriate security solutions.

5.1 Introduction

Internet of Things (IoT) represents the concept of information flow among different kinds of embedded computing devices interconnected through the internet. The aim of IoT is to provide an advanced mode

of communication among the various systems and devices, and also to facilitate the interaction between humans and the virtual world. With this aim, IoT plays a significant role in the modern society and has applications in almost all fields including healthcare systems, automobile, industrial appliances, sports, homes, entertainments, environmental monitoring etc. IoT devices have already outnumbered the number of people at a computerized workplaces, and by 2020, connected “things” based on IoT will be around 212 billion [18, 22]. Those “things” will be daily used appliances like smart-phones, smart-watches, smart television, smart refrigerators, and others. As a result of this expansion, and as most things are connecting to the internet for exchanging information, IoT is vulnerable to various security issues and attacks (e.g., man in the middle attack, eavesdropping attack, denial of service attack, access attack, as well as major privacy concerns for the end users). Despite the advance abilities provided by IoT in the data communication area, its vulnerability implications from a security and privacy standpoint are still of great concern. Therefore appropriate steps in the initial phase of design and development of IoT systems should be taken.

In this paper, we focus on the comprehensive view of the state-of-the-art concerning security and privacy functionalities and requirements for IoT systems. Consequently, we suggest a complementary methodology for analyzing the functionalities in a comprehensive framework that can help both providers and consumers of IoT devices to have a better understanding of the security and privacy aspects. By functionality we mean: “The security and privacy-related features, functions, mechanisms, services, procedures, and architectures implemented within organizational information systems or the environments in which those systems operate” [24]. We explore how the framework can facilitate the process of improving IoT security and privacy, in combination with the security classification method suggested in [4, 26]. Security classification of a system will lead to a better understanding of the value of security and promote the extra cost of securing a system. An IoT system includes different devices, and protecting all of these devices at the same level is costly. It is economically impractical to employ all the security protection mechanisms for all the devices in a system. Dividing security into different classifications is necessary, to secure IoT systems to an appropriate level.

One attractive application area of IoT is health care [16, 21], where

IoT devices are becoming common. Medical applications like remote health monitoring, fitness programs, chronic diseases, elderly care, compliance with treatment and medication at home and by health-care providers are some of the important potential applications that can be facilitated by IoT. IoT-based health-care services can help to reduce costs, increase the quality of life, and enrich the users' experience. Therefore, we choose to demonstrate the framework on a case study concerning health products. Through the development of this framework together with security classification, extensive attention has been given to the requirements and limitations for securing IoT systems.

For this framework, we have investigated the most important IoT-related security baselines and guidelines developed by ENISA [10, 11], OWASP [19], Industrial Internet Consortium [25], Cloud Security Alliance[8], and Broadband Internet Technical Advisory Group [6], etc., as well as security and privacy guidelines from ISO [1, 2] and NIST [24], which could be relevant for securing IoT systems. With respect to privacy, the European Union (EU) has passed the general data protection regulation (GDPR), which regulates who can access private data, how and for what purpose, based on the consent of the data subject [11]. We have extracted the parts of these documents that deal with IoT and security/privacy, and then we have unified them using a common vocabulary, and have then categorized and integrated the resulting guidelines and requirements in a uniform style, and embedded them in a graphical representation by means of a tool based on diagrams.

Having a comprehensive view and taxonomy of security and privacy requirements and functionalities in IoT is a prerequisite for architecting optimal security solutions, designing, and developing secure and privacy-aware IoT systems. To give an understanding of how the framework can help to improve security and privacy in practice, we combine the framework with the security classification method of [4, 26], and demonstrate how the combined methodology can be used on a case study of health products.

The contribution of this paper is to present a new functionality framework for security and privacy of IoT systems, as outlined above, and show how it can be combined with the security classification method to analyse and evaluate the security and privacy weaknesses of IoT systems, and to reduce these weaknesses, as demonstrated in the case study. Systems are often made without the help of security and privacy

experts. Our framework is easy to follow, even for non-experts. The case study shows that by following our guidelines, one can detect security problems and get help in avoiding them.

The remainder of this paper is structured as follows: Section II explains IoT-related standards and guidelines. Section III provides related work. Section IV introduces the IoT security and privacy functionality framework. Section V explains the security classification method. Section VI describes the pacemaker case study, and Section VII concludes the paper.

5.2 IoT-Related Standards and Guidelines

Cloud Security Alliance [8] has provided considerations and guidance for designing and developing secure IoT devices. It aims to reduce some of the more common issues that can be found in the development of IoT devices. A number of activities that will enable a development organization to begin enhancing the security state of IoT devices have been outlined. This document has provided a graphical view of the steps needed in order to develop more secure IoT devices. Although IoT systems are complex, including devices, gateways, mobile applications, appliances, web services, datastores, analytics systems and more, the focus of this guidance is mainly on the “devices”. In contrast, our work summarizes security and privacy functionalities considerations in the whole range of IoT system.

A security framework has been presented in Industrial Internet Consortium [25] which comprises of six interacting building blocks. These building blocks are organised into three layers. The top layer includes four core security functions, which are supported by a data protection layer and a system-wide security model and policy layer. The four core security functions are: endpoint protection, communication and connectivity protection, security monitoring and analysis, and security configuration management. And then they break down each layer into related key functions and explain the responsibility for each function. This document explains and positions security or related architectures, designs, and technologies. It also identifies procedures relevant to trustworthy Industrial Internet of Things (IIoT) systems. Security characteristics, technologies, and techniques that should be applied, and methods for




Figure 5.1: IoT security and privacy functionality framework

addressing security, have been described. However, it lacks some of the security functionalities, and in particular, it does not focus on privacy issues. The layer structure is a bit complicated, and we believe our framework has a more relaxed structure to use.

The ISO [1, 2] and NIST [24] standards are general requirements for establishing, implementing, maintaining and continually improving an information security management system, and protecting the confidentiality of Controlled Unclassified Information (CUI), respectively. We have extracted the IoT-related requirements from these standards and included these parts in our framework, while combining them with baselines and guidelines from other IoT-related documents, including OWASP and ENISA discussed below.

OWASP [19] presents guidance at a basic level, giving builders of IoT products a basic set of guidelines to consider from their perspective. The idea is that ensuring that these fundamentals are covered, will significantly improve the security of any IoT product. ENISA [10] elaborates baseline cybersecurity recommendations for IoT with a focus on Critical Information Infrastructures, which encompass facilities, networks, services, and physical and information technology equipment. Both of these guidelines, OWASP and ENISA, are addressing IoT, but both are presented in textual form, without defining a framework.

5.3 Related Work

In a work presented by Sicari et al. [27], the most relevant available solutions regarding security, privacy, and trust in IoT have been analyzed. Proposals related to security middleware, secure solutions for mobile devices, and ongoing international projects on this subject, have been

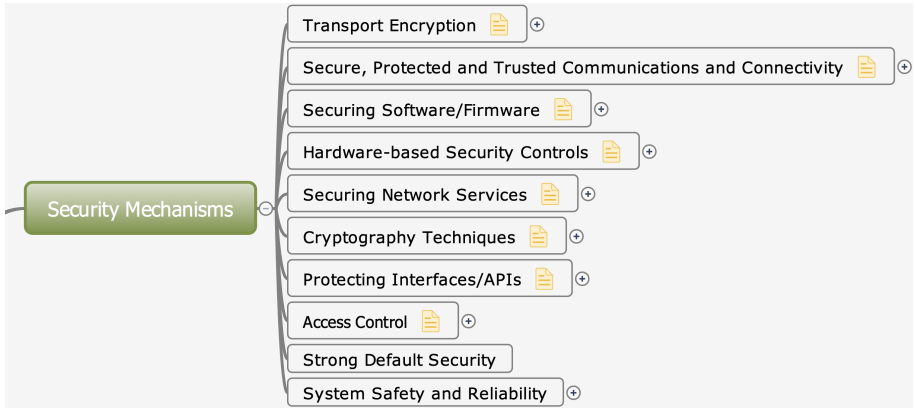


Figure 5.2: Security mechanisms

discussed in their work; however, the focus is more general, addressing authentication, confidentiality and access control, while we break down security and privacy requirements in more detail, using a framework of functionalities for all the baselines.

Main challenges and security threats in smart home networks have been analyzed by Lee et al. [17], and the fundamental requirements in order to provide secure and confidential operations in smart homes are explained from the results of their analysis. Although these requirements have been listed, they still lack practical solutions or recommendations in this matter. In [28], Suo et al. have deeply analyzed security architecture and features, and divided IoT systems into four key levels of architecture. According to this analysis, the security requirements for each level have been summarized. Furthermore, the research status of key technologies including encryption mechanism, communication security, protection of sensor data, and cryptography algorithms, have been discussed.

Roman et al. [23] have discussed threats faced by IoT, as well as security and privacy foundations based on objectives in a scenario involving a smart meter. However, they did not give any details about practical baselines and guidelines showing how to achieve these foundations.

Babar et al. [5] have presented a threat taxonomy and high-level security requirements for IoT, which like most of the other works only highlight these requirements without any practical recommendations for each category. And at the end, they introduced a security model

based on high-level requirements of security, privacy and trust. Related security requirements of IoT systems are discussed by Alqassem and Svetinovic [3], proposing a taxonomy of quality attributes, and some of the existing security mechanisms and policies in this matter have been reviewed, to reduce the identified security attacks and mitigate future vulnerabilities in these systems. They also have applied this taxonomy in a smart grid AMI as an IoT scenario. In contrast, our framework considers both security and privacy requirements and decomposes the related mechanisms, policies, and requirements with more details. In summary, our work provides a comprehensive view and a framework that covers all of the IoT security and privacy baselines, guidelines, and recommendations for every requirement.

5.4 Framework Explanation

In Figure. 5.1, we present an overview of the security and privacy functionality framework, including the top-level security and privacy concepts. The functionalities are separated into two major parts, the *life cycle* aspects of a system and the *management* aspects of security and privacy. The life cycle relates to the different phases in the life of a system, while the management of security and privacy is the ability to put supporting functionality elements in the system. We believe that awareness about where we are in the life cycle is essential, and makes it easier to apply the right functionalities and how to do the appropriate security and privacy management. We use blue color to distinguish subtopics related to the *life cycle* of a system from those associated with the *management* of security and privacy, colored in green. The coloring makes the division clear, and gives a better structure of the framework, separating the two primary concerns.

In the following, we describe each part of the framework and the related subtopics. For brevity, we do not expand the whole framework and just mention some of the aspects. For more details and complete expansion of the framework refer to the long version [12].

Security mechanisms - Different security mechanisms are illustrated in Figure. 5.2. Security mechanisms are processes designed to detect, prevent or recover from a security attack in IoT devices, including:

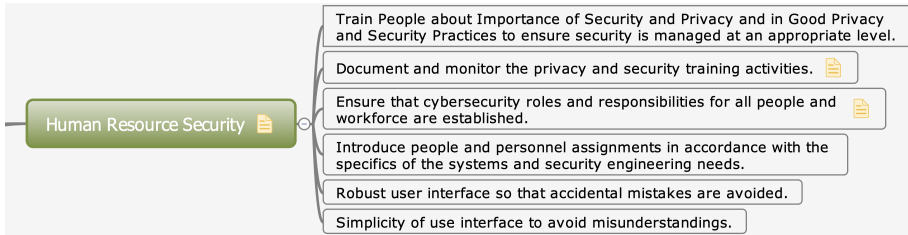


Figure 5.3: Human resource security

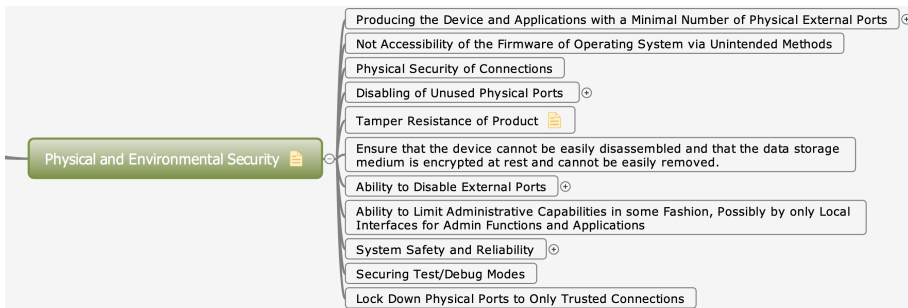


Figure 5.4: Physical and environmental security

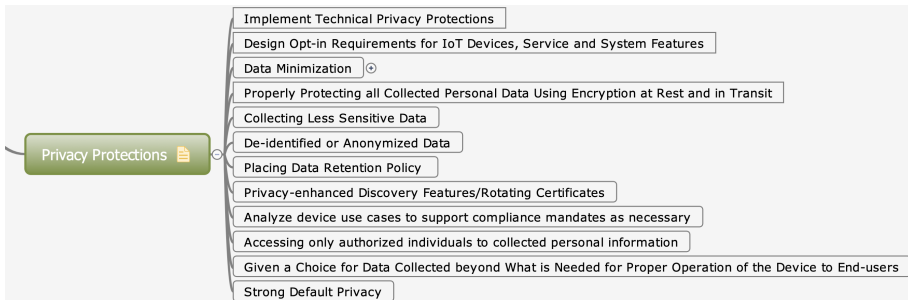


Figure 5.5: Privacy protection

- *Secure, protected and trusted communications and connectivity:* This includes information flow protection, standardising security protocols (i.e., Transport Layer Security (TLS) for encryption) guaranteeing data authenticity, signing data, disabling specific ports and/or network connections for selective connectivity, etc.
- *Hardware-based security controls:* product developers should evaluate and implement hardware protection mechanisms, including the use of Memory Protection Units (MPUs), considering a Trusted

Platform Module (TPM) into IoT Devices, securing physical interfaces, tamper protections, etc.

- *Protecting interfaces/application programming interface (APIs):* Interface security is one of the critical tasks when it comes to developing IoT devices. IoT products interact with so many cloud services, custom-developed smartphone apps and also peer IoT products. If APIs do not protect adequately, service providers might be exposed. APIs must protect adequately against misuse, by techniques like rate-limiting to protect against compromised IoT devices that attempts to flood the service, error handling, embedding time-stamps or counters into messaging to protect against replay attacks, certificate pinning to protect against sensitive data transmission into GET requests, etc.
- *Access control:* only authorized users should have access, applications and services and unauthorized accesses should be prevented, user accountability should be enabled to safeguard their authentication information.

Human resource security - People and contractors should understand the cybersecurity responsibilities suitable for their roles, and be trained about the importance of security and privacy. Further, to avoid misunderstanding, the user interfaces should be simple yet robust enough to avoid accidental mistakes. See Figure. 5.3.

Physical and environmental security - The objectives of this section include prevention of unauthorized physical access, damage, and interference with IoT's information and premises, as well as prevention of loss, damage, theft or compromise of assets and interruption to the activities and operations of IoT devices and systems. All the equipment and processing facilities should be placed in secure areas and protected from physical and environmental threats. The functional requirements of this matter are listed in Figure. 5.4.

Privacy protection - Personally identifiable information (PII) needs to be protected, according to the European General Data Protection Regulation (GDPR) regulations [30]. Privacy protection is also advisable to increase trust in the internet (see Figure. 5.5 for practical requirements).

Operations security - Information processing facilities should ensure correct and secure operation, including protection against attacks. Further, accountability auditing must be enabled for all events to ensure the




Figure 5.6: Operations security

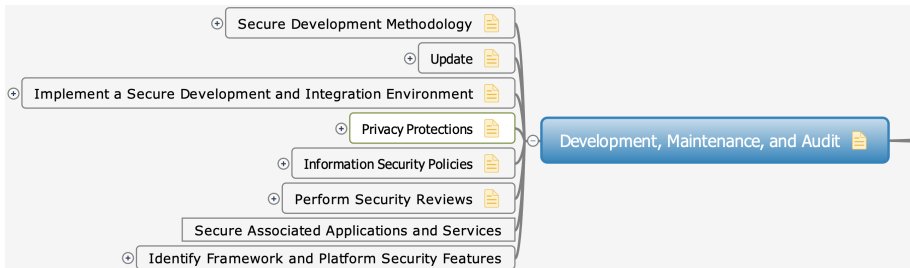


Figure 5.7: Development, maintenance, and audit



Figure 5.8: IoT security and privacy functionality framework – decommissioning

integrity of operational systems, and prevent against exploitation of technical vulnerabilities. See Figure. 5.6.

- *Logging and monitoring:* IoT products should have sufficient observations of occurrences happening on the device. For instance, connection requests, authentications (successful or failed), physical tamper, account updates, etc. It is essential to be able to monitor users' interaction with the system or a device, and in particular when they fail to login. Therefore, to detect possible security and integrity errors or potential threats, data should be captured on the entire state of the system from the endpoints, connectivity traffic and verifying the device behavior, in addition to analysing it.

