Antti Hakkala

On Security and Privacy for Networked Information Society

Observations and Solutions for Security Engineering and Trust Building in Advanced Societal Processes

Turku Centre *for* Computer Science

TUCS Dissertations
No 225, November 2017

ON SECURITY AND PRIVACY FOR NETWORKED INFORMATION SOCIETY

Observations and Solutions for Security Engineering and Trust Building in Advanced Societal Processes

ANTTI HAKKALA

To be presented, with the permission of the Faculty of Mathematics and Natural Sciences of the University of Turku, for public criticism in Auditorium XXII on November 18th, 2017, at 12 noon.

University of Turku
Department of Future Technologies
FI-20014 Turun yliopisto

2017

SUPERVISORS

Adjunct professor *Seppo Virtanen*, D. Sc. (Tech.)
Department of Future Technologies
University of Turku
Turku, Finland

Professor *Jouni Isoaho*, D. Sc. (Tech.)
Department of Future Technologies
University of Turku
Turku, Finland

REVIEWERS

Professor *Tuomas Aura*
Department of Computer Science
Aalto University
Espoo, Finland

Professor *Olaf Maennel*
Department of Computer Science
Tallinn University of Technology
Tallinn, Estonia

OPPONENT

Professor *Jarno Limnéll*
Department of Communications and Networking
Aalto University
Espoo, Finland

ISBN 978-952-12-3607-5 (Online)
ISSN 1239-1883

To my wife Maria,
I am forever grateful for everything.

Thank you.

ABSTRACT

Our society has developed into a networked information society, in which all aspects of human life are interconnected via the Internet — the backbone through which a significant part of communications traffic is routed. This makes the Internet arguably the most important piece of critical infrastructure in the world. Securing Internet communications for everyone using it is extremely important, as the continuing growth of the networked information society relies upon fast, reliable and secure communications.

A prominent threat to the security and privacy of Internet users is mass surveillance of Internet communications. The methods and tools used to implement mass surveillance capabilities on the Internet pose a danger to the security of all communications, not just the intended targets. When we continue to further build the networked information upon the unreliable foundation of the Internet we encounter increasingly complex problems, which are the main focus of this dissertation. As the reliance on communication technology grows in a society, so does the importance of information security. At this stage, information security issues become separated from the purely technological domain and begin to affect everyone in society. The approach taken in this thesis is therefore both technical and socio-technical.

The research presented in this PhD thesis builds security in to the networked information society and provides parameters for further development of a safe and secure networked information society. This is achieved by proposing improvements on a multitude of layers. In the technical domain we present an efficient design flow for secure embedded devices that use cryptographic primitives in a resource-constrained environment, examine and analyze threats to biometric passport and electronic voting systems, observe techniques used to conduct mass Internet surveillance, and analyze the security of Finnish web user passwords. In the socio-technical domain we examine surveillance and how it affects the citizens of a networked information society, study methods for delivering efficient security education,

examine what is essential security knowledge for citizens, advocate mastery over surveillance data by the targeted citizens in the networked information society, and examine the concept of forced trust that permeates all topics examined in this work.

KEYWORDS: Information society, security, privacy, trust, surveillance, information security education, biometrics, Internet voting, cryptography

TIIVISTELMÄ

Yhteiskunta, jossa elämme, on muovautunut teknologian kehityksen myötä todelliseksi tietoyhteiskunnaksi. Monet verkottuneen tietoyhteiskunnan osa-alueet ovat kokeneet muutoksen tämän kehityksen seurauksena. Tämän muutoksen keskiössä on Internet: maailmanlaajuinen tietoverkko, joka mahdollistaa verkottuneiden laitteiden keskenäisen viestinnän ennennäkemättömässä mittakaavassa. Internet on muovautunut ehkä keskeisimmäksi osaksi globaalia viestintäinfrastruktuuria, ja siksi myös globaalin viestinnän turvaaminen korostuu tulevaisuudessa yhä enemmän. Verkottuneen tietoyhteiskunnan kasvu ja kehitys edellyttävät vakaan, turvallisen ja nopean viestintäjärjestelmän olemassaoloa.

Laajamittainen tietoverkkojen joukkovalvonta muodostaa merkittävän uhan tämän järjestelmän vakaudelle ja turvallisuudelle. Verkkovalvonnan toteuttamiseen käytetyt menetelmät ja työkalut eivät vain anna mahdollisuutta tarkastella valvonnan kohteena olevaa viestiliikennettä, vaan myös vaarantavat kaiken Internet-liikenteen ja siitä riippuvaisen toiminnan turvallisuuden. Kun verkottunutta tietoyhteiskuntaa rakennetaan tämän kaltaisia valuvikoja ja haavoittuvuuksia sisältävän järjestelmän varaan, keskeinen uhkatekijä on, että yhteiskunnan ydintoiminnot ovat alttiina ulkopuoliselle vaikuttamiselle. Näiden uhkatekijöiden ja niiden taustalla vaikuttavien mekanismien tarkastelu on tämän väitöskirjatyön keskiössä. Koska työssä on teknisen sisällön lisäksi vahva yhteiskunnallinen elementti, tarkastellaan tiukan teknisen tarkastelun sijaan aihepiirä laajemmin myös yhteiskunnallisesta näkökulmasta.

Tässä väitöskirjassa pyritään rakentamaan kokonaiskuvaa verkottuneen tietoyhteiskunnan turvallisuuteen, toimintaan ja vakauteen vaikuttavista tekijöistä, sekä tuomaan esiin uusia ratkaisuja ja avauksia eri näkökulmista. Työn tavoitteena on osaltaan mahdollistaa entistä turvallisemman verkottuneen tietoyhteiskunnan rakentaminen tulevaisuudessa. Teknisestä näkökulmasta työssä esitetään suunnitteluvuo kryptografisia primitiivejä tehokkaasti hyödyntäville rajallisen laskentatehon sulautetu-

ille järjestelmille, analysoidaan biometrisiin passeihin, kansainväliseen passijärjestelmään, sekä sähköiseen äänestykseen kohdistuvia uhkia, tarkastellaan joukkovalvontaan käytettyjen tekniikoiden toimintaperiaatteita ja niiden aiheuttamia uhkia, sekä tutkitaan suomalaisten Internet-käyttäjien salasanatottumuksia verkkosovelluksissa.

Teknis-yhteiskunnallisesta näkökulmasta työssä tarkastellaan valvonnan teoriaa ja perehdytään siihen, miten valvonta vaikuttaa verkottuneen tietoyhteiskunnan kansalaisiin. Lisäksi kehitetään menetelmiä parempaan tietoturvaopetukseen kaikilla koulutusasteilla, määritellään keskeiset tietoturvatietouden käsitteet, tarkastellaan mahdollisuutta soveltaa tiedon herruuden periaatetta verkottuneen tietoyhteiskunnan kansalaisistaan keräämän tiedon hallintaan ja käyttöön, sekä tutkitaan luottamuksen merkitystä yhteiskunnan ydintoimintojen turvallisuudelle ja toiminnalle, keskittyen erityisesti pakotetun luottamuksen vaikutuksiin.

AVAINSANAT: Tietoyhteiskunta, turvallisuus, yksityisyys, luottamus, valvonta, tietoturvakoulutus, biometriikat, Internet-äänestys, kryptografia

PUBLICATIONS

Certain key aspects, ideas and figures in this thesis have previously appeared – in different phases of maturity – in the following publications, in chronological order:

Hakkala, Antti and Virtanen, Seppo. Accelerating Cryptographic Protocols: A Review of Theory and Technologies. In *Proceedings of The Fourth International Conference on Communication Theory, Reliability, and Quality of Service CTRQ 2011, NexComm 2011 conference*, pages 103 – 109, Budapest, Hungary, April 17.-22.2011.

Heimo, Olli I. and Hakkala, Antti and Kimppa, Kai K. The problems with security and privacy in eGovernment — Case: Biometric Passports in Finland. In *Ethicomp 2011 Conference Proceedings*, Sheffield-Hallam University, UK, September 14.-16.2011.

Heimo, Olli I. and Hakkala, Antti and Kimppa, Kai K. How to abuse biometric recognition systems. *Journal of Information, Communication & Ethics in Society*, Vol. 10 Iss: 2 pp. 68 – 81, 2012.

Hakkala, Antti and Virtanen, Seppo. University-Industry Collaboration in Network Security Education for Engineering Students. In *International Conference on Engineering Education ICEE 2012*, Turku, Finland, July 2012.

Moosavi, Sanaz Rahimi and Hakkala, Antti and Isoaho, Johanna and Virtanen, Seppo and Isoaho, Jouni. Specification analysis for secure RFID implants. *International Journal of Computer Theory and Engineering* 04/2014; 6(2), 2014.

Nigussie, Ethiopia and Hakkala, Antti and Virtanen, Seppo and Isoaho, Jouni. Energy-aware Adaptive Security Management for Wireless Sensor Networks. In *IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks*, Sydney, Australia, July 2014.