- *Security configuration and management*: Includes the control of modifications to the operational functionality of the system (which covers reliability and safety behaviors) along with the security controls ensuring its protection,
- *Trust and integrity management*: Some of the practices in this matter include the following.
 1. Establishing trust in the boot environment before anything else since both the main hardware components and the operating system have been initialized by the boot process,
 2. Signing code cryptographically to prevent tampering,
 3. Controlling the installation of software on operating systems to prevent loading unauthorized software and files onto it, etc.
- *Management of security vulnerabilities and/or incidents*: To ensure a quick, effective and orderly response to information security incidents, management procedures should be established. To address identified vulnerabilities, disclosure of vulnerabilities should be coordinated. Participate in information sharing platforms, in order to report vulnerabilities and receive timely and critical information about current cyber threats and vulnerabilities from public and private partners, is recommended.

Development, maintenance, and audit - To ensure that security controls are efficient, audits and reviews for security controls should be organized periodically. Penetration tests also should be done regularly. Good practices in this area are shown in Figure. 5.7.

- *Secure development methodology*: Documentation, peer reviews, and incorporating security requirements into the product life cycle should be included, in addition to the technological checks. Additionally, essential feedback loops should be included in the engineering process to create more secure IoT products.
- *Update*: Ensuring a secured system update is probably one of the biggest challenges in an IoT life cycle. While initial systems are subject to secure testing on both the producers and the customers, a similar awareness is often missing for system updates. In the absence of sufficiently secured update, an intruder may change legitimate software and firmware, and put new malicious software and firmware into the device. Malicious software and firmware can disable security controls, apply new features or build

data exfiltration mechanisms. End-to-end protection of software and firmware is essential to the whole life cycle. And so are permissions regarding the update process, the integrity of updates, and authentication of update transactions of software and firmware.

- *Information security policies:* Regulatory, organizational and machine levels of security are covered here. The purpose of security in a system come as a security policy, and the security model represents security policies which should formally apply to the system. A security model and policy should state how to protect endpoints, communications and data, and specify what should be monitored, and analyzed, etc.
- *Perform security reviews:* Continuous feedback loops (i.e., the link connecting design, development, and test) and optimization during the life cycle of an IoT product are essential in this practice area. Identified faults/vulnerabilities have to go back to the design and threat modeling process, hardware and software baselines must be updated accordingly, and then be tested again to make sure that the patches do not identify new vulnerabilities. These vulnerabilities might be detected using IoT device security testing processes. Tests like Static Application Security Testing (SAST), Dynamic Application Security Testing (DAST), Interactive Application Security Testing (IAST) are a recommendation.
- *Secure associated applications and services:* Applications (Apps) and services connected to IoT devices must be developed securely. Configuring IoT devices, or interacting with IoT devices are usually being done using smartphone apps. These apps also create gateway functionality to transfer data from IoT devices to the cloud. So developers must use security credentials in order to provide authenticated and integrity protected communications to IoT devices.
- *Implement a secure development and integration environment:* A framework, addressing both physical and IT-security, is required to ensure a controlled environment for software and hardware development.

Decommissioning - To prevent exposing critical information to any possible attacker, the product must be disposed in a secure way at the end of life time. Therefore, secure devices should not be placed into the




Figure 5.9: Basic inputs for defining a security class [26]

supply chain again. A low-cost and high-guaranty way to decommission can be provided by an automated decommissioning procedure. We show the practical considerations of decommissioning in Figure. 5.8.

5.5 Security Classification

Concepts of security classes have been suggested and defined in [4, 26]. In this section, we briefly touch on these concepts to give an overview. There are six classes of security, from A to F, with A representing the best security and F representing the least security. Further, we based these security classes on the *exposure* and the *impact* factors of the possible attacks on the system (see Figure. 5.9 and Table 6.2). A lower exposure level means a lower attack surface. Therefore, attacks that have low exposure are relatively safe and vice-versa. The high impact of attacks on a system affects the security class of the system, and necessary precaution should be considered to protect the system. Consequently the security class of the system will increase. A system with low exposure and low impact is relatively safe.

Impact is a consequence of the possible attacks on a system. When a system is compromised, it can have an impact on several sectors beyond the system itself, including business, government, or society. We divided this *impact* into five levels, namely, Insignificant, Minor, Moderate, Major, and Catastrophic. Defining each of these levels depends on the system under evaluation, the type of impacts it can have (e.g. financial, social), or the application area. Therefore, we base the security classification level assigned to the system on the security/risk analyst responsible for that particular application or system discretion.

Exposure (see Table 6.1) is a consequence of connectivity and the

protection level of a system. According to the connectivity of the system, an appropriate set of security and privacy functionalities is identified to protect the system, while the strength of the identified security and privacy functionalities determines the protection levels (i.e., for authentication we can use passwords or PINs or two/multi-factor authentication). The definition of the protection levels (P1 to P5) is according to the ISA99 standard. However, the protection level evaluation depends on the expert and the particular scenario considered. The connectivity is divided into five levels, C1 to C5, according to ANSSI [4]:

- (C1): a closed and isolated Information & Communication System (ICS)
- (C2): an ICS connected to a corporate Management Information System (MIS) for which operations from outside the network are not allowed
- (C3): an ICS connected to wireless technology.
- (C4): an ICS with *private* infrastructure permitting operations from outside (VPN, APN, etc.)
- (C5): a distributed ICS with *public* infrastructure.

Increasing the protection level or reducing the connectivity level can reduce the exposure (see Table 6.1). From Tables 6.2 and 6.1 we observe that by keeping the protection level in the highest level (P1 is the lowest

Table 5.1: Security classes (from [26])

| | | | | | | |
|--------|---------------|---------|---------|---------|---------|---------|
| Impact | Catastrophic | Class A | Class D | Class E | Class F | Class F |
| | Major | Class A | Class B | Class D | Class E | Class F |
| | Moderate | Class A | Class B | Class C | Class E | Class E |
| | Minor | Class A | Class B | Class B | Class C | Class D |
| | Insignificant | Class A | Class A | Class A | Class B | Class C |
| | E1 | E2 | E3 | E4 | E5 | |
| | Exposure | | | | | |

Table 5.2: Exposure (from [26])

| | | | | | | |
|------------|--------------|----|----|----|----|----|
| Protection | P1 | E4 | E4 | E5 | E5 | E5 |
| | P2 | E3 | E3 | E4 | E4 | E4 |
| | P3 | E2 | E2 | E3 | E3 | E3 |
| | P4 | E1 | E1 | E2 | E2 | E2 |
| | P5 | E1 | E1 | E1 | E1 | E1 |
| | C1 | C2 | C3 | C4 | C5 | |
| | Connectivity | | | | | |

and P5 the highest protection level), exposure will be in the lowest level (E1), therefore resulting in the highest security class. And the way we can provide the highest protection level is by applying an effective and appropriate set of security and privacy functionality criteria in a particular device, for instance, use of multi-factor authentication instead of just passwords or PINs for authentication. Adequate and proper sets of security and privacy functionality criteria used in the case study have been taken from our framework.

5.6 Pacemaker Case Study

We have built a methodology for looking at the functionalities from both security and privacy points of view. In order to understand how we can use the functionality framework to improve security and privacy of IoT systems in practice, we here use a case study of health products involving a pacemaker and related control units such as a mobile phone and a heart rate sensor. In the domain of health services, the highest security class is recommended for devices like pacemakers that directly control life functions of a patient.

A pacemaker is a medical device that is implanted under the skin to help with abnormal behaviors of heartbeats [7, 15]. It consists of a battery, a computerized generator, and wires with sensors at their tips [14]. The generator is powered by the battery, and both are surrounded by a thin metal box. The generator is connected to the heart by the wires. The pacemaker helps to monitor and control the heartbeat. The sensors detect the heart's electrical activity and send data through the wires to the computer in the generator. If the heart rhythm is abnormal, the generator will be directed by the computer to send electrical pulses to the heart, and the pulses travel through the wires to reach the heart. Embedded microprocessors in modern pacemakers have enabled them to do additional tasks like monitoring heart activity and providing a record for the patients and their healthcare providers, as well as collecting data on heart functions to help doctors to identify and diagnose patient conditions, and send required shock signals when needed. Doctors can monitor the patient's heart activities and control the pacemaker using a mobile phone or a computer device, or send required shock signals in case of observation of abnormal behavior.

Beside threatening patients' lives, malicious attackers can get access to patients' medical records through a pacemaker [9, 29], or track a patient's location. In addition, malicious software can be run on pacemakers and cause security and privacy breaks. Therefore, any security or functional weakness can result in a security failure. Like any device that uses remote technology, pacemakers are also vulnerable against cyber attacks, and hackers can break into the pacemaker itself, the back-end systems or the communication between the pacemaker and its surrounding. By breaking into a pacemaker, an attacker can send strong shock signals, disturb the pacemaker setting or heart functions, or disturb them from working properly. One simple example of hacking into a pacemaker is to change the setting from battery-saving "sleep" mode to "standby" mode, and this can quickly drain the battery, which is normally supposed to last for years. Furthermore, an attacker can steal private and personal information of a patient from the device that for instance can later track the patient's location. Hence, security in this kind of device is crucial and should have the highest (best) security class, meaning security class A.

We will below discuss the security and privacy challenges for each of the three devices. In each case we discuss connectivity, protection level and relevant functionality criteria. We use the general criteria given in the functionality framework, select and discuss the parts relevant for each device.

Pacemaker security controller. We consider Connectivity 2 (explained in Section 6.3) for the pacemaker security controller, since it is only connected to the pacemaker. We define below the protection levels which are relevant for the pacemaker security controller:

- Protection level 1 (P1): includes secure authentication, securing software/firmware, secure communication, and human interface security. For authentication we consider: requiring passwords, option to change the default username and password. For communication we consider: data authenticity to enable reliable exchanges from data emission to data reception. For securing software/firmware we consider: update capability for *some* of the system devices and applications, transmitting the files using encryption.
- Protection level 2 (P2): includes P1 and in addition, for authenti-

cation: requiring strong passwords, securing password recovery mechanisms, making sure that default passwords and even default usernames are changed during the initial setup, and that weak, null or blank passwords are not allowed. For communication: verifying any interconnections, discover, identify and verify/authenticate the devices connected to the network before trust can be established, and preserve their integrity for reliable solutions and services, prevent unauthorised connections to it or other devices the product is connected to, at all protocol levels, providing communication security using state-of-the-art mechanisms, standardising security protocols, such as TLS for encryption. For securing software/-firmware: encrypting update files for *some* of the applications.

- Protection level 3 (P3): includes P2 and in addition, for authentication: having options to force password expiration after a specific period, and to change the default username and password, making sure that the password recovery or reset mechanism are robust and do not supply an attacker with information indicating a valid account. The same should apply to key update and recovery mechanisms. For communication: data authenticity to enable reliable exchanges from data emission to data reception.
- Protection level 4 (P4): includes P3 and in addition, for authentication: implementing two-Factor Authentication (2FA), making sure that default passwords and even default usernames are changed during the initial setup. For communication: signing the data whenever and wherever it is captured and stored, making intentional connections, disabling specific ports and/or network connections for selective connectivity. For securing software/firmware: capability of quick updates when vulnerabilities are discovered for *some* of the system devices and applications, and offering an automatic firmware update.
- Protection level 5 (P5): includes P4. In addition, for authentication: using Multi-Factor Authentication (MFA) (considering biometrics for authentication), considering Certificate-Less Authenticated Encryption (CLAE), and User Managed Access (UMA). For communication: rate limiting – controlling the traffic sent or received by a network to reduce the risk of automated attacks. For securing

software/firmware: update capability for *all* system devices and applications, capability of quick updates when vulnerabilities are discovered for *all* system devices and applications, encrypting update files for *all* applications, signing update files and validating by the device before installing, securing update servers, having ability to implement scheduled updates, having backward compatibility of firmware updates.

According to Tables 6.2 and 6.1, we might have protection levels 2, 3, 4, or 5 to obtain security class A; however we choose to use protection level 3 since it is a realistic level in this case. A higher level would be costly, and P2 would give poor protection. We therefore consider how to obtain that protection level, and select the following set of relevant functionality criteria from the functionality framework, adapted to the challenges of pacemakers. In this selection we have used the guidelines for securing pacemakers from [13]. This gives a certain guarantee that we cover all relevant aspects.

- *Secure authentication/authorization*: Access to the pacemaker security controller and mobile phone connected to the doctor, should be limited using the authentication of users (for example user ID and password, Personal Identification Numbers (PINs), biometric authentications). Authorization refers to checking necessary permissions of an identified individual to do an action. Authentication and authorization are completely related to each other. Authorization checks should immediately be followed by authentication of a request. In order to have secure authorization, the roles and permissions of the authenticated user should only be verified through information in backend systems, not through roles or permission information coming from the device. Any incoming identifiers with a request alongside should be verified by backend code independently. Failure in a secure authentication/authorization would give access to unauthorized people and could lead to reputational damages, fraud, unauthorized access to information, information theft, and modification of data.
- *Securing software/firmware*: Before any software or firmware update, user authentication or other suitable controls should be required. Software/firmware updates should be restricted to authorised code. Manufacturers may consider code signature

verification as an authentication method.

- *Secure communication*: Data transmission between the pacemaker security controller and the mobile phone connected to the doctor must be secure enough so that a third party cannot listen to their communication. The communication should not be vulnerable to eavesdropping or interception. Failure in having a secure communication can cause identity theft, fraud, data modification, privacy information leakage. One should consider strong handshaking, correct SSL versions, no clear-text communication of sensitive assets, etc.
- *Human interface security*: Patients should be trained about the importance of security and privacy and how to use the pacemaker properly to ensure security is managed at an appropriate level. For instance, if we consider all the security protections in the highest level in the security controller, but the patient does not know how to deal with error notifications (restart, turn off or low battery errors) in the security controller, all the security considerations in the controller will be useless.
- *Data privacy*: Measures to avoid risk of breaches in connection with long term storage of private information, handling of encryption of private information, and GDPR compliance including consent, purpose, and access rights.

The final discussion of the security class of the pacemaker system depends on the scenario chosen and the other components involved. We next consider the mobile phone.

Mobile phone. In our case study, a mobile phone is used in the communication between the pacemaker and the healthcare provider/doctor. Security in this mobile phone is then important in order to send correct data to/from the pacemaker controller. Hacking or tampering into this mobile phone can result in sending wrong data from the pacemaker to the doctor, something that could result in wrong decisions from the doctor, or possibly sending inappropriate shocks to the patient's heart.

Therefore, it is important to secure the mobile phone properly. However, the security class of the mobile phone is only considered class B, since mobile phones are inherently not of the highest security class

and have a number of possibilities for attacks due to high exposure. For instance, because of all the applications and browsers on the device, as well as software developed by third parties. Therefore, both the impact of attack as well as the connectivity are in a higher level for the mobile phone compared to the pacemaker. Hence the security class of the mobile phone is lower than the security class of the pacemaker security controller.

We consider C4 in the mobile phone connected to the doctor. According to Tables 6.2 and 6.1, we might have protection levels 2, 3, or 4 to obtain security class B, however since mobile phones naturally have moderate impact of attacks as discussed above, we should use protection level 4 in order to reach security class B. So based on that the following set of relevant security and privacy functionality criteria are selected to ensure that the mobile phone has a sufficiently high protection level, following the functionality framework and the OWASP guidelines for securing mobile phones [20]:

- *Secure authentication/authorization and secure communication:* are the same as what we discussed for the pacemaker security controller.
- *Sufficient cryptography techniques:* Appropriate cryptography techniques should be considered for instance using AES instead of DES.
- *Code tampering:* When an application is on the device, the code and data resources are also available there. An attacker modify the code through either malicious apps in third party app stores or trick the user via phishing attacks and install the app on the device. Code tampering could result in revenue loss due to piracy, reputational damage, unauthorized new features, identity theft or fraud.
- *Secure data storage:* Failure in having secure data storage can result in data loss, extraction of sensitive data using malware, modified apps or forensic tools, identity theft, fraud, reputation damage, material loss or external policy violations.
- *Proper platform usage:* Misuse of a platform feature or failure to use platform security controls fall under this category. Android intents, platform permissions, misuse of TouchID, Keychains, or other security controls which are part of the operating system could




Figure 5.10: Scenario 1: pacemaker along with security controller




Figure 5.11: Scenario 2: pacemaker with separate security controller

be included. To prevent the attacks in this category, secure coding and configuration practices must be applied on the server side of the application.

- *Data privacy*: are as discussed above for the pacemaker security controller.

We next consider the protection levels of ISA99, and select the relevant criteria from the functionality framework, and adjust them for the case of the mobile phone connected to the doctor. These are defined below:

- Protection level 1 (P1): includes secure authentication, secure authorization, securing software/firmware, sufficient cryptography techniques, code tampering, secure data storage, secure communication, and proper platform usage. Secure authentication, securing software/firmware, and secure communication are the same as the considerations for protection level 1 for the pacemaker security controller.
- Protection level 2 (P2): includes P1. Secure authentication, securing software/firmware, and secure communication are the same as the considerations for protection level 2 for the pacemaker security controller. In addition, for cryptography techniques we should consider: ensuring proper selection of standard encryption algorithms and keys.
- Protection level 3 (P3): includes P2. Secure authentication,

securing software/firmware, and secure communication are the same as the considerations for protection level 3 for the pacemaker security controller. In addition, for cryptography techniques we should consider: strong encryption algorithms, strong keys.

- Protection level 4 (P4): includes P3. Secure authentication, securing software/firmware, and secure communication are the same as the considerations for protection level 4 for the pacemaker security controller. In addition, for cryptography techniques we should consider: verifying the robustness of the implementation, establishing secure and scalable key management. And cryptographic keys must be securely managed.
- Protection level 5 (P5): includes P4. Secure authentication, securing software/firmware, and secure communication are the same as the considerations for protection level 4 for the pacemaker security controller. In addition, for cryptography techniques we should consider: disabling insecure protocols.

Finally, the security class of the heart rate sensor in our case study is considered below.

Heart rate sensor. The heart rate sensor does not have any external connectivity; it has only connections to the heart and computerized generator inside the pacemaker in order to transfer the captured heart rate from the heart to the pacemaker generator. So the connectivity in the sensor is C1. Furthermore it has only a chip set from the company provider. Hence, the sensor has the least chance of attack. However, there is still a possible vulnerability if the wired connection is not sufficiently shielded and allows eavesdropping by inductive sensors or senders. But this requires very short physical distance. Therefore we may assume only minor impact, and do not need a high protection level for this device, and may use P1. Exposure is then E4 (see Table 6.4). This gives security class C.

Pacemaker scenarios. The challenge in the case study is how to implement the security controls like decryption, authentication, etc. For instance, whether to do data decryption in the mobile phone, and then send the decrypted data to the pacemaker. If the mobile phone is compromised, wrong data and signals could then go to the doctor and pacemaker, since mobile phones can get compromised, as we discussed before. Whereas, if we consider security controls in the pacemaker itself, we only have software provided by one company. Thus it would be much

Table 5.3: Security and privacy challenge comparison of scenario 1 and 2

| Security and Privacy Challenges | | Scenario 1 | Scenario 2 |
|--|---|------------|------------|
| In Sensor | Draining battery by changing battery saving controls | | |
| | Interrupting into heartbeat capturing | | |
| | Data and private information breaches | | |
| In Pacemaker | Sending/Receiving wrong data to/from mobile phone connected to the doctor | | |
| | Changing pacemaker shock settings | | |
| | Draining battery by changing battery saving controls | | |
| | Higher battery usage | | |
| | Data and private information breaches | | |
| | Transparency of GDPR compliance | | |
| | Risk of long-term storage of private information | | |
| In Mobile Phone Connected to the Doctor | Sending/Receiving wrong data to/from pacemaker | | |
| | Sending wrong unnecessary/necessary shocks | | |
| | Changing pacemaker shock settings | | |
| | Draining pacemaker battery by changing battery saving controls | | |
| | Data and private information breaches | | |
| | Transparency of GDPR compliance | | |
| | Risk of long-term storage of private information | | |

Severity of Challenge: Insignificant: Minor: Moderate: Major: Catastrophic:

more secure to do the security controls directly in the pacemaker, and use the mobile phone just as a gateway to transfer the information. However, this will require additional computational power and battery capacity. Another issue is whether we can do all the security controls inside the computerized generator of the pacemaker - or consider a security controller as a separate unit out of the body with a close and secure connection to the pacemaker.

We therefore define two scenarios: In scenario 1 (see Figure. 6.1), the security controls for the pacemaker are done inside the computerized generator of the pacemaker, and in scenario 2 (see Figure. 6.2), a separate security controller unit makes the security controls of the pacemaker, such that this unit is outside the body with a close and secure connection to the pacemaker. In scenario 1, we would need computational power inside the body, needing more storage, stronger CPU, much more battery capacity and so on. This is not desirable since it might increase the potential necessary surgeries in order to change the battery, maintain, or update the pacemaker. Moreover, we might not be able to consider some of the security and privacy functionalities in order to avoid increasing CPU usage, which results in even more battery usage.

Therefore, we might want to have a pacemaker with a simple sensor

inside the body and a security controller out of the body, say in the pocket or at home very close to the pacemaker. Here, it is essential that the security controller be close to the pacemaker, since the pacemaker must have very weak signals, limited interactions and computations, to avoid using too much battery power and resulting battery changes. So, all the security and battery-intensive controls would be done in the external security controller, which can have a stronger and easily rechargeable battery. The controller can maintain the device in case of any problem, update or troubleshoot its computer system, or even increase the level of protection by adding more security and privacy functionalities or appropriate software at any time because of easy access.

According to all the security considerations in scenario 2, we can then have a better security class in the pacemaker security controller: In Table 6.3 we summarize all the security and privacy challenges in the sensor, pacemaker, and mobile phone connected to the doctor, and compare the severity of these security and privacy challenges in scenarios 1 and 2. This table determines the impact. The challenges considered in the table are found by following the functionality framework, for each device, in a top-down manner and in each case determining the problems that may occur. The further we break down the problems according to the functionality framework, the easier it is to find the specific challenges for the given device.

In the following, we see that the severity of the security and privacy challenges in the sensor and pacemaker has been reduced from very high in the worst case to low. Hence, the impact of attacks is reduced from catastrophic to minor.

Discussion. The pacemaker in scenario 1 is a complex device because all security controls are done inside the computerized generator of the pacemaker. The security controller is very close to the sensor, the signals received by the sensor from the pacemaker are very frequent. Interference is possible and that can have negative effect on correct heartbeat capturing. However, by transferring the security controller outside of the body this effect would be low (still there are some frequencies from devices close by such as mobile phones that can affect the sensor) resulting in more precise heartbeat capturing. As mentioned earlier in the heart rate sensor, there are low possibilities of data and private information breaches in both scenarios.

In the pacemaker in scenario 1, because of the complexity of the sys-

tem and having all the security controls inside the computerized generator of the pacemaker, we need higher CPU and memory consumption and then battery usage would be very high, while in scenario 2 this problem would decrease to very low. In scenario 1, because of difficulties in accessing the pacemaker and its security controller on time (in case of maintenance, update, troubleshooting or installing new software/equipment, also not being able to apply all the necessary criteria with high protection level in order to avoid high battery usage), the level of the security and privacy functionality criteria as well as the level of the protection is low.

Therefore the vulnerability of the pacemaker and sensor against attacks compromising the device (that can cause changing battery saving controls so draining the battery more quickly, changing pacemaker shock settings, tampering of transferred information from/to the pacemaker, data and private information breaches, transparency of GDPR compliance, risk of long-term storage of private information) is high; however, this problems decrease to very low when we change to scenario 2, due to easier access to the security controller out of the body.

By compromising the mobile phone connected to the doctor, the problems listed in the last part of Table 6.3 may occur. In both scenarios, we have considered high level of protection for the mobile phone by

Table 5.4: Security class of the sensor

| | | | | | | |
|-------------------|---------------------|----|---------------|-----------------|---------|---------------|
| Protection | P1 | E4 | Impact | Catastrophic | Class F | In Scenario 1 |
| | P2 | E3 | | Major | Class E | |
| | P3 | E2 | | Moderate | Class E | |
| | P4 | E1 | | Minor | Class C | In Scenario 2 |
| | P5 | E1 | | Insignificant | Class B | |
| | C1 | | | E4 | | |
| | Connectivity | | | Exposure | | |

Table 5.5: Security class of the pacemaker security controller

| | | | | | | |
|-------------------|---------------------|----|---------------|-----------------|---------|---------------|
| Protection | P1 | E4 | Impact | Catastrophic | Class D | In Scenario 1 |
| | P2 | E3 | | Major | Class B | |
| | P3 | E2 | | Moderate | Class B | |
| | P4 | E1 | | Minor | Class B | In Scenario 2 |
| | P5 | E1 | | Insignificant | Class A | |
| | C2 | | | E2 | | |
| | Connectivity | | | Exposure | | |

Table 5.6: Security class of the mobile phone

| | | | | | |
|-------------------|---------------------|----|---------------|-----------------|---------|
| Protection | P1 | E5 | Impact | Catastrophic | Class D |
| | P2 | E4 | | Major | Class B |
| | P3 | E3 | | Moderate | Class B |
| | P4 | E2 | | Minor | Class B |
| | P5 | E1 | | Insignificant | Class A |
| | C4 | | | E2 | |
| | Connectivity | | | Exposure | |

In Scenarios 1 & 2

using appropriate set of security and privacy functionality criteria with high protection level, therefore we have reduced the vulnerability of the device to low, but as discussed earlier, these devices still have a number of vulnerabilities for attacks, and therefore the vulnerability is not in a very low level. Furthermore, softening the problems in scenario 1 results in obtaining higher protection level and security class. For instance, we can easily recharge the battery of the security controller, maintain the device in case of any problem to avoid shutting down all the security considerations, and increase the level of protection by updating or troubleshooting its computer system..

In the sensor, as explained above, we have connectivity C1, and protection level P1, therefore exposure is E4 (see Table 6.4), and because of very high problem severity, the impact of attacks is Catastrophic, resulting in security class F in this device. By changing from scenario 1 to 2, the severity of the security and privacy challenges has reduced from very high to low, therefore the impact of attacks is reduced from Catastrophic to Minor, and consequently, the security class has improved from class F to class C.

In the pacemaker, we have connectivity C2, and protection level P5, however in scenario 1, the severity of security and privacy challenges have affect on the protection level, which falls to level 3, therefore exposure is E2 (see Table 6.5), and because of very high security and privacy challenge severity, the impact of attacks is Catastrophic, then we have security class D in this device. By changing from scenario 1 to 2, the severity of the security and privacy challenges has reduced from very high to very low, therefore the impact of attacks has reduced from Catastrophic to Insignificant, and consequently, the security class has improved from class D to A.

And, in the mobile phone connected to the doctor, we have connectivity C4, and protection level P4, therefore the exposure is E2 (see

Table 6.6), and the security is class B in both scenarios. By considering security and privacy challenges, as well as the effect of these challenges on the protection level of each device and the whole system, and also their affect on the impact of attacks, we have changed scenario 1 to 2, and were able to improve the security class in the sensor and pacemaker security controller from class F to C, and class D to A, respectively. Hence the security of the overall system has improved significantly.

In this case study, we started out with security and privacy design requirements to each device: class A for the pacemaker security controller, class B for the mobile phone, and class C for the sensor. We have seen that this is not realistically possible to achieve for the design of scenario 1, while it was possible to satisfy these requirements for scenario 2. The framework was most useful in these design evaluations. Indeed, the case study shows that by following the guidelines given by our framework, one can achieve security easily and decrease the impact of a possible attack.

5.7 Conclusion

The expansion of IoT in the last decade has resulted in several security and privacy issues and attacks against things and people. Unfortunately, the security and privacy functionalities to combat these attacks are not well-recognized in the domain of IoT. This paper summarizes and categorizes IoT security and privacy functionalities, and as the main contribution, the paper presents a new taxonomy framework that organizes the related standards in this area. The proposed framework is oriented towards practical application. We have demonstrated the application of this framework in combination with the security classification method using a case study about pacemakers. Our case study is quite generic and reveals general issues that can be found in other case studies.

The security class of a system is based on two factors: exposure and impact of possible attacks. The exposure is a consequence of the protection level of the system and its connectivity. A lower exposure level means a lower attack surface. Therefore, by reducing the exposure level of a system we can have a more secure system. A higher protection level in a system or lower connectivity can result in lower exposure. At the same time, lower impact of attacks on a system raises the security class

of the system. By applying the appropriate set of security and privacy functionality criteria from our framework, the protection level of a system can increase, while exposure can decrease, as demonstrated by the case study and discussed at the end of Section 6.4.

The main objective of this paper is to give security developers, designers, and end-users an opportunity to understand and explore what the IoT security and privacy functionalities are, and how these functionalities can help to improve the security and privacy of IoT systems. Our approach combines detailed information about security and privacy functionalities with a security classification method. The approach is systematic and structured; it is easy to use and is oriented towards practical engineering. The framework is based on the available recommendations and standards for IoT systems. This should imply that all aspects are covered, but there is no guarantee for that. This can be seen as a limitation of the work. Secondly, the application of the methodology is ultimately depending on the judgements of software engineers or security experts, and is therefore not 100% precise. If their judgement is wrong, for instance if they choose the wrong connectivity or protection level, it will in general give a wrong estimate.

Future work will consider case studies with several kinds of IoT devices and sensors involved. This will be a valuable step toward validating our framework, and possibly allowing the framework to be complemented with even more elements.

Bibliography

- [1] 27000, I. *Information technology — Security techniques — Information security management systems — Overview and vocabulary (fourth edition)*. ISO/IEC 2016. 2016.
- [2] 27001, I. *INTERNATIONAL STANDARD ISO/IEC 27001 Information technology—Security techniques —Information security management systems —Requirements*. ISO/IEC 2013. 2013.
- [3] Alqassem, I. and Svetinovic, D. “A taxonomy of security and privacy requirements for the Internet of Things (IoT)”. In: *2014 IEEE Intern. Conf. on Industrial Engineering and Engineering Management (IEEM)*. IEEE. 2014, pp. 1244–1248.
- [4] ANSSI. *Classification Method and Key Measures*. 2014.
- [5] Babar, S. et al. “Proposed security model and threat taxonomy for the Internet of Things (IoT)”. In: *Intern. Conf. on Network Security and Applications*. Springer. 2010, pp. 420–429.
- [6] BITAG. *Internet of Things (IoT) Security and Privacy Recommendations*. [https://www.bitag.org/documents/BITAG_Report_-_Internet_of_Things_\(IoT\)_Security_and_Privacy_Recommendations.pdf](https://www.bitag.org/documents/BITAG_Report_-_Internet_of_Things_(IoT)_Security_and_Privacy_Recommendations.pdf). 2016.
- [7] Boulos, P. et al. “Pacemakers: a survey on development history, cyber-security threats and countermeasures”. In: *Int. J. Innov. Stud. Sci. Eng. Technol* vol. 2, no. 8 (2016).
- [8] *Future-proofing the Connected World: 13 Steps to Developing Secure IoT Products*. <https://downloads.cloudsecurityalliance.org/assets/research/internet-of-things/future-proofing-the-connected-world.pdf>. 2016.
- [9] Dadourian, M. *Heart alert: Pacemakers can be hacked*. 2018.
- [10] ENISA. *Baseline Security Recommendations for IoT in the context of Critical Information Infrastructures*. European Union Agency for Network and Inf. Sec. 2017.

- [11] ENISA. *Guidelines for SMEs on the security of personal data processing*. <https://www.enisa.europa.eu/publications/guidelines-for-smes-on-the-security-of-personal-data-processing>. 2017.
- [12] Fazeldehkordi, E., Owe, O., and Noll, J. *Security and Privacy Functionalities in the Internet of Things (long version)*. Tech. rep. University of Oslo, 2019.
- [13] FDA. *Content of Premarket Submissions for Management of Cybersecurity in Medical Devices Guidance for Industry and Food and Drug Administration Staff*. U.S. Food and Drug Administration (FDA). 2014.
- [14] Health & Human Services, U. D. of. *Pacemaker*. 2019.
- [15] Healthline. *Heart Pacemaker*. 2019.
- [16] Islam, S. R. et al. “The Internet of Things for health care: a comprehensive survey”. In: *IEEE Access* vol. 3 (2015), pp. 678–708.
- [17] Lee, C. et al. “Securing smart home: Technologies, security challenges, and security requirements”. In: *2014 IEEE Conference on Communications and Network Security*. IEEE. 2014, pp. 67–72.
- [18] MacGillivray, C., Turner, V., and Lund, D. *Worldwide Internet of Things (IoT) 2013–2020 Forecast: Billions of Things, Trillions of Dollars*. International Data Corporation. <http://www.idc.com/getdoc.jsp?containerId=243661f>. 2014.
- [19] OWASP. *IoT Security Guidance*.
- [20] OWASP. *OWASP Mobile Security Project*.
- [21] Pang, Z. “Technologies and Architectures of the Internet-of-Things (IoT) for Health and Well-being”. PhD thesis. KTH Royal Institute of Technology, 2013.
- [22] Pishva, D. “Internet of Things: Security and privacy issues and possible solution”. In: *Advanced Communication Technology (ICACT), 2017 19th Intern. Conf. on*. IEEE. 2017, pp. 797–808.
- [23] Roman, R., Najera, P., and Lopez, J. “Securing the Internet of Things”. In: *Computer* vol. 44, no. 9 (2011), pp. 51–58.