Hakkala, Antti and Virtanen, Seppo. Virtualization of laboratory education in network security engineering. In proceedings of *International Conference on Engineering Education ICEE 2015*, Zagreb, Croatia, 20.-24.7.2015.

Hakkala, Antti and Isoaho, Jouni. Defining and Measuring Key Expertise Areas in Information Security for Higher Education Students. In proceedings of *International Conference on Engineering Education ICEE 2015*, Zagreb, Croatia, 20.-24.7.2015.

Hakkala, Antti and Isoaho, Jouni and Virtanen, Seppo. Towards Adaptive Cryptography and Security with Software Defined Platforms. In Waqar Hussain, Jari Nurmi, Jouni Isoaho, and Fabio Garzia (eds.), *Computing Platforms for Software-Defined Radio*, Springer, 2017.

Hakkala, Antti and Heimo, Olli I. and Hyrynsalmi, Sami and Kimppa, Kai K. Security, Privacy'); DROP TABLE users; – and Forced Trust in the Information Age? In *Ethicomp 2017 Conference*, Turin, Italy, 5.-8.6.2017.

While not directly in the central theme of this dissertation, the following publications have had an influence on this research:

Hakkala, Antti and Virtanen, Seppo. Modern Cryptography Algorithms in the Embedded Environment. Presentation, FRUCT conference, St. Petersburg, Russia, April 2010.

Hakkala, Antti. Fractional Biometrics and the Security of Biometric Identifiers in Remote Identification. Poster paper, Nordsec 2010 conference, Espoo, Finland, October, 2010.

Nigussie, Ethiopia and Guang, Liang and Boyko, Alexey and Hakkala, Antti and Sainio, Petri and Virtanen, Seppo and Isoaho, Jouni. Incubator Platform for Multidisciplinary Innovation in Research and Education. *International Journal of Knowledge Society Research*, 3(3), 29-44, July-September 2012.

ACKNOWLEDGMENTS

Back in 2009, just before I was starting my doctoral studies, I had only a vague idea of what lies ahead for me. I had been offered an opportunity to do research for a living and while I had competing offers from the industry, I chose the perhaps more uncertain path of academia. One of my main motivations was to find out whether I had what it takes to get a doctoral degree. Now, as I am writing this almost exactly 8 years later, I can say that those 8 years have been among the best in my life. I have had the privilege to work and collaborate with some amazing people, and thus I am indebted to many for their help, support, and encouragement.

I want to start by thanking my supervisors, adjunct professor SEPPO VIRTANEN and professor JOUNI ISOAHO, who have guided me from the beginning and given me a lot of support, patience and advice.

I also want to thank the reviewers of my thesis, professor TUOMAS AURA and professor OLAF MAENNEL, for finding the time to read my book and for the insightful comments and helpful suggestions. I also would like to warmly thank professor JARNO LIMNÉLL for being my opponent in my public defense.

I want to thank KAI K. KIMPPA for all the discussions and help over the years; OLLI I. HEIMO for being a good friend and colleague (has it been 15 years already?); SAMI HYRYNSALMI for all the bad jokes and good coffee (or was it the other way around?); KATRI HAVERINEN for all the good ideas, insightful discussions that sometimes had something to do with work, and getting me to take regular breaks; and ETHIOPIA NIGUSSIE for good discussions and bearing with my music choices at the office. I am indebted to you all.

Thank you to all my past and present colleagues at the Communication Systems lab at University of Turku, Department of Future Technologies (previously Department of Information Technology): PETRI SAINIO, ALI FAROOQ, NANDA KUMAR THANIGAIVELAN, SANAZ RAHIMI MOOSAVI, JOHANNA ISOAHO, and ALEXEY BOYKO.

A huge thank you to the great colleagues at the department who have helped me with all possible (and impossible) things over the years: ERKKI KAILA, ROLF LINDÉN, JOHANNES HOLVITIE, ANTTI AIROLA, SAMI NUUTTILA, MARKO LAHTI, PETER LARSSON, ERNO LOKKILA, EINARI KURVINEN, SAMPSA RAUTI, and all the rest. You are awesome.

A warm thank you to all past and present members of SOHON TORWET, for all the wonderful gigs, rehearsals and parties over the years. It has been fun, and I'll try to pop in to play a gig or two now and then. Thank you to all the wonderful students I have had the privilege to work with during my years at the department, and to everyone else I have forgotten to mention here.

Finally, I would like to thank my family; my wife MARIA and son ALEKSI, you guys have made all this worth it; and my parents, for everything.

In Turku
October 20th 2017

Antti Hakkala

CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACRONYMS

AA Active Authentication

AES Advanced Encryption Standard

BAC Basic Access Control

BGP Border Gateway Protocol

BYOD Bring Your Own Device

BYOT Bring Your Own Technology

CCTV Closed-Circuit Television

CGIS Critical Governmental Information System

CIA Central Intelligence Agency

CVE Common Vulnerabilities and Exposures

DDoS Distributed Denial of Service

DES Data Encryption Standard

DNS Domain Name System

DPI Deep Packet Inspection

DRE Direct Recording Electric

DSP Digital Signal Processing

DVB Digital Video Broadcasting

EAC Extended Access Control

ECC Elliptic Curve Cryptography

ECHR European Court of Human Rights

FISA Foreign Intelligence Surveillance Act

FISC United States Foreign Intelligence Surveillance Court

FPGA Field Programmable Gate Array

FRA Försvarets Radioanstalt

FU Functional Unit

GCHQ Government Communications Headquarters

GDR German Democratic Republic

HFST Helsinki Finite-State Transducer Technology

ICAO International Civil Aviation Organization

IRC Internet Relay Chat

IETF Internet Engineering Task Force

IoT Internet of Things

IPS Intrusion Prevention System

IRIS Iris Recognition Immigration System

ISE Instruction Set Extension

ISP Internet Service Provider

LDS Logical Data Structure

MRZ Machine-Readable Zone

NIST National Institute of Standards and Technology

NLP Natural Language Processing

NSA National Security Agency

OMorFi Open Source Finnish Morphology

OWA Open Wireless Architecture

PA Passive Authentication

PACE Password Authenticated Connection Establishment

PC Personal Computer

POODLE Padding Oracle On Downgraded Legacy Encryption

RDP Remote Desktop Protocol

RFID Radio Frequency Identification

RNG Random Number Generator

SCOTUS Supreme Court of the United States

SDR Software Defined Radio

SDSC Software-Defined Secure Communications

SIGINT Signals Intelligence

SM Secure Messaging

SOD Security Object

SSH Secure Shell

SSL Secure Socket Layer

SSO Single Sign-On

Stasi Ministerium für Staatssicherheit

SuPo Suojelupoliisi

TACO Tools for Application-specific hardware/software CO-design

TLS Transport Layer Security

TTA Transport Triggered Architecture

UI User Interface

VPN Virtual Private Network

WEP Wired Equivalent Privacy

WLAN Wireless Local Area Network

WPA Wifi Protected Access

WSN Wireless Sensor Network

WWW World Wide Web

1

INTRODUCTION

Look at you, hacker. A pathetic creature of meat and bone.
Panting and sweating as you run through my corridors.
How can you challenge a perfect immortal machine?

– SHODAN
System Shock (1994)

In the networked information society, data has become one of the most valuable assets. People are connected through communication networks and massive amounts of information is stored in distributed data storage, interconnected and accessible from anywhere in the world. This presents unique opportunities for scientific and economic development. With the advent of big data sciences, our capacity to process data and extract information from data amounts that have been beyond our means in the past is now within the realm of possibility. This in turn presents opportunities for the increase of wellbeing in all layers of society.

As the value of information – and especially well-ordered information – has skyrocketed, security and privacy have simultaneously become important assets in the networked information society. When most information systems that store sensitive data are interconnected, it should also follow that the security and privacy of those systems, and the interconnection network should also be beyond reproach. This is unfortunately not the case in our current society, and this observation is one of the key motivations behind this thesis.

The Internet is at its heart an insecure communications channel. This stems from its distributed nature – no single entity has complete control – and neutrality towards traffic: all packets are at least in principle treated equally, and traffic is routed based on efficiency and politics, not security and privacy. Regardless

of inherent security issues, we as a society have managed to become all but dependent on the Internet. This dependence can be leveraged for various purposes, ranging from the benevolent – such as providing better services to customers – to the nefarious – exploiting vulnerabilities and malicious data collection.

Another problem is that with the advance of data analysis techniques, it is no longer possible for an individual to hide in the masses of people. Modern data mining techniques can extract interesting patterns from datasets previously considered prohibitively large in the past, and this data can in turn be connected to an individual person. By nature, the information collected by governments on their citizens is generally accurate, strongly identifying, and potentially very sensitive from the point of view of privacy. This creates controversy on many levels, including trust in government systems, and the potential use scope of governmental information systems.