- [24] Ross, R. et al. “Protecting controlled unclassified information in nonfederal information systems and organizations”. In: *NIST Special Publication* vol. 800 (2015), p. 171.
- [25] Schrecker, S. and al., et. *Industrial Internet of Things Volume G4: Security Framework*. Industrial Internet Consortium. https://www.iiconsortium.org/pdf/IIC_PUB_G4_V1.00_PB.pdf. 2016.
- [26] Shrestha, M., Johansen, C., and Noll, J. “Security Classification for Smart Grid Infra structures (long version)”. In: (2017).
- [27] Sicari, S. et al. “Security, privacy and trust in Internet of Things: The road ahead”. In: *Computer networks* vol. 76 (2015), pp. 146–164.
- [28] Suo, H. et al. “Security in the internet of things: a review”. In: *2012 international conference on computer science and electronics engineering*. Vol. 3. IEEE. 2012, pp. 648–651.
- [29] Thehackernews.com. *Over 8,600 Vulnerabilities Found in Pacemakers*. <https://thehackernews.com/2017/06/pacemaker-vulnerability.html>.
- [30] Union, E. *REGULATION (EU) 2016/679 OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL*. <https://eur-lex.europa.eu/eli/reg/2016/679/oj>.

Paper 2: Security and Privacy in IoT Systems: A Case Study of Healthcare Products

Authors: Elahe Fazeldehkordi, Olaf Owe, Josef Noll

Publication: Presented at 13th International Symposium on Medical Information and Communication Technology (ISMICT), 8-10 May, 2019 (pp. 1-8). Published in IEEE. DOI: <https://doi.org/10.1109/ISMICT.2019.8743971>

Abstract

Internet of Things (IoT) is facilitated by heterogeneous technologies, which contribute to the providing of innovative services in a large number of application domains. The satisfaction of security and privacy requirements in this scenario are becoming a main challenge for IoT systems and their developers. Nevertheless, most works on IoT security and privacy requirements look at these requirements from a high level view. Hence, important aspects of security and privacy functionalities will be disregarded, causing wrong design decisions. To combat this problem, in our previous work, we summarized the most current documents related to security and privacy functionalities in the setting of IoT and provided a new taxonomy framework that organizes all aspects of security and privacy baselines, guidelines and recommendations. To give an understanding of how the framework can help to improve security and privacy of IoT products, and help to facilitate developing and designing secure and privacy-aware IoT systems, in this paper we delve deeper and demonstrate the usefulness of the framework by a case study of healthcare products, in combination with a recent security classification method.

6.1 Introduction

Internet of Things (IoT) involves information flow between different kinds of embedded computing devices interconnected through a network. The aim of IoT is to enable an advanced mode of communication among the various systems and devices, and also to facilitate the interaction between humans and the virtual world. With this aim, IoT plays a significant role in the modern society and has applications in almost all fields of the modern society including healthcare systems, automobile, industrial appliances, sports, homes, entertainments, environmental monitoring etc. IoT devices have already outnumbered the number of people at computerized workplaces, and by 2020, connected "things" based on IoT will be around 212 billion [18, 22]. Those "things" include daily used appliances like smart-phones, smart-watches, smart television, smart refrigerators and others. As a result of this expansion, and as most things are connecting to the internet for exchanging information, IoT is vulnerable to various security issues and attacks, e.g., man in the middle attack, eavesdropping attack, denial of service attack, access attack, as well as major privacy concerns for the end users. Even though IoT provides advanced abilities in the data communication area, it is vulnerable from the security and privacy points of view. Therefore appropriate steps in the initial phase of design and development of IoT systems should be taken.

In our previous work [12], we gave a comprehensive view of the state of the art with respect to security and privacy functionalities and requirements for IoT systems, and suggested a complementary methodology for analyzing the functionalities in a comprehensive framework that can help both providers and consumers of IoT devices to have a better understanding of the security and privacy aspects [12]. By functionality we mean: "The security and privacy-related features, functions, mechanisms, services, procedures, and architectures implemented within organizational information systems or the environments in which those systems operate" [24]. For this framework we investigated IoT-related security baselines and guidelines developed by ENISA [10], OWASP [19], Industrial Internet Consortium [25], Cloud Security Alliance [8], and Broadband Internet Technical Advisory Group [6], as well as IoT related security guidelines from ISO [1, 2] and NIST [24]. With respect to privacy, EU has passed

the general data protection regulation (GDPR), which regulates who can access private data, how and for what purpose, based on consent of the data subject [11]. We then categorized and integrated these guidelines and requirements in a uniform style, and embedded them in a graphical representation by means of diagrams. Having a comprehensive view and taxonomy of security and privacy requirements and functionalities in IoT is a prerequisite for architecting optimal security solutions, designing, and developing secure and privacy-aware IoT systems.

In this paper, we explore how the framework can facilitate the process of improving IoT security and privacy, in combination with the security classification method suggested in [4, 26]. Through the development of this framework, together with the security classes, extensive attention has been given to the requirements and limitations for securing IoT systems. The security classification of a system will lead to better understanding of the value of security and promote the extra cost of securing a system. An IoT system includes different kinds of devices, and protecting all of these devices at a same level is costly. It is economically impossible to employ all the security protection mechanisms for all the devices in a system. Dividing security into different classifications is necessary in order to secure IoT systems to an appropriate level.

To give an understanding of how the framework can help to improve security and privacy in practice, we give a comprehensive discussion on the details of two scenarios in a case study of healthcare products, compare security and privacy challenges of the two scenarios, and then show how to soften these challenges using security classifications and functionalities of our previous framework. One of the most attractive application areas of IoT is health care [16, 21]. Medical applications like remote health monitoring, fitness programs, chronic diseases, elderly care, compliance with treatment and medication at home and by healthcare providers, are some of the important potential applications that IoT can bring. IoT-based healthcare services can help to reduce costs, increase the quality of life, and enrich the user's experience. Therefore, in this paper we demonstrate the framework on a case study concerning healthcare products.

The remainder of this paper is structured as follows: Section 2 provides related work. Section 3 gives a background about the security classification method. Section 4 describes the pacemaker case study, and Section 5 concludes the paper.

6.2 Related Work

In a work presented by Sicari et al. [27], the most relevant available solutions regarding security, privacy, and trust in IoT have been analyzed. Proposals related to security middleware, secure solutions for mobile devices and ongoing international projects on this subject have also been discussed in their work; however, the focus is more general, addressing authentication, confidentiality and access control, while we break down security and privacy requirements in more detail, using the framework of functionalities for all the baselines.

Major challenges and security threats in smart home networks have been analyzed by Lee et al. [17], and fundamental requirements in order to provide secure and confidential operations in smart homes are explained from the results of their analysis. However, these requirements are just listed up and there is no practical solutions or recommendations in this matter. In [28], Suo et al. have deeply analyzed security architectures and features, and divided IoT systems into four key levels of architecture. According to this analysis, the security requirements for each level have been summarized. Furthermore, the research status of key technologies including encryption mechanism, communication security, protection of sensor data, and cryptography algorithms, have been discussed in this paper.

Roman et al. [23] have discussed threats faced by IoT, as well as security and privacy foundations based on objectives in a scenario involving a smart meter. However, they did not give any details about practical baselines and guidelines showing how to achieve these foundations.

Babar et al. [5] have presented a threat taxonomy and high level security requirements for IoT, which like most of the other works just number these requirements without any practical recommendations for each category. And at the end, they introduced a security model based on high-level requirements of security, privacy and trust. Related security requirements of IoT systems are discussed by Alqassem and Svetinovic [3], proposing a taxonomy of quality attributes, and some of the existing security mechanisms and policies in this matter have been reviewed, in order to reduce the identified security attacks and mitigate future vulnerabilities in these systems. They also have applied

this taxonomy in a smart grid AMI as an IoT scenario. In contrast, the framework [12] considers both security and privacy requirements and decomposes the related mechanisms, policies, and requirements with more details.

6.3 Security Classification

Concepts of security classes have been suggested and defined in [4, 26]. Here, we briefly touch on these concepts to give an overview. There are six classes of security, from A to F, with A representing the best security and F representing the least security. Security classes are based on two factors: *exposure* and *impact* of the possible attacks on the system (see Table 6.2). A lower exposure level means a lower attack surface, therefore, attacks that have low exposure are relatively safe. And vice versa, high impact of attacks on a system affects the security class of the system and necessary precaution should be considered to protect the system. Consequently the security class of the system will be raised. A system with low exposure and low impact is relatively safe.

Table 6.1: Exposure (from [26])

| | | | | | | |
|-------------------|---------------------|----|----|----|----|----|
| Protection | P1 | E4 | E4 | E5 | E5 | E5 |
| | P2 | E3 | E3 | E4 | E4 | E4 |
| | P3 | E2 | E2 | E3 | E3 | E3 |
| | P4 | E1 | E1 | E2 | E2 | E2 |
| | P5 | E1 | E1 | E1 | E1 | E1 |
| | | C1 | C2 | C3 | C4 | C5 |
| | Connectivity | | | | | |

Table 6.2: Security classes (from [26])

| | | | | | | |
|---------------|-----------------|---------|---------|---------|---------|---------|
| Impact | Catastrophic | Class A | Class D | Class E | Class F | Class F |
| | Major | Class A | Class B | Class D | Class E | Class F |
| | Moderate | Class A | Class B | Class C | Class E | Class E |
| | Minor | Class A | Class B | Class B | Class C | Class D |
| | Insignificant | Class A | Class A | Class A | Class B | Class C |
| | | E1 | E2 | E3 | E4 | E5 |
| | Exposure | | | | | |

Impact is a consequence of the possible attacks on a system. When a system is compromised, it can have impact on several sectors beyond the system itself, including business, government, or society. The *impact* is divided into 5 levels: Insignificant, Minor, Moderate, Major, and Catastrophic. Defining each of these levels depends on the system under evaluation, the type of impacts it can have (e.g. financial, social), or the application area. So, it is based on the judgment of the security/risk analysts in that special application or system.

Exposure (see Table 6.1) is a consequence of the connectivity and the protection level of a system. According to the connectivity of the system, an appropriate set of security functionalities is identified to protect the system, and the strength of the identified security functionalities determines the protection levels (for instance for authentication we can use passwords or PINs or two/multi-factor authentication). The definition of the protection levels (P1 to P5) is according to the ISA99 standard. However, the protection level evaluation is depending on the expert and the particular scenario considered.

The connectivity is divided into five levels, C1 to C5, according to ANSSI [4]:

- (C1): a closed and isolated Information & Communication System (ICS)
- (C2): an ICS connected to a corporate Management Information System (MIS) for which operations from outside the network are not allowed
- (C3): an ICS connected to wireless technology.
- (C4): an ICS with *private* infrastructure permitting operations from outside
- (C5): a distributed ICS with *public* infrastructure.

Increasing the protection level or reducing the connectivity level can reduce the exposure (see Table 6.1). As we can observe from Tables 6.1 and 6.2, by keeping the protection level in the highest level (P1 is the lowest and P5 the highest protection level), exposure will be in the lowest level (E1), therefore resulting in the highest security class. And the way we can provide the highest protection level is by applying an effective and appropriate set of security functionality criteria in a particular device, for instance, use of multi-factor authentication instead of just passwords or PINs in authentication. Effective and appropriate sets of security functionality criteria used in the case study have been

taken from our framework in [12], which is based on the most relevant standards: ENISA [10], OWASP [19], Industrial Internet Consortium [25], Cloud Security Alliance[8], and Broadband Internet Technical Advisory Group [6], as well as security and privacy guidelines from ISO [1, 2], ENISA [11], and NIST [24].

6.4 Pacemaker Case Study

In previous work [12], we have built a methodology for looking at the functionalities from both security and privacy points of view. In order to understand how we can use the functionality framework to improve security and privacy of IoT systems in practice, we here use a case study of health products involving a pacemaker and related control units such as a mobile phone and a heart rate sensor. In the domain of health services, the highest security class is recommended for devices like pacemakers that directly control life functions of a patient.

A pacemaker is a medical device that is implanted under the skin to help with abnormal behaviors of heartbeats [7, 15]. It consists of a battery, a computerized generator, and wires with sensors at their tips [14]. The generator is powered by the battery, and both are surrounded by a thin metal box. The generator is connected to the heart by the wires. The pacemaker helps to monitor and control the heartbeat. The sensors detect the heart's electrical activity and send data through the wires to the computer in the generator. If the heart rhythm is abnormal, the generator will be directed by the computer to send electrical pulses to the heart, and the pulses travel through the wires to reach the heart. Embedded microprocessors in modern pacemakers have enabled them to do additional tasks like monitoring heart activity and providing a record for the patients and their healthcare providers, as well as collecting data on heart functions to help doctors to identify and diagnose patient conditions, and send required shock signals when needed. Doctors can monitor the patient's heart activities and control the pacemaker using a mobile phone or a computer device, or send required shock signals in case of observation of abnormal behavior.

For each device we will below discuss the security and privacy challenges for each of the three devices. In each case we discuss connectivity, protection level and relevant functionality criteria. We

use the general criteria given in the functionality framework, select and discuss the parts relevant for each device.

Pacemaker security controller. Beside threatening patients' lives, malicious attackers can get access to patients' medical records through a pacemakers [9, 29], or track a patient's location. In addition, malicious software can be run on pacemakers and cause security and privacy breaks. Therefore, any security or functional weakness can result in a security failure. Like any device that uses remote technology, pacemakers are also vulnerable against cyber attacks, and hackers can break into the pacemaker itself, the back-end systems or the communication between the pacemaker and its surrounding. By breaking into a pacemaker, an attacker can send strong shock signals, disturb the pacemaker setting or heart functions, or disturb them from working properly. One simple example of hacking into a pacemaker is to change the setting from battery-saving "sleep" mode to "standby" mode, and this can quickly drain the battery, which is normally supposed to last for years. Hence, security in this kind of device is crucial and should have the highest (best) security class, meaning security class A.

We consider Connectivity 2 (explained in Section. 6.3) in the pacemaker security controller, since it is only connected to the pacemaker. We define below the protection levels which are relevant <for the pacemaker security controller:

- Protection level 1 (P1): includes Secure authentication, Securing Software/Firmware, Secure Communication, and Human Interface Security. For authentication we consider: requiring passwords, option to change the default username and password. For communication we consider: data authenticity to enable reliable exchanges from data emission to data reception. For securing software/firmware we consider: update capability for *some* of the system devices and applications, transmitting the files using encryption.
- Protection level 2 (P2): includes P1 and in addition, for authentication: requiring strong passwords, securing password recovery mechanisms, making sure that default passwords and even default usernames are changed during the initial setup, and that weak, null or blank passwords are not allowed. For communication: verifying any interconnections, discover, identify and verify/authenticate the devices connected to the network before trust can be established,

and preserve their integrity for reliable solutions and services, prevent unauthorised connections to it or other devices the product is connected to, at all protocol levels, providing communication security using state-of-the-art mechanisms, standardising security protocols, such as TLS for encryption. For securing software/firmware: encrypting update files for *some* of the applications.

- Protection level 3 (P3): includes P2 and in addition, for authentication: having options to force password expiration after a specific period, and to change the default username and password, making sure that the password recovery or reset mechanism are robust and do not supply an attacker with information indicating a valid account. The same should apply to key update and recovery mechanisms. For communication: data authenticity to enable reliable exchanges from data emission to data reception.
- Protection level 4 (P4): includes P3 and in addition, for authentication: implementing two-Factor Authentication (2FA), making sure that default passwords and even default usernames are changed during the initial setup. For communication: signing the data whenever and wherever it is captured and stored, making intentional connections, disabling specific ports and/or network connections for selective connectivity. For securing software/firmware: capability of quick updates when vulnerabilities are discovered for *some* of the system devices and applications, and offering an automatic firmware update.
- Protection level 5 (P5): includes P4. In addition, for authentication: using Multi-Factor Authentication (MFA) (considering biometrics for authentication), considering Certificate-Less Authenticated Encryption (CLAE) and User Managed Access (UMA). For communication: rate limiting – controlling the traffic sent or received by a network to reduce the risk of automated attacks. For securing software/firmware: update capability for *all* system devices and applications, capability of quick updates when vulnerabilities are discovered for *all* system devices and applications, encrypting update files for *all* applications, signing update files and validating by the device before installing, securing update servers, having ability to implement scheduled updates, having backward compatibility of firmware updates.

According to Tables 6.1 and 6.2, protection level 4 or 5 are needed

to obtain security class A. We therefore consider how to obtain that protection level, and select the following set of relevant functionality criteria from the functionality framework given in [12], adapted to the challenges of pacemakers. In this selection we have used the guidelines for securing pacemakers from [13]. This gives a certain guarantee that we cover all relevant aspects.

- *Secure authentication/authorization*: Access to the pacemaker security controller and mobile phone connected to the doctor, should be limited using authentication of users (for example user ID and password, Personal Identification Numbers (PINs), biometric authentications). Authorization refers to checking necessary permissions of an identified individual to do an action. Authentication and authorization are completely related to each other. Authorization checks should immediately be followed by authentication of a request. In order to have secure authorization, the roles and permissions of the authenticated user should only be verified through information in back-end systems, not through roles or permission information coming from the device. Any incoming identifiers with a request alongside should be verified by the back-end code independently. Failure in a secure authentication/authorization would give access to unauthorized people and could lead to reputational damages, fraud, unauthorized access to information, information theft, and modification of data.
- *Securing software/firmware*: Before any software or firmware update, user authentication or other suitable controls should be required. Software/firmware updates should be restricted to authorised code. Manufacturers may consider code signature verification as an authentication method.
- *Secure communication*: Data transmission between the pacemaker security controller and the mobile phone connected to the doctor must be secure enough so that a third party cannot listen to their communication. The communication should not be vulnerable to eavesdropping or interception. Failure in having a secure communication can cause identity theft, fraud, data modification, or privacy information leakage. One should consider strong handshaking, correct SSL versions, no clear-text communication of sensitive assets, etc.

- *Human interface security*: Patients should be trained about the importance of security and privacy and how to use the pacemaker properly to ensure security is managed at an appropriate level. For instance, if we consider all the security protections in the highest level in the security controller, but the patient does not know how to deal with error notifications (restart, turn off or low battery errors) in the security controller, all the security considerations in the controller will be useless.
- *Data privacy*: Measures to avoid risk of breaches in connection with long term storage of private information, handling of encryption of private information, and GDPR compliance including consent, purpose, and access rights.

The final discussion of the security class of the pacemaker system depends on the scenario chosen and the other components involved. We next consider the mobile phone.

Mobile phone. In our case study, a mobile phone is used in the communication between the pacemaker and the healthcare provider/doctor. Security in this mobile phone is then important in order to send correct data to/from the pacemaker controller. Hacking or tampering into this mobile phone can result in sending wrong data from the pacemaker to the doctor, something that could result in wrong decisions from the doctor, or possibly sending inappropriate shocks to the patient's heart.

Therefore, it is important to secure the mobile phone properly. However, the security class of the mobile phone/computer device is only considered class B, since mobile phones are inherently not of the highest security class and have a number of possibilities for attacks due to high exposure. For instance, because of all the applications and browsers on the device, as well as software developed by third parties. Therefore, both the impact of attack as well as the connectivity are in a higher level for the mobile phone compared to the pacemaker. Hence the security class of the mobile phone is lower than the security class of the pacemaker security controller. We consider C4 in the mobile phone connected to the doctor, and based on that the following set of relevant security functionality criteria are selected to ensure that the mobile phone has a high protection level, following the functionality framework of [12] and the OWASP guidelines for securing mobile phones [20]:




Figure 6.1: Scenario 1: pacemaker along with security controller

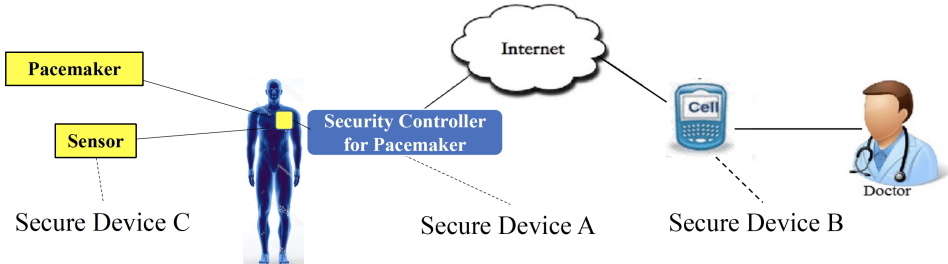


Figure 6.2: Scenario 2: pacemaker with separate security controller

Table 6.3: Security and privacy challenge comparison of scenarios 1 and 2

| | Security and Privacy Challenges | Scenario 1 | Scenario 2 |
|---|---|--------------|---------------|
| In Sensor | Draining battery by changing battery saving controls | Moderate | Insignificant |
| | Interrupting into heartbeat capturing | Catastrophic | Minor |
| | Data and private information breaches | Minor | Minor |
| In Pacemaker | Sending/Receiving wrong data to/from mobile phone connected to the doctor | Moderate | Insignificant |
| | Changing pacemaker shock settings | Moderate | Insignificant |
| | Draining battery by changing battery saving controls | Moderate | Insignificant |
| | Higher battery usage | Catastrophic | Insignificant |
| | Data and private information breaches | Moderate | Insignificant |
| | Transparency of GDPR compliance | Moderate | Insignificant |
| | Risk of long-term storage of private information | Moderate | Insignificant |
| In Mobile Phone Connected to the Doctor | Sending/Receiving wrong data to/from pacemaker | Minor | Minor |
| | Sending wrong unnecessary/necessary shocks | Minor | Minor |
| | Changing pacemaker shock settings | Minor | Minor |
| | Draining pacemaker battery by changing battery saving controls | Minor | Minor |
| | Data and private information breaches | Minor | Minor |
| | Transparency of GDPR compliance | Minor | Minor |
| | Risk of long-term storage of private information | Minor | Minor |

Severity of Challenge: Insignificant: Minor: Moderate: Major: Catastrophic:

- *Secure authentication/authorization and secure communication* are the same as what we discussed for the pacemaker.
- *Sufficient cryptography techniques*: Appropriate cryptography techniques should be considered for instance using AES instead of

DES.

- *Code tampering*: When an application is on the device, the code and data resources are also available there. An attacker can modify the code through either malicious apps in third party app stores or trick the user via phishing attacks and install the app on the device. Code tampering could result in revenue loss due to piracy, reputational damage, unauthorized new features, identity theft or fraud.
- *Secure data storage*: Failure in having secure data storage can result in data loss, extraction of sensitive data using malware, modified apps or forensic tools, identity theft, fraud, reputation damage, material loss or external policy violations.
- *Proper platform usage*: Misuse of a platform feature or failure to use platform security controls fall under this category. Android intents, platform permissions, misuse of TouchID, Keychains, or other security controls which are part of the operating system could be included. To prevent the attacks in this category, secure coding and configuration practices must be applied on the server side of the application.
- *Data privacy*: are as discussed above for the pacemaker.

We next consider the protection levels of ISA99, and select the relevant criteria from the functionality framework, and adjust them for the case of the mobile phone connected to the doctor. These are defined below:

- Protection level 1 (P1): includes Secure Authentication, Secure Authorization, Securing software/firmware, Sufficient Cryptography Techniques, Code Tampering, Secure Data Storage, Secure Communication, and Proper Platform Usage. Secure authentication, Securing software/firmware, and Secure Communication are the same as the considerations for protection level 1 for the pacemaker security controller.
- Protection level 2 (P2): includes P1. Secure authentication, Securing software/firmware, and Secure Communication are the same as the considerations for protection level 2 for the pacemaker security controller. In addition, for cryptography techniques we should consider: ensuring proper selection of standard encryption algorithms and keys.

- Protection level 3 (P3): includes P2. Secure authentication, Securing software/firmware, and Secure Communication are the same as the considerations for protection level 3 for the pacemaker security controller. In addition, for cryptography techniques we should consider: strong encryption algorithms, strong keys.
- Protection level 4 (P4): includes P3. Secure authentication, Securing software/firmware, and Secure Communication are the same as the considerations for protection level 4 for the pacemaker security controller. In addition, for cryptography techniques we should consider: verifying the robustness of the implementation, establishing secure and scalable key management. And cryptographic keys must be securely managed.
- Protection level 5 (P5): includes P4. Secure authentication, Securing software/firmware, and Secure Communication are same as the considerations for protection level 4 for the pacemaker security controller. In addition, for cryptography techniques we should consider: disabling insecure protocols.

Finally, the security class of the heart rate sensor in our case study is considered below.

Heart rate sensor. The heart rate sensor does not have any external connectivity; it has only connections to the heart and computerized generator inside the pacemaker in order to transfer the captured heart rate from the heart to the pacemaker generator. So the connectivity in the sensor is C1. Furthermore it has a chip set from the company provider. Hence, the sensor has the least chance of attack. However, there is still a possible vulnerability if the wired connection is not sufficiently shielded and allows eavesdropping by inductive sensors or senders. But this requires very short physical distance. Therefore we may assume only minor impact, and do not need a high protection level for this device, and may use P1. Exposure is then E4 (see Table 6.4). This gives security class C.

Pacemaker scenarios. The challenge in the case study is how to implement the security controls like decryption, authentication, etc. For instance, whether to do data decryption in the mobile phone, and then send the decrypted data to the pacemaker. If the mobile phone is compromised, wrong data and signals could then go to the doctor and pacemaker, since mobile phones can get compromised, as we discussed before. Whereas, if we consider security controls in the pacemaker itself,

we only have software provided by one company. Thus it would be much more secure to do the security controls directly in the pacemaker, and use the mobile phone just as a gateway to transfer the information. However, this will require additional computational power and battery capacity. Another issue is whether we can do all the security controls inside the computerized generator of the pacemaker – or consider a security controller as a separate unit out of the body with a close and secure connection to the pacemaker.

We therefore define two scenarios: In scenario 1 (see Figure. 6.1), the security controls for the pacemaker are done inside the computerized generator of the pacemaker, and in scenario 2 (see Figure. 6.2), a separate security controller unit makes the security controls of the pacemaker, such that this unit is outside the body with a close and secure connection to the pacemaker.

In scenario 1, we would need computational power inside the body, needing more storage, stronger CPU, much more battery capacity and so on. This is not desirable since it might increase the potential necessary surgeries in order to change the battery, maintain, or update the pacemaker. Moreover, we might not be able to consider some of the security functionalities in order to avoid increasing CPU usage, which results in even more battery usage. Therefore, we might want to have a pacemaker with a simple sensor inside the body and a security controller out of the body, say in the pocket or at home very close to the pacemaker. Here, it is essential that the security controller is close to the pacemaker, since the pacemaker must have very weak signals, limited interactions and computations, to avoid using too much battery power and resulting battery changes. So, all the security and battery-intensive controls would be done in the external security controller, which can have a stronger and easily rechargeable battery. The controller can maintain the device in case of any problem, update or troubleshoot its computer system, or even increase the level of protection by adding more security and privacy functionalities or appropriate software at any time because of easy access. According to all the security considerations in scenario 2, we can then have a better security class in the pacemaker security controller: In Table 6.3 we summarize all the security and privacy challenges in the sensor, pacemaker, and mobile phone connected to the doctor, and compare the severity of these security and privacy challenges in scenarios 1 and 2. In the next section we see that the severity of the security and

privacy challenges in the sensor and pacemaker has been reduced from very high in the worst case to low. Hence, the impact of attacks is reduced from catastrophic to minor.

6.5 Discussion

The pacemaker in scenario 1 is a complex device because all security controls are done inside the computerized generator of the pacemaker. The security controller is very close to the sensor, the signals received by the sensor from the pacemaker are very frequent. Interference is possible and that can have negative effect on correct heartbeat capturing. However, by transferring the security controller outside of the body this effect would be low (still there are some frequencies from devices close by such as mobile phones that can affect the sensor) resulting in more precise heartbeat capturing. As mentioned earlier in the heart rate sensor, there are low possibilities of data and private information breaches in both scenarios.

In the pacemaker in scenario 1, because of the complexity of the system and having all the security controls inside the computerized generator

Table 6.4: Security class of the sensor

| | | | | | | |
|------------|--------------|----|----------|---------------|---------|---------------|
| Protection | P1 | E4 | Impact | Catastrophic | Class F | In Scenario 1 |
| | P2 | E3 | | Major | Class E | |
| | P3 | E2 | | Moderate | Class E | |
| | P4 | E1 | | Minor | Class C | |
| | P5 | E1 | | Insignificant | Class B | |
| | | C1 | | | E4 | In Scenario 2 |
| | Connectivity | | Exposure | | | |

Table 6.5: Security class of the pacemaker

| | | | | | | |
|------------|--------------|----|----------|---------------|---------|---------------|
| Protection | P1 | E4 | Impact | Catastrophic | Class D | In Scenario 1 |
| | P2 | E3 | | Major | Class B | |
| | P3 | E2 | | Moderate | Class B | |
| | P4 | E1 | | Minor | Class B | |
| | P5 | E1 | | Insignificant | Class A | |
| | | C2 | | | E2 | In Scenario 2 |
| | Connectivity | | Exposure | | | |

Table 6.6: Security class of the mobile phone

| | | | | | | |
|-------------------|---------------------|----|-----------------|---------------|---------|-----------------------|
| Protection | P1 | E5 | Impact | Catastrophic | Class D | In Scenarios 1 & 2 |
| | P2 | E4 | | Major | Class B | |
| | P3 | E3 | | Moderate | Class B | |
| | P4 | E2 | | Minor | Class B | |
| | P5 | E1 | | Insignificant | Class A | |
| | C4 | | E2 | | | |
| | Connectivity | | Exposure | | | |

of the pacemaker, we need higher CPU and memory consumption and then battery usage would be very high, while in scenario 2 this problem would decrease to very low. In scenario 1, because of difficulties in accessing the pacemaker and its security controller on time (in case of maintenance, update, troubleshooting or installing new software/equipment, also not being able to apply all the necessary criteria with high protection level in order to avoid high battery usage), the level of the security and privacy functionality criteria as well as the level of the protection is low. Therefore the vulnerability of the pacemaker and sensor against attacks compromising the device (that can cause tampering of transferred information from/to the pacemaker, changing pacemaker shock settings, changing battery saving controls, data and private information breaches, transparency of GDPR compliance, risk of long-term storage of private information) is high; however, this problems decreases to very low when we change to scenario 2, due to easier access to the security controller out of the body.

By compromising the mobile phone connected to the doctor, the problems listed in the last part of Table 6.3 may occur. In both scenarios, we have considered high level of protection for the mobile phone by using appropriate set of security and privacy functionality criteria with high protection level, therefore we have reduced the vulnerability of the device to low, but as discussed earlier, these devices still have a number of vulnerabilities for attacks, and therefore the vulnerability is not in a very low level.

Furthermore, softening the problems in scenario 1 results in obtaining higher protection level and security class. For instance, we can easily recharge the battery of the security controller, maintain the device in case of any problem to avoid shutting down all the security considerations, and increase the level of protection by updating or troubleshooting its

computer system.

In the sensor, as explained above, we have connectivity C1, and protection level P1, therefore exposure is E4 (see Table 6.4), and because of very high problem severity, the impact of attacks is Catastrophic, resulting in security Class F in this device. By changing from scenario 1 to 2, the severity of the security and privacy challenges has reduced from very high to low, therefore the impact of attacks is reduced from Catastrophic to Minor, and consequently, the security class has improved from Class F to Class C.

In the pacemaker, we have connectivity C2, and protection level P5, however in scenario 1, the severity of security and privacy challenges have affect on the protection level, which falls to level 3, therefore exposure is E2 (see Table 6.5), and because of very high security and privacy challenge severity, the impact of attacks is Catastrophic, then we have security Class D in this device. By changing from scenario 1 to 2, the severity of the security and privacy challenges has reduced from very high to very low, therefore the impact of attacks has reduced from Catastrophic to Insignificant, and consequently, the security class has improved from Class D to A.

And, in the mobile phone connected to the doctor, we have connectivity C4, and protection level P4, therefore the exposure is E2 (see Table 6.6), and the security is Class B in both scenarios.

By considering security and privacy challenges, as well as the effect of these challenges on the protection level of each device and the whole system, and also their affect on the impact of attacks, we have changed scenario 1 to 2, and were able to improve the security class in the sensor and pacemaker security controller from class F to C, and class D to A, respectively. Hence the security of the overall system has improved significantly.

6.6 Conclusion

The expansion of IoT in the last decade has resulted in several security and privacy vulnerabilities and attacks against IoT devices and people. Unfortunately, the security and privacy functionalities to combat these attacks are not well-recognized in the domain of IoT. We previously presented a new taxonomy framework that addresses all of the IoT

security and privacy functionalities. In this paper, we demonstrate the application of this framework in combination with the security classification method, using a case study of pacemaker as medical device communicating with the surrounding, and discuss the various security and privacy challenges associated with two scenarios.

Our analysis results demonstrate that using our security functionality framework, the security class in the sensor and pacemaker has been improved from class F to C, and class D to A, respectively. Therefore, the security of the overall system has improved. The main objective of this paper is to give security developers, designers, and end-users an opportunity to explore what the IoT security and privacy functionalities are, and how these functionalities can help to improve security and privacy of IoT systems. In future work, we will investigate the applications of our framework to other IoT domains.

Acknowledgment

This work has been supported by the research project *IoTSec*, Security in IoT for Smart Grids, with number 248113/O70 part of the IKTPLUS program funded by the Research Council of Norway. It is also supported by the *SCOTT* project (www.scott-project.eu), funded by the Electronic Component Systems of the European Leadership Joint Undertaking under grant agreement No 737422, with support from the EU H2020 research and innovation programme.