In this thesis these problems of the networked information society are mitigated through improving design tools for secure embedded systems, developing better methods and tools for security teaching in higher education, researching password authentication systems, and providing a deeper understanding of the relationship between trust, surveillance, security, and privacy in the networked information society.

1.1 RESEARCH PROBLEMS

In this thesis we identify, examine, explore, and provide *research-based solutions and suggestions for the following problems related to security and privacy of the citizens of the networked information society.*

From the point of view of privacy, systematic information gathering and surveillance are inherently problematic. We study various aspects of privacy and surveillance, especially focusing on topics and themes closely related to Internet surveillance. Surveillance and information gathering are valuable and widely used tools for a society to provide better security and services to its citizens. We postulate that *systematic information gathering and surveillance in today's networked information society limit the willingness for societal participation.* This effectively reduces the potential benefits of information gathering and surveillance similarly as

has been observed in historical examples of pre-digital age surveillance societies. From the point of view of the infrastructure of the networked information society, *the technological methods used to conduct systematic information gathering and surveillance endanger the security and safety of all Internet users.* By pursuing strategic advantage in information warfare, governments and influential organizations use tactics that highlight and even introduce weaknesses in communication infrastructure that is vital to the networked information society.

From the perspective of a citizen of the networked information society, there is very little privacy and security built in to the system. Instead, upon closer inspection the networked information society appears to rely on societal trust, just as previous societies have done as well. We observe the phenomenon of *forced trust* in this context, and postulate that the citizens of the networked information society are in an unsustainable position. On a societal level, people are *forced to use and trust their personal information in the hands of a system they have very little control over.* When this forced trust is combined with security issues with the insecure underlying infrastructure of the networked information society these systems are based on, the users have a limited set of behavioral responses, as the damage potential to privacy and data security of citizens is significant. One of the few avenues of influence citizens of a society have on such matters of policy is voting in elections to choose new politicians. Ironically, online voting systems are precisely the kind of systems that are at the heart of the issues arising from forced trust. These problems are examined later in this thesis.

From the perspective of the networked information society, *key societal processes are negatively affected by data collection, surveillance and reduced infrastructure security.* In this thesis we examine two prominent processes: border control and societal decision making. Biometric passport systems leverage biometric identifiers for secure machine-readable international travel documents (i.e. passports and similar documents), at least at face value providing more security. At the same time *biometric passports introduce new problems and vulnerabilities related to issues with biometric information.* Combining these with the previously observed problems of forced trust on information systems, insecure infrastructure, and surveillance further highlights these issues.

Decision making on a societal level requires the arrangement of elections. Various electronic voting systems and Internet-based voting protocols have been extensively researched in the past. These systems have been found to have serious inherent vulnerabilities and shortcomings that make them ill-suited for actual large-scale voting. Regardless of these identified vulnerabilities, the pressure to implement electronic and Internet voting protocols on a societal level is intense. When the previously observed problems of forced trust on information systems, insecure infrastructure, and surveillance are combined with Internet voting, new threat scenarios involving infrastructure attacks against Internet voting systems emerge. Based on what we know of infrastructural security of the networked information society, these potential threats are credible and in the worst case potentially destabilizing to the whole society. Existing frameworks and taxonomies for Internet voting do not allow for an adversary with the capabilities that actors capable of mass network surveillance possess. Ultimately in the networked information society, in elections with major significance *it is reasonable to assume the existence of an adversary with the short-term capability to compromise the security offered by cryptography.*

User authentication is vital to operational security of information systems. The most common user authentication method – the password – has been previously found to be insecure in many use cases. Natural language passwords have been suggested in order to make password authentication more resilient. *Natural languages with sufficient linguistic complexity could provide increased security against attackers.* We examine web user passwords with Natural Language Processing (NLP) methods for natural language patters, and explore password behavior in students.

Information security, and by extension cyber security, does not rely only on technological solutions to engineering problems. The citizens of a networked information society are in a key position to adapt good security practices into their behavior. For this to happen, people need to be educated on security issues. *Proper teaching methods for security education and also definition of essential security knowledge for different demographics are needed.*

The networked information society has brought ubiquitous computing to the forefront with Internet of Things (IoT). An

increasing number of devices connect to the Internet, many of them with restricted computational capacity. In order for these devices to comply with the security requirements of the modern Internet, they must still be able to perform cryptographic operations efficiently. *Potential weaknesses — whether intentional or otherwise — in established cryptographic protocols threaten the security and privacy of Internet communications, and for devices with limited computational capability this can be a serious threat.*

When all these observations are combined, from the result we can postulate that *the current technology and infrastructure of the networked information society cannot be trusted to preserve the security and privacy of all its users in its current form.*

1.1.1 *Research problems and perspective*

The research problems are approached from different perspectives in this thesis. These are illustrated in Figure 1. From the figure we can observe that the human being has clearly been placed at the heart of the research problems. Indeed, a key observation in this thesis is that human beings are central to the issues we have observed in the networked information society due to human-derived data being a common theme in all observed research problems. *Potential solutions should be considered from both technical and socio-technical perspectives.* Technology-centered solutions will not be able to address all issues with problems that have socio-technological dimensions, some of which could be characterized as *wicked problems* [1]; there are many conflicting interests and definitions, and potential "brute force" solutions may be even worse than the original problem. Similarly, it is difficult to use "soft" measures to solve something that is at its heart an engineering problem, solvable with technology instead of legislation.

1.2 OBJECTIVES

Based on the identified research problems that were outlined in the previous section, the following research objectives are pursued in this thesis.

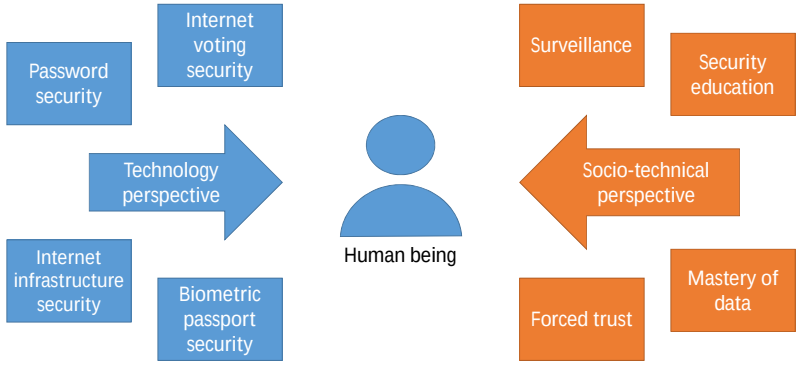

Figure 1: Themes discussed in this thesis and their relation to technological and socio-technological perspectives.

RO1: BUILDING SECURITY IN TO THE NETWORKED INFORMATION SOCIETY. Our basic network infrastructure – and thus the infrastructure of the networked information society – has been built on unsustainable assumptions regarding privacy and security. These assumptions are highlighted through identifying forced trust relationships at the infrastructure level, analyzing natural language use patterns of web user passwords, and examining central information society processes of biometric border control and voting over the Internet through these highlighted issues. Security is built in the information society by describing new security design and optimization methods for embedded systems, better security education, and providing a framework for data ownership of person-derived data.

RO2: PROVIDING DIRECTIONS FOR FURTHER DEVELOPMENT OF A SAFE AND SECURE NETWORKED INFORMATION SOCIETY. We must assume a holistic view of the problem with simultaneously providing privacy, security and meaningful and ethically sustainable services to the citizens of the networked information society. The objective is to present a convincing argument for surveillance having a detrimental effect on the networked information society. The objective is to suggest large-scale approaches that can alleviate the con-

> sequences of built-in surveillance. While it has legitimate uses, the present way surveillance is being used in our society is counterproductive to societal progress, and a clear and present threat to security and privacy of Internet users.

We adopt a data-driven approach in this study. Data is at the heart of the described research problems, and by observing them from the point of view of data gives a clear view of the challenges posed to the networked information society. A central theme is human-derived data, and whether the origin of data has any significance.

The research objectives can be clearly partitioned into technical and socio-technical domains. Technology does not exist in a vacuum with regard to its surrounding society, and therefore any research into the problems and opportunities generated by technology should also at least attempt to take both aspects into consideration. RO1 is clearly approached from the technological point of view, while RO2 has embedded social dimensions that bring additional complexity to the objective.

1.2.1 *Scope of the thesis*

Given the expanse of topics appearing in the defined research problems and the associated research objectives, it is clear that the scope of this thesis must be limited. While we aim to provide a holistic view of the presented problems, we recognize that it is impossible and that a tighter scope must be defined. Surveillance is observed both as a concept and as an activity that is undertaken using technical tools and methods, but the scope is limited to how surveillance relates to the abuse of forced trust on network infrastructure. Any further discussion on the sociological dimensions of surveillance is outside the scope of this thesis.

The scope on human mobility is limited to the use of biometric passports in border control, and how the use of biometrics in such a vital societal process causes problems and potential threats to security and privacy of citizens in the networked information society. Internet voting systems are limited in scope to adversary models in existing frameworks and taxonomies. Other

issues with Internet voting systems, and electronic voting in general, are outside the scope of this thesis.