Bibliography

- [1] 27000, I. *Information technology — Security techniques — Information security management systems — Overview and vocabulary (fourth edition)*. ISO/IEC 2016. 2016.
- [2] 27001, I. *INTERNATIONAL STANDARD ISO/IEC 27001 Information technology—Security techniques —Information security management systems —Requirements*. ISO/IEC 2013. 2013.
- [3] Alqassem, I. and Svetinovic, D. “A taxonomy of security and privacy requirements for the Internet of Things (IoT)”. In: *2014 IEEE Intern. Conf. on Industrial Engineering and Engineering Management (IEEM)*. IEEE. 2014, pp. 1244–1248.
- [4] ANSSI. *Classification Method and Key Measures*. 2014.
- [5] Babar, S. et al. “Proposed security model and threat taxonomy for the Internet of Things (IoT)”. In: *Intern. Conf. on Network Security and Applications*. Springer. 2010, pp. 420–429.
- [6] BITAG. *Internet of Things (IoT) Security and Privacy Recommendations*. [https://www.bitag.org/documents/BITAG_Report_-_Internet_of_Things_\(IoT\)_Security_and_Privacy_Recommendations.pdf](https://www.bitag.org/documents/BITAG_Report_-_Internet_of_Things_(IoT)_Security_and_Privacy_Recommendations.pdf). 2016.
- [7] Boulos, P. et al. “Pacemakers: a survey on development history, cyber-security threats and countermeasures”. In: *Int. J. Innov. Stud. Sci. Eng. Technol* vol. 2, no. 8 (2016).
- [8] *Future-proofing the Connected World: 13 Steps to Developing Secure IoT Products*. <https://downloads.cloudsecurityalliance.org/assets/research/internet-of-things/future-proofing-the-connected-world.pdf>. 2016.
- [9] Dadourian, M. *Heart alert: Pacemakers can be hacked*. 2018.
- [10] ENISA. *Baseline Security Recommendations for IoT in the context of Critical Information Infrastructures*. European Union Agency for Network and Inf. Sec. 2017.

- [11] ENISA. *Guidelines for SMEs on the security of personal data processing*. <https://www.enisa.europa.eu/publications/guidelines-for-smes-on-the-security-of-personal-data-processing>. 2017.
- [12] Fazeldehkordi, E., Owe, O., and Noll, J. *Security and Privacy Functionalities in the Internet of Things (long version)*. Tech. rep. University of Oslo, 2019.
- [13] FDA. *Content of Premarket Submissions for Management of Cybersecurity in Medical Devices Guidance for Industry and Food and Drug Administration Staff*. U.S. Food and Drug Administration (FDA). 2014.
- [14] Health & Human Services, U. D. of. *Pacemaker*. 2019.
- [15] Healthline. *Heart Pacemaker*. 2019.
- [16] Islam, S. R. et al. “The Internet of Things for health care: a comprehensive survey”. In: *IEEE Access* vol. 3 (2015), pp. 678–708.
- [17] Lee, C. et al. “Securing smart home: Technologies, security challenges, and security requirements”. In: *Communications and Network Security (CNS), 2014 IEEE Conf. on*. IEEE. 2014, pp. 67–72.
- [18] MacGillivray, C., Turner, V., and Lund, D. *Worldwide Internet of Things (IoT) 2013–2020 Forecast: Billions of Things, Trillions of Dollars*. International Data Corporation. <http://www.idc.com/getdoc.jsp?containerId=243661f>. 2014.
- [19] OWASP. *IoT Security Guidance*.
- [20] OWASP. *OWASP Mobile Security Project*.
- [21] Pang, Z. “Technologies and Architectures of the Internet-of-Things (IoT) for Health and Well-being”. PhD thesis. KTH Royal Institute of Technology, 2013.
- [22] Pishva, D. “Internet of Things: Security and privacy issues and possible solution”. In: *Advanced Communication Technology (ICACT), 2017 19th Intern. Conf. on*. IEEE. 2017, pp. 797–808.
- [23] Roman, R., Najera, P., and Lopez, J. “Securing the Internet of Things”. In: *Computer* vol. 44, no. 9 (2011), pp. 51–58.

- [24] Ross, R. et al. “Protecting controlled unclassified information in nonfederal information systems and organizations”. In: *NIST Special Publication* vol. 800 (2015), p. 171.
- [25] Schrecker, S. and al., et. *Industrial Internet of Things Volume G4: Security Framework*. Industrial Internet Consortium. https://www.iiconsortium.org/pdf/IIC_PUB_G4_V1.00_PB.pdf. 2016.
- [26] Shrestha, M., Johansen, C., and Noll, J. “Security Classification for Smart Grid Infra structures (long version)”. In: (2017).
- [27] Sicari, S. et al. “Security, privacy and trust in Internet of Things: The road ahead”. In: *Computer networks* vol. 76 (2015), pp. 146–164.
- [28] Suo, H. et al. “Security in the Internet of Things: a review”. In: *Computer Science and Electronics Engineering (ICCSEE), 2012 Intern. Conf. on*. Vol. 3. IEEE. 2012, pp. 648–651.
- [29] Thehackernews.com. *Over 8,600 Vulnerabilities Found in Pacemakers*. <https://thehackernews.com/2017/06/pacemaker-vulnerability.html>.

Paper 3: A Language-Based Approach to Prevent DDoS Attacks in Distributed Financial Agent Systems

Authors: Elahe Fazeldehkordi, Olaf Owe, Toktam Ramezanifarkhani

Publication: Presented at Security for Financial Critical Infrastructures and Services (FINSEC), part of the 24th European Symposium on Research in Computer Security (ESORICS) Conference, 23-27 September 2019. Published in Lecture Notes in Computer Science, vol 11981 (pp. 258-277). Springer. DOI: https://doi.org/10.1007/978-3-030-42051-2_18

Abstract

Denial of Service (DoS) and Distributed DoS (DDoS) attacks, with even higher severity, are among the major security threats for distributed systems, and in particular in the financial sector where trust is essential.

In this paper, our aim is to develop an additional layer of defense in distributed agent systems to combat such threats. We consider a high-level object-oriented modeling framework for distributed systems, based on the actor model with support of asynchronous and synchronous method interaction and futures, which are sophisticated and popular communication mechanisms applied in many systems today. Our approach uses static detection to identify and prevent potential vulnerabilities caused by asynchronous communication including call-based DoS or DDoS attacks, possibly involving a large number of distributed actors.

7.1 Introduction

Today distributed and service-oriented systems form critical parts of

infrastructures of the modern society, including financial services. In the financial sector security and trust are essential for users of financial services [19]. Security breaches may lead to significant loss of assets, such as physical or virtual money, including cryptocurrencies and bitcoins. In addition, successful attacks on services of a financial institution may damage the trust of customers, which indirectly may hurt the institution [18]. According to [1, 12, 17], a main threat on financial institutions is Distributed Denial of Service (DDoS) attacks. Protection against DoS/DDoS attacks is therefore crucial for financial institutions. Slow website responses caused by targeted attacks, can imply that customers cannot access their online banking and trading websites during such attacks. Both network layer and application layer DDoS attacks continue to be more and more persistent according to a report from the Global DDoS Threat Landscape Q4 2017. Based on this report, it becomes easier and easier to launch DDoS attacks, and one may even purchase botnet-for-hire services that provide the basis for starting a hazardous DDoS attack. Financial institutions are recommended to monitor the internet traffic to their websites in order to detect and react to possible threats. However, this kind of run-time protection may slow down or temporarily shut down the websites.

Unintended attacks on customers from a financial institution may easily destroy the customer's trust and confidence and result in reputation damage. If customers cannot trust an institution, they may quickly shift to a different institution, due to competition between the many different financial institutions and service providers. Even a single unfortunate incident of a financial service provider could be enough to influence customers. One should make sure that the software is not harmful for the customers before running it, and this makes static (compile-time) detection more important than in other areas. Thus, in the financial sector static detection is a valuable complement to run-time detection methods, and seems underrepresented.

Call-based *flooding* is commonly seen in the form of application-based DDoS attacks [6]. To prevent DoS/DDoS flooding attacks in a manner complementary to existing approaches, we propose an additional layer of defense, based on language-based security analysis. We focus on DDoS attacks that try to force a (sub)system out of order by flooding applications running on the target system, or by using such applications to drain the resources of their victim.




Figure 7.1: Distributed communication (s-obj stands for server object and c-obj for consumer object)

In this paper, we consider a high-level imperative and object-oriented language for distributed systems, based on the actor model with support of asynchronous and synchronous method interaction. This setting is appealing in that it naturally supports the distribution of autonomous concurrent units, and efficient interaction, avoiding active waiting and low-level synchronization primitives such as explicit signaling and lock operations. It is therefore useful as a framework for modeling and analysis of distributed service-oriented systems. Our language supports efficient interaction by features such as asynchronous and non-blocking method calls and first-class futures, which are popular features applied in many distributed systems today. However, these mechanisms make it even easier for an attacker to launch a DDoS attack, because undesirable waiting by the attacker can be avoided with these mechanisms.

We propose an approach consisting of static analysis. We identify and prevent potential vulnerabilities in asynchronous communication that directly or indirectly can cause call-based flooding of agents. More precisely, we adapt a general algorithm for detecting call flooding [14] to the setting of security analysis and for detection of distributed denial of service attacks adding support for many-to-one attacks. The algorithm detects call cycles that might overflow the incoming queues of one or more communicating agents. Each cycle may involve any number of agents, possible involving the attacked agent(s).

The high-level framework considered here is relevant for a large class of programming languages and service-oriented systems.

Outline. Section. 7.2 describes the background of the problem. Related work is discussed in Section. 7.3. The active object framework is

explained in Section. 7.4. Our static analysis to prevent attacks is described in Section. 7.5. Examples of possible DoS/DDoS attacks are given in Section. 7.6. The final section concludes and suggests future work.

7.2 Overview

In distributed system communication there is an underlying distributed object system as shown in Figure. 7.1. In such a distributed system, classes such as server or client classes would be instantiated by objects, and communication is established in the form of method calls, usually wrapped in XML or other forms. Therefore, communication in a distributed system is implemented by method calls between objects. If there is a possibility of flood of requests to the service provider (S-Obj) from the consumer object(s) (C-Obj) in this figure, a DoS attack is probable.

Call-based flooding attacks. To launch a DoS attack, the attacker may try to submerge the target server under many requests to saturate its computing resources. To do so, flooding attacks [6, 20] by method calls are effective, especially when the server allocates a lot of resources in response to a single request. Therefore, we detect:

- call-flooding: flooding from one object to another.
- parametric-call-flooding: flooding from one object to another when the target object allocates resources or consumes resources for each call.

In the case of call-flooding, communications are just simple requests like a simple call without parameters or parameters that do not lead to resource consumption. Parametric-call-flooding is when requests usually include parameters in a non-trivial manner. Such requests usually trigger relatively complex processing on the server such as access to a database. Parametric-call-flooding is more effective than call-flooding because it takes fewer requests to drown the target system. However, call-flooding are more common and easier for attackers to exploit.

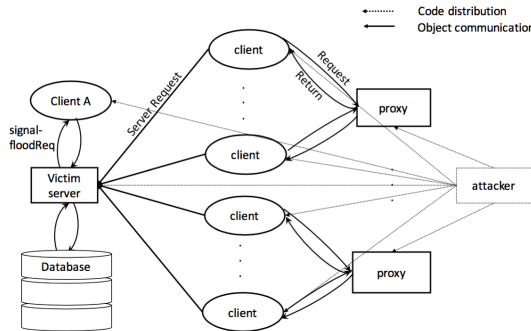


Figure 7.2: Distributed object communication in DDoS

Categories of call-based flooding attacks: DoS or DDoS possibilities.

One-to-one (OTO): If thousands of requests every single second come from one source object to a target object, then it is a *one-to-one (OTO) DoS attack*. The intent of the flooding might be malicious, or even undeliberate call cycles. Communication between Client A and the victim server in Figure. 7.2 is an example of this attack.

Many-to-one (MTO): If the incoming flooding traffic originates from many distinguishable different sources, then it is a *many-to-one (MTO) DDoS attack*. Figure. 7.2 shows a distribution of code between clients and a server, and proxies.

One-to-many (OTM): A *one-to-many attack* appears if a system makes an unlimited number of requests to many objects simultaneously. Such an attack can be serious since many target objects are attacked at the same time.

Static attack detection and prevention. For any set of methods that call the same target method, a call cycle could be harmful. The methods might belong to the same or different objects with the same or different interfaces. In the case of normal blocking calls, where the caller is blocking while waiting for the response, making a flood of requests also means receiving a flood of responses. And thus in the case of OTO, it may cause a self DoS for the attacker. With the possibility of non-blocking calls in a distributed setting, it is more cost-beneficial for an OTO attacker to launch a DoS, because then undesirable blocking by the attacker is

avoided. By means of futures and asynchronous calls, a caller process can make non-blocking method calls.

The possibility of unbounded object creation, referred to as *instantiation flooding* [8], could cause resource consumption and DoS that could be detected statically, especially if those objects and their communication can cause flooding requests from the bots, such as the customers in our example. It is even worse if there is instantiation flooding on the target side of the distributed code. This can be detected by static analysis of the target. (See the example in Figure. 7.10.) Our static analysis detects explicit or implicit call-flooding. Static detection is accomplished by static analysis at compile time and informs the programmer about the possibility of program exploitation at runtime.

7.3 Related Work

A DoS attack, or its distributed version, happens when access to a computer or network resource is intentionally blocked. Considering the exploited vulnerabilities, these attacks might be classified by resource consumption attacks or flooding attacks, of which the latter category is the most common [6]. In this paper, we aim to prevent distributed code to be exploited by attackers to launch a DoS attack by detection of possible call-based flooding in both of the target and zombie sides. To do so, we analyze the distributed code to make an additional layer of defense against DoS or DDoS attacks.

In the following, we discuss related works for preventing application-based DDOS attacks using static detection. In the paper presented by Chang et al. [3], a novel static analysis approach was introduced in order to detect semantic vulnerabilities in networked software that might cause denial of service attacks because of resource exhaustion. Their approach is implemented in a tool named SAFER: Static Analysis Framework for Exhaustion of Resources. SAFER integrates taint analysis (in order to compute the group of program values that are data-dependent on network inputs) and control dependency analysis (for computing the group of program statements whose execution can affect the execution of a given statement) toward detecting high complexity control structures that can be caused by untrusted network inputs. The tool applies the CIL static analysis framework and combines different heuristics for recognizing

loops and recursive calls. Compared to our work the SAFER approach is oriented toward detecting server attacks from within the server code, whereas our approach is mainly targeting server attacks from an external attacker, or a combination of external agents. An attacker needs to understand the code of the server in order to find weaknesses that can be triggered by specific inputs. In contrast, our approach is detecting attacks caused by coordination of several agents and/or servers in a distributed setting.

Another work that detects resource attacks from within the server code is presented by Qie et al. [15]. In their toolkit, they check for possible “rule” violations at runtime. This work is complementary to ours, since our work is oriented toward static detection. Gulavani and Gulwani [7] describe a precise numerical abstract domain. This domain can be used to prove the termination of a large class of programs and also to estimate valuable information such as timing bounds. In order to make linear numerical abstract domains more precise, they make use of two domain lifting operations: One operation depends on the principle of *expression abstraction*. This describes a set of expressions and determines their semantics by use of a selection of directed inference rules. It works by picking up an abstract domain and a group of expressions, such that their semantics are described by a group of rewrite rules, in order to construct a more precise abstract domain. The second domain constructor operation picks up a linear arithmetic abstract domain and constructs a new arithmetic domain that is able to represent linear relations through introduction of *max* expressions. Another approach to estimate worst-case complexity is presented by Colon and Sipma [4]. These approaches [4, 7], in which the complexity of loops and recursive calls has been estimated using structural analysis, are widely complementary to our work.

Zheng and Myers [21] propose a framework for using static information flow analysis in order to specify and enforce end-to-end availability policies in programs. They extend the decentralized label model to include security policies for availability. This work presents a simple language with fine-grained information security policies described by type annotations. In addition, this language has a security type system to reason about end-to-end availability policies. Various examples have been discussed, in which abuse of an availability policy can represent denial of service attacks.

In a work by Meadows [13], a formal analysis has been developed in

order to apply the maximum benefit of tools and approaches that have already been used to strengthen protocols against denial of service attack. This analysis has been done at the protocol specification level. Also, different ways in which existing cryptographic protocol analysis tools can be modified for the purpose of operating in this formal framework, have been demonstrated. In contrast, we do a detailed static analysis of source code both inside and outside a server. The class of software vulnerabilities that we can detect is more complicated than what appears just at the network-protocol specifications level. Moreover, vulnerable sections of the source code have been identified in our work.

The current work shows how the static analysis method for detection of flooding can be used for detection of DoS and DDoS attacks. This general idea was also outlined by the same authors in an extended workshop abstract [16]. Moreover we here discuss why this is particularly harmful in the financial sector, where both economic assets and customer trust are at risk. Furthermore we simplify and adapt the static analysis method of [14] to the setting of financial service systems, extending it to detect many-to-one attacks involving unbounded creation of objects (as demonstrated in example 7.10), as well as hidden attacks, neither of which were detected by the original method of [14].

7.4 Our Framework for Active Object Systems

The setting of concurrent objects communicating by asynchronous method calls combines the Actor model and object-orientation, and is referred to as *active objects*. Active object languages are suitable for modeling and implementing distributed applications, letting a distributed system be modeled by a number of active objects that interact via asynchronous method calls. The active object model provides natural description of autonomous agents in a distributed system, and the *future mechanism* provides an efficient communication primitives [2, 11], allowing results computed in a distributed setting to be referred to and shared. Moreover, the addition of *cooperative scheduling*, as suggested in the Creol language [9], allows further communication efficiency, by adding process scheduling control in the programming

language, and passive waiting. This is achieved by including statements for suspension control, and letting each object have a process queue for holding suspended processes. We consider a core language for active objects with future-based communication primitives, inspired by Creol and ABS [10]. The objects are concurrent units distributed over a network, and their identity is globally unique. An object has a process queue, as well as a queue for incoming method call requests, and can perform at most one process (i.e., remaining part of method call) at a time. A process can be suspended by an **await** statement, allowing other (enabled) processes to continue. When a process is ended or suspended, the object may continue with an incoming call request or other enabled process from the process queue (if any). The **await** statement allows a process to wait for a Boolean condition to be satisfied, or for a future value to be available. The statement is enabled when the waiting condition/future is satisfied/available. The **await** statement enables high-level process control, instead of low-level process synchronization statements such as signaling and lock operations.

Our core language is a typed, imperative language. An assignment has the form $x := e$ where the expression e is without side-effects. All object variables (i.e., object references) are typed by an interface, and an interface specifies the set of methods that are visible through that interface. The interfaces of a class protect and limit the object communication, and in particular shared variable interaction is forbidden. Local data structure is made by data type declarations, indicated by **data**, and a functional data type sublanguage is used to create and manipulate data values. Data values are passed by value, while object variables are passed by reference. The language supports first-class futures. The basic interaction mechanisms (by method calls/futures) are as follows:

- $f := o!m(\bar{e})$ – the current object calls method m on object o with actual parameters \bar{e} . A globally unique identity u identifying the call is assigned to the future variable f . A message is then sent over the network from the current object to object o . When object o eventually performs the method and the method gives a result defined by a **return** statement, that result is placed in a (globally accessible) future with identity u , and the future u is then said to be *resolved*. Any process of any object that knows u may access the future value or wait for it to be available.

- $x := \mathbf{get} f$ – this statement blocks until the value of the future f is available, and then that value is assigned to the variable x . (Here f may be an expression resulting in a future identity.)
- $\mathbf{await} c$ – this statement suspends if the Boolean condition c is not satisfied, and is enabled when c is satisfied,
- $\mathbf{await} x := \mathbf{get} f$ – this statement suspends if the value of f is not yet available, and is enabled when the future is available. Then the future value is assigned to x .

The statement sequence $f := o!m(\bar{e}); x := \mathbf{get} f$ corresponds to a traditional blocking call, and is abbreviated $x := o.m(\bar{e})$ using the conventional dot-notation. The statement sequence $f := o!m(\bar{e}); \mathbf{await} f; x := \mathbf{get} f$ is abbreviated $\mathbf{await} x := o.m(\bar{e})$ and corresponds to a non-blocking call, since the **await**-statement ensures that the future is available before the **get**-statement is performed. If the result value is not needed, we may simplify the syntax to $o.m(\bar{e})$ for blocking calls and $\mathbf{await} o.m(\bar{e})$ for non-blocking calls. And if the future is not needed, $f := o!m(\bar{e})$ may be abbreviated to $o!m(\bar{e})$, in which case the future cannot be accessed (since it is not stored in a future variable).

Object creation has the syntax $x := \mathbf{new} C(\bar{e}) \mathbf{at} o$, where the class parameters behave like fields (initialized to the values of \bar{e}) except that they are read-only, and the new object is created locally at the site of object o . With the syntax $x := \mathbf{new} C(\bar{e})$ the new object is located anywhere in the distributed system.

One may refer to the current object by `this` and to the caller object by (the implicit method parameter) `caller`. Self calls are possible by making calls to `this`, and recursion is allowed. *Active behavior* is possible by making a recursive self call in the constructor method (given as a nameless method). By means of suspension, the active self behavior may be interleaved with execution of incoming calls from other objects, thereby combining active behavior and passive behavior. If- and while-statements are as usual.

We assume all class parameters and method parameters (including `this` and `caller`) are read-only. This helps the static analysis by reducing the set of false positives. For methods that return no information we use a predefined type *Void* with only one value, *void*. For simplicity we omit **return void** at the end of *Void* methods in the examples.

7.5 Static Analysis to Prevent Attacks

We base our approach on the static analysis of flooding presented in [14] for detection of flooding of requests, formalized for the Creol/ABS setting with futures. We adapt this notion of flooding to deal with detection of DDoS attacks, which have a similar nature. The static analysis will search for flooding cycles in the code, possibly involving several classes. According to [14] (unbounded) flooding is defined as follows:

Definition 7.5.1 (Flooding). An execution is *flooding with respect to a method m* if there is an execution cycle C containing a call statement to a method m at a given program location, such that this statement may produce an unbounded number of uncompleted calls to method m , in which case we say that the call is *flooding with respect to C* in the given execution.

Like in [14], we distinguish between weak flooding and strong flooding. Strong flooding is flooding under the assumption of so-called *favorable* process scheduling, i.e., enabled processes are executed in a fair manner.

Definition 7.5.2 (Strong and weak flooding). A call is *weakly flooding* with respect to a cycle C if there is an execution where the call is flooding with respect to C . And a call is *strongly flooding* with respect to a cycle C if there is an execution with fair scheduling of enabled processes where the call is flooding with respect to C .

Strong flooding reflects the more serious flooding situations that persist regardless of the underlying scheduling policy. In the detection of strong flooding, a *statically enabled* node is considered strongly reachable if each of its predecessor flow nodes are strongly reachable. All statements are statically enabled, apart from **get/await** statements. A **get** statement or an **await** on a future/call is statically enabled if the corresponding future/result is available, detected statically if the corresponding return statement is strongly reachable or another **get/await** statement on the same future is strongly reachable. We rely on a static under-detection of the correspondence between return statements and futures. In the examples this detection is straight forward. With respect to DDoS, weak flooding of a server is in general harmless unless the flooding is caused by a large enough number of objects. Strong flooding

1. Make separate *control flow graphs* (CFGs) for each method. Include a node for each *call*, *get*, *await*, *new* (for object creation), *if* and *while* statement, as well as an initial *starting* and a final *return* node.
2. Add call edges from call nodes to the start node of a *copy* of the called method. In case the call is recursive, simply add a call edge to the existing start node.
3. Identify any cycles in the resulting graph (including all copies of the CFGs).
4. Assign a unique label to each call node, and assign this label to the *start* and *return* node of the corresponding copy of the method CFG.
5. Make *put* edges from the *return* nodes to the corresponding *get/await* nodes. This requires static flow analysis, possibly with over-approximation of *put* edges.

Figure 7.3: Control flow graph

is dangerous even from a single attacker.

Following [14], flooding is detected by building the control flow graph (CFG) of the program, locating control flow cycles as outlined in Figure. 7.3, and then analyzing the sets of weakly reachable calls, denoted *calls*, and the set of strongly reachable call completions, denoted *comps*, in each cycle. Flooding is reported for each cycle with a nonempty difference between *calls* and *comps*, as explained in Figure. 7.4. Note that the abbreviated notations for synchronous calls and suspending calls are expanded to the more basic call primitives, as explained above. We assume that assignments (other than calls) will terminate efficiently and therefore ignore them in the CFG. For each method the CFG begins with a start node and ends with a return node (even for void methods) – the latter helps in the analysis of method completion. We next define *weakly and strongly reachable nodes*. The detection of strongly reachable nodes uses a combination of forward and backward analysis, and is simplified compared to [14]:

Definition 7.5.3 (Weakly and strongly reachable nodes). Consider a

Consider a cycle C in the control flow graph G resulting from Figure. 7.3:

1. Mark all nodes in C as *strongly-reachable (SR)*, and the rest as (initially) *not reachable*.
2. From the entry point to the cycle, follow all flow and call edges in a depth-first traversal of G and mark the nodes as *weakly-reachable (WR)*, *strongly-reachable (SR)*, or neither, as defined in Def. 7.5.3.
3. If the previous step results in any changes to the *SR* or *WR* node sets, go to step 2.
4. Report flooding of call n if $n \in (calls - comps)$ where $calls = \{n \mid call_n \in WR\}$ and $comps = \{n \mid return_n \in SR \vee get_n \in SR\}$.

Figure 7.4: Algorithm for detecting flooding by means of *calls* and *comps* sets in a given cycle

given cycle C .

Weakly reachable (WR) nodes are those that are on the cycle or reachable from the cycle by following a flow edge or a call edge.

A node is *strongly reachable (SR)* if it is in the given cycle or is reachable from an SR node without entering an if/while node nor passing a wait node (**get/await**) outside the cycle, unless the return node of the corresponding call is strongly reachable. A return node is SR if there is a SR **get/await** node on the same future. And a node is SR if all its predecessor nodes are SR.

We consider two versions of SR, the *optimistic*, where we follow call edges (as indicated above), and the *pessimistic*, where we follow a call edge n only when the call is known to complete, i.e., when $n \in SR$ before following the call edge. (As above we follow flow and put edges without restrictions.)

The optimistic version is used to find unbounded flooding under the assumption of favorable scheduling, i.e., *strong flooding*. The pessimistic version is used to detect unbounded flooding without this assumption, i.e., *weak flooding*. Detection of strong flooding implies detection of weak flooding, but with less precise details about which calls that possibly may cause flooding. If there is a call causing weak flooding wrt. a given cycle,

pessimistic detection will report this call or a call leading to this call. If there is a call causing strong flooding for a given cycle C , optimistic detection will report this call. Our notions of optimistic and pessimistic reachability cover a wider class of nodes than in [14]. The soundness of [14] can be generalized to our setting.

7.6 Examples of Possible DoS/DDoS Attacks

An example of flooding cycles. We consider here an example of a possible DoS attack on customers caused by a financial institution. The attack may be unintended by the institution, but may result from an update supposed to give better efficiency, by use of the future mechanism to reduce the amount of data communicated over the network.

We imagine that the financial institution has a subscription *service* for customers, such that customers can register and receive the latest information about shares and funds, through data of type “newsletter”, here simply defined as a product type consisting of a *content* and a *date*. The financial institution uses a method `signal` to notify the customers about new information about shares. In the first version, each call to `signal` has the newsletter as a parameter. This may result in heavy network traffic and many of the newsletters may not be read by the customers. In the “improved” solution, each call to `signal` contains a reference (by means of a future) to the newsletter rather than the newsletter itself. However, this allows the subscription service system to send `signal` calls even before the newsletter is available, and as we will see, this can cause a DoS attack on the subscribing customers.

In order to handle many customers, a (dynamic) number of proxies are used by the service object, and an underlying newsletter producer is used for the sake of getting newsletters, using suspension when waiting for news. The proxies are organized in a list (*myCustomers*), growing upon need. In both solutions, futures are used by the service object to avoid delays while waiting for a newsletter to be available. In this way the service object can continuously respond to customers. The interfaces are shown in Figure. 7.5. We abbreviate “Newsletter” to “News”. Figure. 7.6 represents a high-level implementation of the publish/subscribe model, adapted from [5]. A multi-cast to each object in the *myCustomers* list is made by the statement *myCustomers!signal(ns)* in line 13. If we shift


```

data News == (String content, Int date) // a product data type
interface ServiceI{
    Void subscribe(CustomerI cl) // called by Clients
    Void produce() // called by Proxies
}
interface ProxyI{
    ProxyI add(CustomerI cl) // called by Service
    Void publish(Fut[News]fut) // called by Service
}
interface ProducerI{
    News detectNews() // called by Service
}
interface NewsProducerI{
    Void add(News ns) // called when news arrives
    News getNews() // called by Producers
}
interface CustomerI{
    Void signal(News ns) // called by Proxies
}

```

Figure 7.5: The interfaces of the units in the subscription example

requiring the actual newsletter to have arrived, from the Proxy (as shown by the statements `ns:=get fut; myCustomers!signal(ns)` in the original `publish` method) to the Customer (i.e., `news:=get fut` in the modified `Customer.signal` method). This change in the program causes flooding of customers.

`Service.produce` asynchronously calls `Producer.detectNews` (Pd), line 5 of Figure. 7.6

`Service.produce` asynchronously calls `Proxy.publish` (Xb), line 5 of Figure. 7.6

`Proxy.publish` asynchronously calls `Customer.signal` (Cs), line 6 of Figure. 7.7

`Proxy.publish` asynchronously calls `Service.produce` (Sp) line 6 of Figure. 7.7 Each iteration of this cycle generates an asynchronous call to `Proxy.publish`, which again produces an asynchronous call to `Producer.detectNews`, which is not processed as part of this cycle, nor is its processing synchronized call with the cycle. An unbounded number of suspended calls to `Producer.detectNews` can be produced by this cycle.

```

class Service(Int limit, NewsProducerI np) implements ServiceI {
  ProducerI prod; ProxyI proxy; ProxyI lastProxy; // declaration of fields
  { // constructor
    prod := new Producer(np); proxy:= new Proxy(limit,this);
    lastProxy:=proxy; this!produce() }
  Void subscribe(CustomerI cl) { lastProxy:=lastProxy.add(cl) }
  Void produce() { Fut[News] fut :=prod!detectNews();
    proxy!publish(fut) } // sends future
}

class Proxy(Int limit,ServiceI s) implements ProxyI {
  ProxyI nextProxy;
  List[CustomerI] myCustomers:=empty; // fields
  ProxyI add(CustomerI cl) { ProxyI lastProxy=this;
    if length(myCustomers)<limit
      then myCustomers:=append(myCustomers,cl)
      else if nextProxy=null
        then nextProxy:= new Proxy(limit,s) fi;
      lastProxy:=nextProxy.add(cl) fi;
    return lastProxy }
  Void publish(Fut[News]fut){ News ns := get fut; // wait for the future
    myCustomers!signal(ns); // multi-cast the result
    if nextProxy=null then s!produce() else
      nextProxy!publish(fut) fi }
}

class Producer(NewsProducerI np) implements ProducerI {
  // Wrapper for NewsProducer
  News detectNews(){ News news; news:=np.getNews(); return news }
}

class NewsProducer()implements NewsProducerI{ List[News] nl;
  Void add(News ns){nl:=append(nl,ns)}
  News getNews(){News n;
    await nl /= empty; n:=first(nl); nl:=rest(nl); return n }
}

class Customer implements CustomerI{ // Consumer of news items:
  News news; // the latest news
  Void signal(News ns){ news:=ns }
}

```

Figure 7.6: Classes providing an implementation of the subscription example

```

class Proxy(Int limit,ServiceI s) implements ProxyI{
    ProxyI nextProxy; List[CustomerI] myCustomers:=empty;
    ProxyI add(CustomerI cl){ ... }
    Void publish(Fut[News]fut){
        myCustomers!signal(fut); // send future, no waiting
        if nextProxy=null then s!produce() else nextProxy!publish(fut) fi}
    }

class Customer implements CustomerI{ News news; ...
    Void signal(Fut[News] fut) { news:= get fut } // blocking wait
    }

```

Figure 7.7: DoS attack by a variation of the subscription example

We then say that the cycle is flooding. The flooding cycle identified above is harmless provided the customers are able to process their signal calls as fast as the cycle iterations. The programmer will be warned by our algorithm about each possible flooding, and should determine whether it is a real problem.

In contrast, the modified program version (Figure. 7.7) does not wait in *Proxy.publish* (doing *ns:=get fut*) until the newsletter is produced. Instead the future is directly passed to another asynchronous call (*myCustomers!signal(fut)*) in line 6 of Figure. 7.7 through the method *Proxy.publish*, this removes any progress dependency between the cycle producing the *Producer.detectNews* and *Customer.signal* calls and the processing of those calls. The completion of the *Producer.detectNews* and *Customer.signal* calls does not only depend on the speed of code execution, but depend on the rate of newsletter items arrivals. Practically, this flooding cycle generates a number of unprocessed calls that quickly grows to system limits.