1.3 INFORMATION SOCIETY

As this concept is central to the theme of this thesis, the first definition that we establish is for *information society*. Dictionary definitions from various sources differ from one another quite significantly. BusinessDictionary.com gives the following definition:

> [a] post-industrial society in which information technology (IT) is transforming every aspect of cultural, political, and social life and which is based on the production and distribution of information.[1]

Another definition for information society is given by Scott and Marshall [2]:

> The information society is one in which information is the defining feature, unlike the industrial society where steam power and fossil fuels were distinguishing elements.

These give a general but vague impression on the nature of information society. Indeed, Scott and Marshall observe that the term is widely used, but imprecise [2]. But if we aim to use it as context for observing societal processes, we must use a more robust and precise definition. Moore [3, pp. 271–272] identifies the three following key characteristics of an information society.

1. Information is used as an economic resource
2. It is possible to identify greater use of information amongst the general public
3. There exists a developed information sector within the economy

1 BusinessDictionary.com http://www.businessdictionary.com/definition/information-society.html

He goes further to note that even though we can identify these characteristics, providing a definition in quantitative terms for an information society seems to be exceedingly difficult. He approaches the problem from an economy viewpoint by observing that it is possible to define the information sector of the economy, and proceeds further in this direction. We cannot define the information society based solely on economic definitions, however.

Webster [4, p. 8] points out that many writers who write about the information society do so without any clear definition or understanding of what information society is, and consequently fail to provide new views on the matter. Webster himself identifies the following five definitions of an information society: *technological, economic, occupational, spatial, and cultural*. All of these can be used as lenses of observation to identify new aspects, and they are not mutually exclusive. He also proceeds to describe a sixth definition, *transformation*, where information or theoretical knowledge is at the core of our conduct and that information transforms everything we do.

An information society is clearly something that transforms all aspects associated with it. We are witnessing and examining the transformation of society's core processes; how we are being watched and monitored, cross national borders, and make decisions in society. We must – if it is within our power – steer this transformation towards an ethically sustainable direction, and at the same time do not have to make compromises on security and privacy of people, organizations and nations. This is a certainly a worthy goal, even if it sometimes is on the idealistic side of things.

Whenever we discuss technology and its role in the wider society, there are also certain possible criticisms to combining these two seemingly unrelated domains. It can be argued that technology should kept in the technological domain, without trying to "mud the waters" with questions of ethics and societal interactions. On this, the author strongly disagrees. Technology can be separated from the context in which it is used, but this is not recommendable if we wish to gain insight into more complex questions that appear when we consider ethical and social ramifications for use and development of new technology. To put it another way, technology does not exist in a vacuum. How it develops shapes not only technological processes, but also ev-

erything that comes into contact with it. Technology is transformative by nature, one needs only to look at how the automobile changed the whole society and shaped many aspects of the future of western civilization. Technocracy leads to looking at problems and only seeing technological or engineering solutions, while a true solution would require a more holistic approach – something that is aspired to within this thesis.

1.3.1 *Networked information society*

We currently are living in a *networked information society*, where the above observations of technology transformation and information being a central resource have realized in multiple aspects of daily life. As is implied by the name, the networked information society is interconnected, with practically unlimited connectivity between citizens, organizations, and nations. This interconnectivity is implemented with the global information network, i. e. the Internet.

1.4 CLOUD COMPUTING AND BIG DATA

Most devices that connect to the Internet have been "traditional" computers and equivalent mobile devices such as smartphones and tablets, but this is changing rapidly. The advent of IoT has brought increasing numbers of connected devices and sensor nodes to the Internet, and the growth is projected to increase in the future. Estimations on the number of IoT devices in 2020 range from 20 million.[2] to 200 million[3]

Handheld platforms have become significantly more popular in the last years. Between 2011 and 2015 in the United States, the percentage of adults that own a smartphone has grown from 35% to 68%, and tablets have seen significant growth in market share in the same time period, from 4% to 45%. The percentage of adults that own traditional desktop or laptop computers

2 Gartner http://www.gartner.com/newsroom/id/3165317

3 Intel http://www.intel.com/content/www/us/en/internet-of-things/infographics/guide-to-iot.html

has stayed at around 70%. [5] More recent reports[4] show market penetration for smart phones to be nearly 90% in the US. We now possess more computational power in a smaller size than ever before, but still an increasing amount of the actual computational work is done on remote platforms that scale well to massive tasks, i.e. the cloud.

Moore's Law has yet to fail us.

Cloud computing has become commonplace on most platforms as a method for transferring load to another platform and accessing those remote resources via a network connection. More often than not, with cloud storage application integration to phones and tablets, user data is nowadays stored on a device that the users do not personally control. The concept of personal computing devices has been transformed by this rise of cloud computing. A vast number of services and software ecosystems would not even exist without cloud computing.

Two of the most common examples of cloud computing are cloud storage and virtualized platforms for general purpose computing. Both have become nearly ubiquitous in the 2010's. Storage services such as Dropbox,[5] Microsoft OneDrive[6], and Google Drive[7] provide users with data storage space and the capability to access that data via the Internet. Virtual server hosting platforms such as Amazon Elastic Compute Cloud (EC2)[8] and Microsoft Azure[9] provide cloud virtual platforms as a service.

1.4.1 *Security concerns for Internet infrastructure*

The concept of *big data* has been a significant driver for advancing cloud services and data analysis tools and techniques, forming a self-enhancing loop of research, innovation and progress. The expansion of *non-ownership*, storing even personal and private data into systems not controlled by owners of the data in question quickly becomes a problem, however, when we consider the situation from the point of view of an adversary want-

4 Nielsen. "Millennials are top smartphone users" Online, availabe at http://www.nielsen.com/us/en/insights/news/2016/millennials-are-top-smartphone-users.html Accessed 24.9.2017.

5 https://www.dropbox.com/

6 https://onedrive.live.com/

7 https://www.google.com/drive/

8 https://aws.amazon.com/ec2/

9 https://azure.microsoft.com

ing to obtain as much of people's personal data as possible. This development makes it significantly easier to gain access to increasingly large sets of data stored in or in transit through a single place.

Security is one of the key success factor of cloud computing platforms. Cloud data storage and processing solutions consolidate a lot of data into a relatively small space with a cohesive attack surface, putting the proverbial eggs in one basket. Users need to trust the service providers that they are handling their data in a secure manner, and that their privacy is upheld. People do not in general appreciate their personal data being accessed without their permission, so the user must be able to trust the cloud provider that they will treat user data accordingly.

Web services in general know a lot about us; For example, both Google and Facebook have a comprehensive picture on who we are, what we do online, who we associate with, and what our interests, passions and even secrets are. Being able to consolidate personal data from multiple sources makes this possible. This access to not only user provided data, but also data combined form all other sources further makes certain online companies lucrative targets for those trying to access confidential data on people.

An important factor in the overall security of cloud services is the underlying infrastructure of the Internet. All organizations and companies that rely on cloud services also rely on the Internet working as intended. This assumption can be exploited for malicious purposes by a sufficiently sophisticated attacker. We shall examine the ways this can be done later in this thesis.

1.4.2 *Governmental information systems*

When the information people voluntarily give for others to store in cloud services is further combined with data stored in government information systems, it opens even more possibilities for malicious activities. A prime example is surveillance. Heimo, Koskinen and Kimppa [6] define a Critical Governmental Information System (CGIS) as *"an information system developed for governmental needs including data or functionality which is critical in nature to the security or wellbeing of individuals or the society as a whole."* Such systems are by definition of critical importance to

the information society in general. Even though it can be argued that the government does not need to involve itself in the critical information systems of the society and let private entities control everything, the opinion of the author is that this is not a very feasible approach to take in a real-world information society.

Connecting CGISs to the Internet expose them to attacks by malicious entities, but also facilitates efficient combining of information stored on CGISs with other data sources. Cross-referencing data about people between different data sets yields significantly more information about the subjects. When considered from surveillance and privacy standpoint, it is clearly a problematic practice [7, p. 40–42].

The two main information society processes examined in this thesis — biometric passport systems and Internet voting systems — are both definitely CGISs; the information managed in these processes is critical to security and privacy of those involved in the process, and thus these processes and their potential implementations must be tightly scrutinized for any security issues.

1.5 ONLINE SURVEILLANCE

The modern world relies heavily upon secure communication of information. All areas of industry and society encounter situations that require security and privacy. For organizations in specific fields with tight security requirements, confidentiality, integrity and authenticity of data are non-negotiable conditions for success. Industries such as finance and health care have a considerable impact on our society, even further underlining strict security requirements. And as people increasingly use the Internet for both personal and private matters, the same requirements of confidentiality and integrity of communications also apply here. To put it concisely, security and privacy are extremely important concepts to people in both personal and business settings, and the Internet has a key role in how the security and privacy are actually realized for users.

Concerns about this progress are often casually dismissed with a quip such as "who else than me cares about my data anyway?", but in reality there are several actors that are very keenly interested in you, your data, and your behavior. These include companies that gather, analyze and sell data on people for various

purposes; In these cases the motive is financial and their goal is to make more money by either selling things to you or trying to affect how you behave or think. The more ominous actors in the business of information gathering and analysis are national intelligence agencies and other organizations with equivalent resources and mandate.

The Internet has significantly affected the development of our society in recent decades, catalyzing the birth of the information society we currently live in. The Internet shapes how we do business, how we communicate, and how we participate in society. This progress is far from done, and therefore it is important to highlight existing problems within the structure of Internet. Network surveillance is a serious threat to security, privacy and trust on Internet infrastructure.

1.6 ORGANIZATION OF THE THESIS

This work is divided into three distinct parts. Part I – *Surveillance, Trust, and Information Society* – focuses on surveillance. We begin by defining what surveillance is, and examine various ways to implement it. Next, we explore how surveillance affects people and the surrounding society by exploring concepts from surveillance studies and observing them in the context of the networked information society. After examining the theoretical background, we proceed to observe real-world surveillance on the Internet. We discuss potential actors, techniques and consequences of real world mass surveillance. We also examine trust, specifically forced trust, how these concepts are modeled in literature, and how they apply to Internet surveillance and the trustworthiness of information society infrastructure.

We begin with Chapter 2, where we examine theoretical foundations of surveillance and how it affects its subjects. In Chapter 3 we examine mass Internet surveillance as a modern phenomenon, describe methods used for Internet surveillance in its current form, and also which entities are performing such surveillance actions. We also discuss what potential consequences effects global Internet surveillance has for different affected parties. In Chapter 4 we examine trust, and use the concept of forced trust to model the trust relationships in an information society

that is reliant on a compromised central infrastructure used for systematic mass surveillance.

In Part II – *Evolving Processes in the Information Society* – we focus on analyzing two essential functions of an information society: traveling across sovereign state borders and voting in elections. Both are affected by computerization and the Internet, and subsequently their potential misuse. In Chapter 5 we begin by examining the standardized biometric passport, how it is structured and secured, and how its security can be compromised. After discussing the security of the biometric passport standard, we continue the analysis from the point of view of surveillance, forced trust, and monitoring the movements of people. In Chapter 6 we examine Internet voting as a societal process, and how the previously identified problems in such voting systems manifest in the context of mass Internet surveillance.

In part III – *Security Engineering Solutions for Future Internet* – we present solutions and suggestions based on research on numerous aspects of information security to improve overall security and privacy for all members of information society, and draw conclusions based on previous observations and suggest potential future research directions. In Chapter 7 we begin by discussing a potential framework for management of personally identifiable information used in surveillance. This framework has previously been used to provide ethically sustainable methods for data management in other application areas. We continue by examining one of the most common authentication methods, the password. We analyze a set of leaked passwords from Finnish web sites and provide insight on how web users use natural language in passwords. The observations give new and more accurate information on user password behavior and also provide a better understanding on how much a small and difficult language protects against malicious attackers.

In Chapter 8 we discuss the importance of security education. We present findings on using efficient methods for network security engineering education, and present results of using the method for teaching security at university level. We identify key security knowledge for engineering students in various fields, and examine efficient methods to transfer this key knowledge to students. We further provide a preliminary analysis on security

awareness of university level IT students on various topics related to information security and the themes of this dissertation

In Chapter 9 we examine advanced encryption techniques as a solution for the requirement of ubiquitous encryption, a foreseeable direction for communication security development. The main contribution in this chapter is an efficient design flow for embedded systems. The design flow allows for rapid redesign and prototyping to meet changing requirements for encryption algorithms and other functionality. Finally, in Chapter 10, we discuss potential future work and draw concluding remarks for this thesis.

1.7 CONTRIBUTION OF THE THESIS

The contributions to the field of information security presented within this thesis are summarized in this section. A short description of the problem field and the presented contribution is provided. It is also explicitly mentioned whether the contribution has been previously published elsewhere, or it appears for the first time in this thesis. The contributions of this thesis are as follows.

1.7.1 *Part I - Surveillance, Trust and Information Society*

An overview of surveillance is presented in Chapter 2, focusing on the consequences and effects of surveillance, both on the individual and the society. Mass Internet surveillance is analyzed through the presented concepts. Implementation of mass Internet surveillance is discussed in Chapter 3, where various technologies and practices and how they are used to break security and privacy in the networked information society are presented and analyzed. All scientific contribution appears here for the first time.

1.7.1.1 *Forced trust*

The contribution in Chapter 4 is based on the author's previous work [8]. The concept of forced trust is introduced and used to describe the complex interactions between users, administrators and owners of information systems. This concept is further used

to illustrate problematic relationships in a wider scope in the networked information society.

1.7.2 *Part II - Evolving Processes in the Information Society*

1.7.2.1 *Biometric passports and border crossings*

In Chapter 5 a comprehensive discussion on the vulnerabilities of the International Civil Aviation Organization (ICAO) e-passport is presented, covering both ethical and technical aspects. The focus is on how to abuse existing and possible future systems. The problems with the biometric passport system are tied into the larger scale dilemma of forced trust in critical infrastructure and systems, and how the yet unrealized threat scenarios for biometric identifiers can become reality later, given recent developments on expansion in the use of biometrics. The contributions in this part have, in significant part, been previously published in [9, 10].

1.7.2.2 *Societal interaction and Internet Voting*

In Chapter 6 we discuss frameworks and taxonomies for Internet voting systems, and how current frameworks lack an adversary model capable of modeling potential adversaries that have the capabilities and technological exploitation possibilities on the infrastructure level of the Internet required for mass Internet surveillance. All scientific contributions in this chapter appear here for the first time.

1.7.3 *Part III - Security Engineering Solutions for Future Internet*

1.7.3.1 *Mastery over data*

In Chapter 7, the concept of Datenherrschaft — previously used by Koskinen [11] as a model for managing patient information — is applied to sensitive data that is used for surveillance purposes.

1.7.3.2 *Password security*

A comprehensive analysis on typical Finnish web user passwords is presented in Chapter 7. The analysis is focused on common

characteristics of Finnish web user passwords, including linguistic properties, as Finnish is a relatively rare, difficult and thus an interesting language from a security perspective. We observe that even though users do indeed already use natural language in passwords for real-life systems, the potential complexity provided by natural language is not realized in user passwords. In other words, people still use bad passwords, and the natural language found in passwords is quite generic, consisting mostly of base form words and names, leaving the increase in bit strength debatable.

Finally, a survey on password behavior of university level computer science and engineering students is also presented in Chapter 7. The survey was conducted using an e-learning platform, and the results are analyzed post-hoc using appropriate statistical methods. The results are borderline inconclusive due to small sample size, and more data is needed in order to make more generalizing conclusions on the password behavior of students.

All scientific contributions in Chapter 7 are previously unpublished and appear here for the first time.

1.7.3.3 *Security education*

Essential knowledge areas in information security are defined, and associated key learning profiles are drafted. Teaching information security knowledge using online education platforms to higher education engineering students is discussed, and results from three courses on the topic are analyzed.

Teaching advanced network security course on firewalls and intrusion detection, in cooperation with an industry partner, is discussed, and results from hands-on courses that have been taught for nearly a decade are presented. Efficient methods for providing state-of-the-art network security education over virtual environments are discussed, and experiences using such a system are presented.

Content analysis techniques are used to examine several aspects of higher education information security students, their perceptions, practices and motivations. These findings are analyzed and discussed.

This research has previously appeared, in significant parts, in [12, 13, 14]. An exception is the content analysis of student secu-

rity perceptions and practices, which appears here for the first time.

1.7.3.4 *Design methodology for secure hardware and adaptive cryptography*

A description of a design flow and methodology for secure communications platform is presented. The concept is explored from the point of view of the requirements placed on modern communications platforms, and the TACO framework [15] is used as the basis for the presented design flow, which provides a clear procedure for reuse and refining of functional units according to application requirements. Adaptive cryptography is discussed as a solution to the changing security requirements of modern communications, and ramifications for future implementations and the applicability to different application areas and problem sets are drafted. This contribution has been published previously [16, 17].

The concept of security dimensions is presented and a set of justifiable dimensions for sensor networks is defined. A method based on these dimensions, used for modeling a volatile security environment from the point of view of a self-aware security access control system for sensor networks, is presented and discussed. This work has previously been published in [18].

Part I

SURVEILLANCE, TRUST, AND INFORMATION SOCIETY

2

SURVEILLANCE AND THE INFORMATION SOCIETY

Scientia potentia est.

– Thomas Hobbes
The Leviathan (1651)

The Internet is the communication backbone of the 21st century society. Because of this, it is of critical importance for everyone connected to it. This is true whether you are an individual, a business, or a nation state. All layers of society use the Internet, and a significant part of communications traffic is routed through it. This makes it a truly critical component of modern society. This also poses new and increasingly difficult challenges for us as a society. Its various aspects and functions are being transformed by the Internet, adapting to the new way we communicate with each other. One of these aspects is surveillance, and this is in the focus of this chapter.

While surveillance in itself is an old concept, for most of us in the western world it has not been a visible part of our lives. Surveillance has rather been something that happens to someone else: criminals, spies, and people that live in totalitarian nations. As communication technologies have advanced, surveillance methods have similarly evolved to keep pace with the advancement of technology. This progress has stayed out of sight for most of western society.

As the Internet has grown significantly in importance during the last 20 years, surveillance and monitoring of Internet traffic has also become more commonplace, albeit this development has been shrouded in secrecy. Surveillance programs have been made visible on a global scale by whistleblowers, most notably Edward Snowden and his 2013 revelations on the United States

National Security Agency (NSA) surveillance of the Internet.[1] While the NSA is not the only agency engaged in the business of mass Internet surveillance, it can argued to be – by far, on most metrics – the most prominent one. We will discuss NSA and other actors in the global surveillance enterprise later in Chapter 3.

We begin by establishing definitions for surveillance, which we will use throughout this thesis. We discuss how surveillance can be divided into different categories and examine examples of surveillance societies in the past and present. Then we proceed to examine how surveillance is modeled and studied in science, and discuss one of the prevailing models for surveillance. We provide discussion on how surveillance affects its targets, and how all of this pertains to mass surveillance in the era of the Internet.

2.1 DEFINITION OF SURVEILLANCE

The goal of surveillance is to gather information which can subsequently be used to an advantage. In the case of *individual* or *targeted surveillance*, the goal is limited to a single known target, and the information gathering effort is premeditated. With the case of *mass surveillance*, its main goal is to gather as much information as possible on the populace. Data collected with mass surveillance can be used for various purposes, but often it has been used to exert control over the populace. The relationship between surveillance, power and control is a very complicated issue. Surveillance in itself is an active research area in social sciences. *Surveillance studies* aim to examine, define and analyze the wide field of surveillance. Providing a comprehensive review of the research within this field is outside the scope of this thesis, but we will briefly summarize some of the relevant research regarding mass Internet surveillance in this chapter.

Examples of mass surveillance and control are discussed later in this thesis.

In this dissertation we will use the same definition as Roger A. Clarke in his 1988 paper "Information technology and Dataveillance". Clarke defines *Surveillance* as [19, p. 499]

> *"[...] the systematic investigation or monitoring of the actions or communications of one or more persons. Its*

1 Electronic Frontier Foundation, "NSA Primary Sources" Online, available at https://www.eff.org/nsa-spying/nsadocs Accessed 29.1.2016.

> *primary purpose is generally to collect information about them, their activities, or their associates. There may be a secondary intention to deter a whole population from undertaking some kinds of activity."*

The concept of *dataveillance* encompasses surveillance performed through data instead of direct observation of a surveillance subject. The concept seems to predate the modern information society, as mentions of dataveillance can be found in (predominantly US law and civil rights) literature at least in the 1970's [20]. In this dissertation we will use the following definition, also by Clarke [19, p. 499]:

> *"Dataveillance is the systematic use of personal data systems in the investigation or monitoring of the actions or communications of one or more persons."*

In this context, a "personal data system" is assumed to be a device that *a)* stores, transfers and processes data, and *b)* is personal in nature and use. We assume that when Clarke wrote his paper in 1988, he was referring to a Personal Computer (PC). Back then, PCs had been available for only a relatively short time, and his understanding of a personal data system would most likely be a traditional home computer system. The development of smartphones has brought "personal data systems" into our pockets. In a 2015 study made in Finland, 69% of the Finnish population had a smartphone and 42% had a tablet computer in their personal use [21]. When compared to the PCs of 1980's, handheld devices of 2017 have more raw processing power by several orders of magnitude, and the communication capabilities have advanced similarly. The position that modern smartphones and similar devices are a very powerful potential platform for dataveillance is quite well justifiable in this context.

There also exists a difference between personal and mass surveillance. *Personal surveillance* is [19, p. 499]

> *"[. . .] the surveillance of an identified person. In general, a specific reason exists for the investigation or monitoring."*

In contrast, *Mass surveillance* is [19, p. 499]

> *"[. . .] the surveillance of groups of people, usually large groups. In general, the reason for investigation or monitor-*

> *ing is to identify individuals who belong to some particular class of interest to the surveillance organization."*

With this definition, mass surveillance can be seen as less focused on gathering information about a certain target, and more focused on gathering as much information as possible about a group of targets that fit certain criteria. Some of the key problems with mass surveillance lie here, of which we will briefly examine three important issues.

The first issue is defining the criteria that are used for selecting the targets of surveillance. In this thesis we shall refer to them as surveillance parameters. They include, but are not limited to, information such as the identity and nature of the target (e. g. the name of an individual, organization, or physical location), the methods to be used, any ramifications on the use of those methods (spatial, temporal, or contextual; e. g. personal phone calls are not included in a financial surveillance case, surveillance is limited by time or physical boundaries). Optimal choice of surveillance parameters is not trivial, and if they are chosen poorly, the surveillance effort can cause a significant amount of *false positives*; instances where surveillance provides erroneous information. This in turn reduces the effectiveness of surveillance, making poorly planned surveillance an exercise in futility. The second issue to consider is who gets to define these parameters. Those who have the opportunity to define how everyone is watched and on what criteria wields tremendous power. The third issue is what is actually to be done with the information gathered with mass surveillance. What are we allowed to do, and what is forbidden? Answers to these questions are not simple, and are also quite dependent on which side of the surveillance operation one is.

Various other definitions for surveillance do also exist. The Merriam-Webster dictionary defines[2] surveillance as *"the act of carefully watching someone or something especially in order to prevent or detect a crime"*, the Oxford dictionary definition[3] is *"close observation, especially of a suspected spy or criminal"*, and the United States military defines[4] surveillance as *"systematic observation"*,

2 http://www.merriam-webster.com/dictionary/surveillance

3 http://www.oxforddictionaries.com/definition/english/surveillance

4 US Joint Chiefs of Staff, "Joint operations," Joint publication 3-0.

for example. We can observe that in the language used to describe surveillance, criminal activities and prevention thereof are a common theme. This connection, while it is hard to show definitely, we consider its existence to be a reasonable assumption, has quite probably contributed to the negative connotation that surveillance has [22, 23, 24].

2.1.1 *Traditional surveillance*

"Traditional surveillance" as a concept can evoke mental images of a spy dressed in a long coat and a wide-brimmed hat, following a person and noting down their activities, sitting next to their targets in a crowded restaurant and attempting to eavesdrop a privileged discussion. Other possibilities include an informant quietly making observations on their neighbor, reporting to their superiors any suspicious activity that they have witnessed. While these characters are more often found in fictional agent movies and books, they do have their real-world counterparts.

Before the advent of digital communications, surveillance operations were more costly to implement, harder to coordinate and inherently risky. Mass surveillance on this scale required people who physically performed the actual monitoring and observation tasks, and reported their findings back to their superiors. The former increased the risk factor of surveillance, and the latter contributed significantly to the work factor; spying is risky and paperwork can be slow, tedious, and thus take a significant amount of human resources.

In this thesis we are more interested in mass surveillance than in personal surveillance, as the former is clearly in the scope of the thesis, while the latter is examined here shortly as an example of how surveillance used to be in the past. As we noted earlier in Section 2.1, there is a profound difference between the two. And to be even more precise, in this thesis we are interested in mass electronic surveillance, i.e. *dataveillance performed by large organizations or nation states*. Next we briefly discuss traditional surveillance on a nation state level, and two other examples of contemporary mass surveillance. While, as noted above, out of the main scope of this thesis, it is purposeful to examine these examples from the past in order to appreciate the dramatic dif-

ference in the cost function for surveillance between the past and modern times, i.e., how much more tedious, difficult and resource-consuming it used to be.

2.1.2 *East Germany*

Arguably the most classic case of a surveillance society was the German Democratic Republic (GDR), or more colloquially East Germany, during the cold war. The GDR national security service, Ministerium für Staatssicherheit (Stasi), was one of the most efficient secret services during the cold war, easily competing with and occasionally exceeding ones like the United States' Central Intelligence Agency (CIA) and the KGB of the Soviet Union. Stasi was the self-described sword and shield of the communist party, responsible for national security and suppressing any opposition. Stasi was exceptionally good at one of its primary functions, mass surveillance of East German citizens. Methods for surveillance used by Stasi included bugging apartments, hotels, public buildings, telephones, and an extensive network of informants feeding information on their fellow citizens to their Stasi handlers [25, p. 9].

The widespread surveillance practiced in East Germany permeated all levels of society and helped to create a paranoid atmosphere. A very reasonable assumption for a citizen in the GDR was that one was almost constantly under surveillance. Very few people could actually be trusted, and speaking one's mind was a very dangerous act (unless one *really* knew the people who were listening). After the fall of the Berlin wall and the unification of Germany in 1990, the Stasi archives have been opened to researchers. The surveillance society of East Germany has been a topic of intense research and also a source of controversy [26], as opening the records also gives people an opportunity to access their own records and find out that their neighbors and loved ones have actually been Stasi informants.

A significant part of East German citizens were targets of surveillance. The Stasi employed 102,000 employees and almost two hundred thousand informants [25, p. 8] to monitor a population of 17 million. They all provided Stasi information on their neighbors, colleagues, friends and even their own family. This

information was used to control the populace and to keep an eye on possible dissenters.

In the end, even the Stasi failed to change the minds of citizens. Pfaff [27] discusses coercive surveillance regimes, focusing on the GDR. In her article she concludes that even though a regime can exact obedience from its populace through surveillance and coercion, there are limits to this power. Should the credibility of an oppressive regime suffer significant damage, a revolution is more than possible – as evidenced by the fall of the GDR. While people maintained an outward veneer of conformity and compliance toward the state, in their private lives — which the secret police were unable to penetrate, even with concentrated mass surveillance efforts — they had already abandoned the message proclaimed by those in power, and their allegiance to the regime [27, p. 402].

2.1.3 *Other mass surveillance cases*

While the GDR is perhaps one of the best known examples of a surveillance society, it is definitely not the only one known to history. Mass Internet surveillance by NSA and its affiliates is examined in Chapter 3, but to keep the focus of this thesis in check, we will not examine other prominent examples in detail. Contemporary mass surveillance cases are common, though, as many nations have an interest in surveillance activities – some arguably legitimate, others perhaps not.

One of the more common mass surveillance methods is Closed-Circuit Television (CCTV). They are automated devices that capture video and audio, which is commonly saved for potential later examination. The goal of CCTV systems is to reduce crime by providing a deterrent for potential criminals, who know that their actions will be recorded. The effectiveness of CCTV is subject to a lively debate. Its effectiveness can depend on context, and some other methods, such as improved lighting and people observing spaces instead of cameras, can provide even better results than video surveillance in some cases [28].

The United Kingdom has been used as a textbook example of mass surveillance due to the extent of the CCTV network covering the country. Already in 2003, the UK was far ahead of other European nations in CCTV surveillance [29]. In 2013, there were

an estimated 11 people per surveillance camera in the UK,[5] making it one of the most strictly monitored countries in the world. Norris, McCahill and Wood [30] note that CCTV surveillance will continue to increase in quantity, but the actual benefits of this progress are more or less questionable.

The UK has also approved[6] one of the broadest surveillance laws in the world. It gives government officials the power to legally hack devices such as computers and smartphones, forces Internet Service Provider (ISP)s to store all traffic records for 12 months, and allows for mass gathering of data.

The Chinese Internet has evolved to a strictly controlled subnet separated from the rest of the world by what is called the *"Great Firewall of China"*. Also known as project Golden Shield, it has been developed in cooperation with western companies and its goal is complete surveillance of the national communication network [31, p. 15]. It is a vast system of firewalls and traffic analyzers that examine network communications within the Chinese Internet, with the capability to examine, block and modify traffic. Content that has been deemed to be inappropriate is removed from search results, and many western online services and web sites are blocked outright. Essentially all traffic that goes through Chinese borders is routed through the Great Firewall, which performs traffic and context analysis. The goals of the Great Firewall are to discourage criticism of the government, catch dissidents and make sure that Chinese Internet users are not exposed to uncontrolled information.

The Great Firewall is an interesting case of surveillance. As the Chinese Internet was designed to be controlled and monitored from its inception, the infrastructure could be designed around the paradigm of surveillance and control. Also, the existence of the Great Firewall is not a secret, neither to the Chinese themselves, or foreigners. Its presence is a known quantity, and the goals and aims of it are at least relatively well-known.

5 David Barrett. "One surveillance camera for every 11 people in Britain, says CCTV survey" (2013) *The Telegraph*, Online, available at: http://www.telegraph.co.uk/technology/10172298/One-surveillance-camera-for-every-11-people-in-Britain-says-CCTV-survey.html Accessed 2.2.2017.

6 Matt Burgess. "Snooper's Charter is set to become law: how the Investigatory Powers Bill will affect you". *Wired*, 16.11.2016. Online, available at http://www.wired.co.uk/article/ip-bill-law-details-passed Accessed 22.11.2016.

2.2 STUDYING SURVEILLANCE

The goal of surveillance, based on the definitions discussed in the previous section, can be stated as gathering information that can subsequently be used to the advantage of the surveying party. Individual or targeted surveillance is limited to a single known target, and the information gathering effort is premeditated. Mass surveillance is targeted indiscriminately, and the ultimate goal is to find some pattern of interest in the observed data that justifies individual or targeted surveillance. It is obviously counterproductive indeed to keep watch on the whole populace, collect massive amounts of data on everyone, and finally never use that data for any particular purpose. By this criteria, mass surveillance is a precursor to targeted surveillance, an effective method for canvasing predefined triggers through a set of potential targets in the hopes of finding something of interest.

The relationship between surveillance, power and control is a very complicated issue, and surveillance is an active research area in social sciences. It is examined in *surveillance studies*, in all its forms and manifestations in society. While a full discussion on surveillance studies, its central research topics, questions, and problems is definitely outside the scope of this thesis, we will briefly outline a key theory pertinent to our examination of mass Internet surveillance. For a more in-depth introduction to these topics, the reader is encouraged to see, e.g., the comprehensive introductions to surveillance studies by Lyon [32, 33].

2.2.1 *The Panopticon theory*

One of the central theories of surveillance studies is the Panopticon theory of surveillance. It is based on the works of Jeremy Bentham, an 18th century philosopher, who introduced the Panopticon as a prison that relies on continuous surveillance and complete lack of privacy for inmates [34]. The concept of Panoptic surveillance was refined to its current form by Michel Foucault in *Disclipine and Punish: The birth of the prison* [35].

In essence, the Panopticon is a model for efficient surveillance and self-imposed control. The key idea is that perceived surveillance is as good as actual surveillance, and that surveillance – whether perceived or real – changes the behavior of the sub-

ject towards whatever is "desirable". This combination makes it possible to control and modify the behavior of surveillance subjects even with perceived surveillance, making the collective surveillance effort easier and less resource intensive. When subjects perceive surveillance to be constant, they adapt their behavior to match this perception, regardless of whether that perception is founded in truth or not.

The Panopticon was initially presented as a prison design that would make the inmates supervise themselves, and lead to change in inmate behavior. In a Panopticon prison, the cells are continuously illuminated and arranged in a circle. The wall facing towards the center is transparent, making everything in the cell visible to an observer. The prison guards are stationed in a tower, placed in the structure's center. Windows on the tower are opaque when viewed from the outside, and transparent when viewed from the inside. This makes it possible for the guards to observe the prisoners at any time. Because of the opaque windows, the prisoners are not able to see the guards at all, and cannot know whether they are being watched or not. This perception of surveillance essentially forces the prisoners to adjust their behavior to the knowledge that they can be under observation any given moment. In this manner, panoptic surveillance changes the behavior of its subjects by forcing self-censorship and adaptation to conform to the expected behavior standard.

A classic example of a Panopticon in literature is the telescreen in George Orwell's classic book *Nineteen eighty-four*.[7] In the book, a telescreen is a television-like device installed in every home. It provides the same functionality as a television screen, but also makes it possible to for the authorities to observe those who are in the room. This observation happens without the knowledge of and any indication to those in the room. People are in turn aware that someone could be watching them through the telescreen *at any moment*. This kind of ubiquitous surveillance without any notification makes the telescreen an excellent example of a panoptic surveillance device.

Devices with similar capabilities as the telescreen do exist in real life, even though they are not used in a similar manner. For example, Microsoft's Xbox 360 and Xbox One gaming consoles have the capability, through the Kinect motion detection sensor,

7 Orwell, George. *Nineteen Eighty-four* (1949).

to identify whether there are people in the room, what they are doing, and record audio of what they are saying. The motion detection accuracy of the Kinect sensor [36] has been tested to be sufficient for applications such as gait analysis [37] and Parkinson disease monitoring [38], so it can be considered sufficiently accurate to conduct surveillance. When Microsoft introduced Xbox One in 2013, one of the pre-release announcements was that the Kinect sensor was mandatory for the device to function. This requirement was later rescinded, partly due to resistance from consumers and privacy advocates.

The presence of smart devices in living rooms is growing steadily. Analysts predict 50% market penetration for smart TVs for households in Japan, US, and Europe.[8] As the prevalence of such smart devices increases, security and privacy aspects also become salient. In March 2017, Wikileaks published classified documents[9] that allegedly described the hacking capabilities of the United States CIA. In these documents an attack codenamed "Weeping Angel" that exploits vulnerable Samsung smart TVs is described.[10] According to this leaked document, Weeping Angel has the capability to turn a compromised TV into a remote microphone that is capable of monitoring its surroundings with the built-in microphone. This data can be further processed and sent to the attacker. This kind of surveillance capability and method is eerily similar to the previously described telescreen example from Orwell's classic book.

The Panopticon theory provides us a compelling background upon which to examine the different manifestations of surveillance. The Panopticon indeed evokes mental images of an "all-seeing eye", a concept perhaps best conceptualized by fictional entities such as Sauron in *The Lord of the Rings*.[11] Sauron was depicted as a demigod capable of observing anything and anyone in the world with his giant eye, and it took the heroes great

8 Hisakazu Torii, IHS Markit. "More Than Half of All Households in Japan, US and Europe Will Have Smart TVs by 2019" Online, available at https://technology.ihs.com/571765/more-than-half-of-all-households-in-japan-us-and-europe-will-have-smart-tvs-by-2019 Accessed 31.3.2017

9 Wikileaks press release. "Vault 7: CIA Hacking Tools Revealed" Online, available at https://wikileaks.org/ciav7p1/ Accessed 31.3.2017.

10 "Weeping Angel (Extending) Engineering Notes" https://wikileaks.org/ciav7p1/cms/page_12353643.html Accessed 31.3.2017.

11 Tolkien, J.R.R. *The Lord of the Rings* (1955).

effort to avoid his gaze. Drawing parallels between the Panopticon and modern mass Internet surveillance is not unreasonable due to the sheer information processing capability of modern computers. We will examine this parallel further in Section 2.4.1, where we discuss how the Panopticon concept can be realized with information systems.

2.2.2 *Criticism for the Panopticon theory*

The Panopticon as a model has been criticized in literature to be both inaccurate and unsuitable for modeling real world surveillance. Bauman [39] argues that the Panopticon is not valid in a hedonistic Western civilization, and that it would be only applicable to a "clockwork" society, where everything must serve a certain purpose and there is no room for deviation from the norm. Boyne [40] examines several socio-theoretic arguments against the Panopticon as a model for surveillance. Among them, he notes [40, p. 299] that the goal of Panopticism is to make the subjects to police themselves, making physical manifestations of surveillance eventually obsolete. In the end Boyne concludes that while there are persuasive theoretical arguments against the Panopticon, at the same time a development towards a general panoptic society can be observed [40, p. 302].

Boyne also mentions an emerging paradigm shift from observation to prediction, and thus away from the constant surveillance of the Panopticon [40]. This paradigm shift can be seen in mass surveillance targeting potential criminals or terrorists. It can be described as *semi-Panoptic surveillance*, in which the goal is to gather all possible information, but without affecting the behavior of subjects by making them aware of the surveillance. Studying past events and observations in order to ascertain the motives and probable future actions of a subject are at the heart of prediction. If we monitor our subjects all the time and they are aware of it happening, their behavior will change. The goal then, is to make this surveillance invisible, or even better, deniable. In this way the original behavior of subjects does not change.

Haggerty [41] argues that the Panopticon is no longer a suitable model of surveillance, based on the vast number of different refinements to the model, and because *"[e]ach new 'opticon' points to a distinction, limitation, or way in which Foucault's model does not*

completely fit the contemporary global, technological or political dynamics of surveillance." Indeed, the sheer number of different refinements[12] indicates that the Panoptic model is not fully capable of capturing all nuances of surveillance. It would be tempting to create "Yet Another Opticon" for the purpose of this dissertation, but that would quite likely be counter-productive. Instead of trying to define a Panopticon capable of describing observed characteristics of mass Internet surveillance, we will leave this to the social scientists.

All in all, the Panopticon is a convenient tool for us to imagine what complete and total surveillance could encompass. It stands to reason that a model prison envisioned in the 18th century – a prison that was never actually built – would be inadequate in modeling 21st century mass surveillance. We will rather use the Panopticon as a backdrop in the discussion about mass surveillance and acknowledge its shortcomings as a complete model for surveillance.

2.3 SURVEILLANCE AND BEHAVIOR

Human beings have a tendency to behave rationally only in economic models, where the idea of people making rational, informed choices is assumed to be the norm. Actual human beings behave in unforeseeable ways and make decisions that are not optimal or rational for them, at least apparently. We all make "bad" choices all the time.

Certain things about behavior can be measured, however, and the effect of surveillance is one of them. It has varying effects on our behavior, some are prominent, some are more subtle. The actions of a person tend to be dependent on whether that person is aware of being under surveillance or not. In general, people prefer not to act in a manner that reflects unfavorably on them if other people are witnessing said actions. This means that taking the action contains some sort of inherent risk due to others' possible reaction. This makes the action unfavorable, and makes people to avoid it, unless the risk is deemed to be worth the gain. Then again, in many instances people tend not to think

12 See [41, p. 26] for a non-exhaustive list.

about the risks and end up doing unfavorable actions regardless of surveillance, peer pressure or reputation.

While decision making is in may ways relevant to how and why surveillance affects us, it is certainly not the focus of this dissertation and therefore this matter is only approached in this very cursory fashion. For further reading in the area of decision making under risk, as a starting point the reader is directed towards work by Kahneman and Tversky [42]. We will, however, briefly discuss how the knowledge of being under surveillance affects a person.

2.3.1 *Chilling effect*

Assume a scenario where certain actions are forbidden in society and punishments for that behavior are also strictly enforced. It would then stand to reason that anything even *resembling* this forbidden behavior would be actively shunned by people, motivated by the fear of association. This effect is exemplified in discussion about free speech, where this kind of behavior effectively translates to self-censorship. If one knows that speaking about particular subjects in a certain manner, however mildly, is considered taboo, one will prefer to avoid the subject altogether.

This is known as the *chilling effect*. As a concept "chill" in free speech discussion – to which surveillance is clearly intertwined – appeared in a 1950's case[13] of the Supreme Court of the United States (SCOTUS). The concept "chilling effect" appears to have been first mentioned by SCOTUS in the 1950's.[14] Schauer [43] discusses the chilling effect, examining it through the perspective of the US judicial system and cases relating to the first amendment.[15] He notes that in an ideal world the judicial system would be perfect, and that there would be no mistakes made or uncertainty in the legal process. Because this is not the case, Schauer notes that the outcome of the justice system cannot be reliably predicted, and that *"even lawful conduct may nonetheless be punished because of the fallibility inherent in the legal process"* [43, p.

13 Wieman v. Updegraff, 344 U.S. 183, 195 (1952)

14 Gibson v. Florida Legis. Investigation Comm., 372 U.S. 539, 556-57 (1963)

15 The first amendment to the US constitution guarantees the right to free speech in the US.

695], especially if it falls close to the border between accepted and unaccepted conduct.

The chilling effect can be argued to apply directly to network surveillance. When one is aware that all Internet traffic is monitored for "suspicious" content, it is reasonable that one will avoid any behavior that could be interpreted as even close to being suspicious. One relatively common example of this is trying to avoid browsing topics related to terrorism on the Internet. Being labeled a terrorist due online behavior such as browsing habits and online messages is a realistic risk, and the consequences of being erroneously targeted as a terrorist can be dire to an individual. Especially in the social atmosphere of the 2010s, these consequences can range from minor inconvenience to imprisonment, or worse.

Recently, Stoycheff [44] has examined the connection of surveillance and the chilling effect. She studied what kind of effect perceived surveillance has on the willingness of people to speak out on public forums, especially when related to minority political views. She uses Noelle-Neumann's theory of the Spiral of Silence [45] to examine human behavior when communicating in a modern hostile communication environment – the comment section of news articles on Facebook. The Spiral of Silence theory describes the behavior of people in a group with radically differing opinions, and postulates that when people are aware of potential backlash to their opinions, they rather keep silent. This silence in turn begets more silence from others who know their ideas are in the minority.

In her study Stoycheff found that a perceived hostile opinion climate does indeed negatively affect the willingness to speak out online. Another finding was that if people are aware of surveillance and feel that the surveillance is justified, they tend to conform their opinions to the perceived majority, and staying silent otherwise. Stoycheff also identified a group of people with different behavior. Those who strongly feel that government surveillance is not justified were not affected by perceived surveillance and opinion climate on sharing their views and opinions publicly. They rather made their choices on sharing their views regardless of potential surveillance. Stoycheff speculates that this group could be comprised of well-educated and vocal individuals who are not affected by surveillance, and also of those who

are so cautious that they refrain from sharing their views publicly at all.

A recent example of the Spiral of Silence potentially in action can be seen in the 2016 presidential elections in the United States. While pre-election polls and media discussion favored Hillary Clinton, sometimes quite strongly, the final result was that the election went to Donald Trump. Voters who were aware of the hostile opinion climate – similarly to those examined by Stoycheff in her study – and yet ended up voting for Trump were seemingly reluctant to respond to polls, but were more active on election day, rather saving their opinion for the ballot box.[16] One potential explanation is the