Applying the algorithm to the example. Following [14], the *call* and *comps* sets for the two publish/subscribe versions are shown in Figures 7.8 and 7.9. Method names are abbreviated with two letters as indicated above, letting *Ng* abbreviate method *getNews* of interface *NewsProducerI*. There are two cycles in Figure. 7.8, i.e., cycle *A* and *B*. We have a flooding on the call to *Customer.signal* (*Cs*) in both cycles. However, this flooding does not reflect an actual flooding since the Customer objects

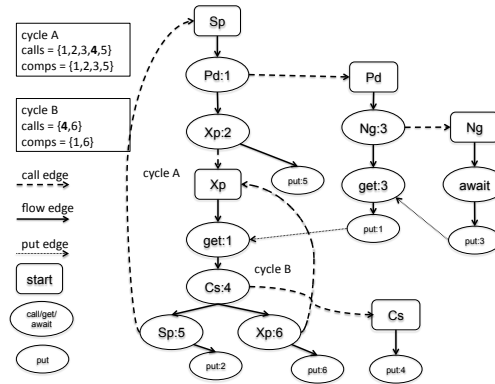


Figure 7.8: The graph and *call/comp* sets for the original version of the program (Figure. 7.6)

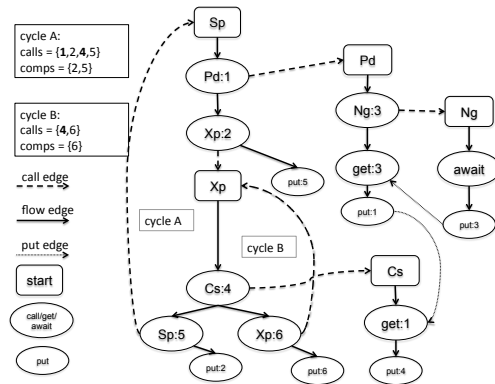


Figure 7.9: The graph and *call/comp* sets for the modified version of the program (Figure. 7.7)

easily keep up with the calls since the amount of work required by the Customer to complete a signal call is trivial. The execution rate is restricted with respect to the actual arrival of new items from the *NewsProducer* (by the blocking call in the proxies), and therefore, the rate of produced asynchronous calls to *Customer.signal* by this cycle is limited. Thus this is an example of *weak flooding* that is harmless.

```

class Attacker(ServerI s) {
  { this!run(); } // initialization
  Void run() { ClientI c := new Client(); c!connect(s);
    this!run() } // terminate and make recursive call
}
class Client() implements ClientI{
  Nat connect(ServerI s){
    Nat n := s.register();
    // blocking call, so a single client will not cause flooding
    return n }
}
class Server(DataBase db) implements ServerI{ {...} // Initialization
  Nat register(){Nat n :=0;
    if okcheck(caller) then Bool ok := db.open();
    if ok then n:=db.add(caller);
    db.query(...); db.close() fi fi; return n }
  // register requires time and resources
  ... }

```

Figure 7.10: Flooding by unbounded creation of innocent clients targeting the same server

Furthermore, cycle *B* is not infinite since it goes through the chain of Proxies. The modified version of the program is shown in Figure. 7.9. This version is displaying *strong flooding*. The flooding-cycle of *Pd* (Producer.detectNews) through both cycles is dangerous and will cause flooding of the system instantly. In version 1, there is a *get* in cycle *A* that regulates the speed of this cycle, whereas in the modified version there is no *get* in cycle *A*.

An example of instantiation flooding. The example in Figure. 7.10 shows how a *ClientDistribution* object can cause an attack by using an unbounded number of clients to flood the same server *s*, due to an unbounded recursion of the *run* method. The initialization of the attacker object of class *ClientDistribution* connects to a client, and the client do the registration of the server object. The attacker may start such a communication with lots of clients to register at the same server. (For simplicity, interfaces are omitted here.) Each client is innocent in the sense that it does not cause any attack by itself. By finding such a

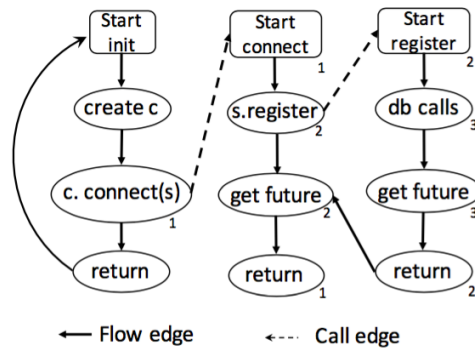


Figure 7.11: Static detection of flooding using unbounded creation

vulnerability in the `ClientDistribution`, an attacker can cause the flooding attack by calling `run()`. In addition, the non-blocking call in this method helps the attacker because the method does not wait for the `connect` calls to complete, therefore it is able to create more and more workload for the server s in almost no time. The execution of `f:=c!connect(s)` causes an asynchronous call and assigns a future to the call. Thus no waiting is involved. The `run` method recursively creates more and more objects, located somewhere in the distributed network. Therefore, the attacker creates flooding by rapidly creating clients that each performs a resource-demanding operation on the same server. Static analysis detects such attacks by finding a call loop (in this case inside `run`) which is also targeting the same server.

In this example, if the object creation in `run` had happened locally, an explicit instantiation flooding that consumes all the resources in an object will happen, which is a self DoS attack. However, since the object creation is distributed, the example in Figure. 7.10 shows an implicit attack because of targeting the same server by different clients.

Static analysis of the instantiation example. Consider the example in Figure. 7.10. For the `run` method of class `ClientDistribution`, the following cycle is detected:

```

the initialization of the attacker calls run
run creates a client object c
run calls c!connect(s)

```

run terminates and calls itself recursively in an asynchronous call.

The *run* call has a call edge to the flow graph of *connect* (call 1), and *connect* has a call edge to the flow graph of *register* (call 2). The call to *register* waits for completion of *register* since it is a blocking call, and the database calls (call 3) made by *register* wait for the completion of these database calls. The code for the database is not given, and therefore the analysis will be worst-case by considering the termination of such calls non-reachable (unless indirectly found strongly reachable). The control flow graph is given in Figure. 7.11. The set of weakly reachable call nodes of the cycle, i.e., *calls*, are $\{1, 2, 3\}$ with optimistic detection, and $\{1\}$ with pessimistic detection. And the set of strongly reachable calls, i.e., *comps*, is empty in both cases. This gives that the set of potentially flooding calls, given by *calls* – *comps*, is $\{1\}$ (c.connect) with pessimistic detection and $\{1, 2, 3\}$ with optimistic detection. However, in this case, call 1 does not reflect a real flooding since each call is on a separate object, but call 2 (s.register) and call 3 (the db calls) do. We detect strong flooding. The example may be improved by using suspending calls (using **await**) on the database operations.

7.6.1 Modification of the Static Detection to DDoS

As seen in Figure. 7.2, a DDoS attack on a server is often made through many innocent clients. This is hard to detect from the server side at runtime since each client may behave in an acceptable manner, and since the real attacker is hidden behind the clients. The original detection method [14] is not oriented towards such attacks, since it is not aware of the number of generated objects of a class. Moreover, the approach is using as an assumption that *an execution has a bounded number of objects*. Nevertheless, if applied to the example in Figure. 7.10, it will report a possible attack on the clients (treating all customers as one object), but not an attack on the Service server, which is the real attack. A draw-back is that there could be reported more false positives due to over-approximation.

A weakness with pessimistic detection is that the *connect* call, but not the *register* call, would be reported. Although the former call leads to the second, the detection result is not appropriate since a harmless

call is reported an not the harmful one. Another weakness is that the attack would not be discovered when removing the `connect` call from the attacker class and instead letting the `register` call be caused by the `init` method of class `Client`. (In this case the method parameter `s` should be transferred as a class parameter.) The reason for this is that indirect calls due to object generation are ignored (since by assumption there cannot be unboundedly many such calls). To compensate these weaknesses, we make two modifications wrt. [14], described below:

First, we modify the static analysis by viewing a **new** `C` statement as a special kind of a call statement with its own associated call number and a call edge to a copy of the `init` code of class `C`, which again may have further calls, treated as usual. More precisely, we treat **new** `C` as a *simple* call statement (like `new!C(classparameters)`) except that the *new* object is not known before the call) since the **new** statement does not wait for the `init` to complete. This allows us to see the generation of objects and to follow all implicit calls from the initialization code. Thus we can detect instantiation flooding attacks depending on call indirectly caused by object initialization. We may assume that an initialization cannot generate flooding in itself since each initialization is on a new object. Thus the call numbers associated with object creation can be included in *comps*. Furthermore, one more call on a new object cannot generate flooding on this object (unless in a cycle after the object creation). The same goes for a finite number of calls on a new object, if it can be detected statically that all these calls have the same new object as callee. These calls can also be included in *comps*. In the example of Figure. 7.10, we detect that the call `c!connect` is to the new `Client` object, and this call will then not be reported with the improved static detection.

Secondly, since implicit calls are important in DDoS attacks, we follow all call edges in the calculation of WR nodes, even in the pessimistic version. The resulting improved static detection method can then also detect the hidden attacker in all versions of the instantiation example, as shown below. Since the improved static detection method depends on static detection of same callee, we briefly discuss how to incorporate this: Two calls in the same method activation have the same callee if the callee is the same variable, and it is either

- a read-only variable (such as a parameter),

- a local variable and there are no updates on this variable between the calls, or
- a field variable and there are no updates on it nor suspension between the calls.

The first of these calls may be an object creation $x := \mathbf{new} C(\dots)$, and the second a call with x as callee (provided one of the conditions above are satisfied for x). This suffices for the example with the two calls $c := \mathbf{new} Client()$ and $f := c.connect(s)$. This detection could be improved in several ways. In particular, we may detect that the actual parameter s in the latter call refers to the same object for all activations of *run* since s is a read-only class parameter and the recursive *run* call is on the same object (since this is read-only). This could be used in the detection algorithm to see that all *s.register* calls refer to the same server s , which gives a clear indication of a DDoS attack.

Instantiation example revisited. We reconsider the example in Figure. 7.10, using the improved static detection algorithm. Now the *create c* node of Figure. 7.11 is represented as a call, say call 0. For the original version of the example, the initialization is empty so call 0 is considered terminating ($0 \in comps$). We get $calls = \{0, 1, 2, 3\}$ where 0 corresponds to the creation of the new C object. But call 0 ($c := \mathbf{new} Client$) and call 1 ($c!\mathbf{connect}$) do not generate flooding since they are on a new object. Thus $comps = \{0, 1\}$. This shows that there is a possibility of call flooding through call 2 ($s.register$) and call 3 (the db calls); and these correspond to actual attacks. For the modified version of Figure. 7.10, where the register call is caused by the Client initialization, we get a similar analysis except that there is no call 1 ($c!\mathbf{connect}$) since this is incorporated in the Client initialization. Thus we get that $calls = \{0, 2, 3\}$ and $comps = \{0\}$. Here the presence of call 0 enables us to detect call 2 in the (modified) initialization code and thereby also call 3 (and both correspond to possible flooding).

7.7 Conclusion

In this paper we have considered denial of service attacks, formulated in a high-level imperative language based on concurrent objects communicat-

ing by asynchronous calls and futures, thereby supporting asynchronous as well as synchronous communication. The language includes mechanisms for process control allowing non-trivial process synchronization by means of cooperative scheduling. We adapt a static detection algorithm developed for analysis of flooding to this setting, in order to detect possible denial of service attacks. This kind of static analysis is useful in the financial sector, because the aspect of trust between customers and service providers is essential, perhaps more so than in other application areas, and therefore static detection is valuable.

We have illustrated the approach on examples of distributed systems in the financial sector, including versions of a one-to-many attack and a many-to-one attack. In the first example a financial institution notifies a number of subscribing customers. We have seen that a revision of the basic notification software used by the financial institution, intended to be more efficient, actually implies a one-to-many attack on the subscribing customers. In this example, the financial institution was responsible for the attack, which could lead to loss of reputation and of customers. Static detection solved the situation here since the detection is made before the program is run. The underlying detection algorithm is sound for call-based coordinated attacks, provided the source code of the objects involved in the coordinated attack is available. In the many-to-one example, an attacker object causes an attack by using an unbounded number of clients, each innocent, to flood the same server s , letting a new client be created in each cycle.

In this paper we have adapted a general algorithm for detecting flooding [14] to the setting of DDoS and improved it to deal with unbounded object generation and to better reveal hidden attacks. Our framework can deal with advanced programming mechanisms including suspension and first-class futures considering distributed systems at a high-level of abstraction. It is therefore relevant for high-level modeling and prototyping of distributed software solutions. In future work, we suggest to complement the static checking with dynamic runtime checking since static detection methods give a degree of over-estimation. This could give a more precise combined detection strategy.

Acknowledgments. We thank the reviewers for significant feedback. This work is supported by the *IoTSec* project, the Norwegian Research

Council (No. 248113/O70), and by the *SCOTT* project, the European Leadership Joint Undertaking under EU H2020 (No. 737422).

Bibliography

- [1] Ashford, W. *DDoS is most common cyber attack on financial institutions*. <https://www.computerweekly.com/news/4500272230/DDoS-is-most-common-cyber-attack-on-financial-institutions/>. 2016.
- [2] Boer, F. D. et al. “A Survey of Active Object Languages”. In: *ACM Comput. Surv.* vol. 50, no. 5 (Oct. 2017), pp. 1–39.
- [3] Chang, R. et al. “Inputs of coma: Static detection of denial-of-service vulnerabilities”. In: *2009 22nd IEEE Computer Security Foundations Symposium*. IEEE. 2009, pp. 186–199.
- [4] Colón, M. and Sipma, H. “Synthesis of Linear Ranking Functions”. In: *Proceedings of the 7th International Conference on Tools and Algorithms for the Construction and Analysis of Systems*. TACAS 2001. London, UK, UK: Springer-Verlag, 2001, pp. 67–81.
- [5] Din, C. C. and Owe, O. “A sound and complete reasoning system for asynchronous communication with shared futures”. In: *J. Logical and Alg. Methods in Prog.* vol. 83, no. 5 (2014), pp. 360–383.
- [6] Douligeris, C. and Mitrokotsa, A. “DDoS Attacks and Defense Mechanisms: Classification and State-Of-The-Art”. In: *Computer Networks* vol. 44, no. 5 (2004), pp. 643–666.
- [7] Gulavani, B. S. and Gulwani, S. “A numerical abstract domain based on expression abstraction and max operator with application in timing analysis”. In: *International Conference on Computer Aided Verification*. Springer. 2008, pp. 370–384.
- [8] Jensen, M., Gruschka, N., and Herkenhöner, R. “A survey of attacks on web services”. In: *Computer Science-Research and Development* vol. 24, no. 4 (2009), p. 185.
- [9] Johnsen, E. B. and Owe, O. “An asynchronous communication model for distributed concurrent objects”. In: *Software & Systems Modeling* vol. 6, no. 1 (2007), pp. 39–58.

- [10] Johnsen, E. B. et al. “ABS: A Core Language for Abstract Behavioral Specification”. In: *Formal Methods for Components and Objects*. Ed. by Aichernig, B. K., Boer, F. S. de, and Bonsangue, M. M. Springer, 2012, pp. 142–164.
- [11] Karami, F., Owe, O., and Ramezanifarkhani, T. “An evaluation of interaction paradigms for active objects”. In: *J. Logical and Alg. Methods in Prog.* vol. 103 (2019), pp. 154–183.
- [12] Lambert, K. *Protecting Financial Institutions from DDoS Attacks*. <https://www.imperva.com/blog/protecting-financial-institutions-from-ddos-attacks>. 2018.
- [13] Meadows, C. “A formal framework and evaluation method for network denial of service”. In: *Computer Security Foundations Workshop. Proceedings of the 12th IEEE* (1999), pp. 4–13.
- [14] Owe, O. and McDowell, C. “On detecting over-eager concurrency in asynchronously communicating concurrent object systems”. In: *Journal of Logical and Algebraic Methods in Programming* vol. 90 (2017), pp. 158–175.
- [15] Qie, X. and Larry Peterson, R. P. and. “Defensive programming: Using an annotation toolkit to build DoS-resistant software”. In: *CM SIGOPS Operating Systems Review, 36(SI)* (2002), pp. 45–60.
- [16] Ramezanifarkhani, T., Fazeldehkordi, E., and Owe, O. “A Language-Based Approach to Prevent DDoS Attacks in Distributed Object Systems”. In: *29th Nordic Workshop on Programming Theory*. (extended abstract, 3 pages). Turku Centre for Computer Science, Nov. 2017.
- [17] Urrico, R. *Denial of Service Attacks Overwhelmingly Target Financial Services: Verisign*. <https://www.cutimes.com/2018/07/03/denial-of-service-attacks-overwhelmingly-target-fi/?sreturn=20190713065814/>. 2018.
- [18] Wilczek, M. *Why Banks Shouldn't Be in Denial About DDoS Attacks*. <https://www.globalbankingandfinance.com/why-banks-shouldnt-be-in-denial-about-ddos-attacks/>. 2018.

- [19] Zahoor, Z., Ud-din, M., and Sunami, K. “Challenges in Privacy and Security in Banking Sector and Related Countermeasures”. In: *International Journal of Computer Applications* vol. 144, no. 3 (2016), pp. 24–35.
- [20] Zargar, S. T., Joshi, J., and Tipper, D. “A survey of defense mechanisms against distributed denial of service (DDoS) flooding attacks”. In: *IEEE Comm. Surveys Tutorials* vol. 15, no. 4 (2013), pp. 2046–2069.
- [21] Zheng, L. and Myers, A. C. “End-to-end availability policies and noninterference”. In: *Computer Security Foundations, 2005. CSFW-18 2005. 18th IEEE Workshop* (2005), pp. 272–286.

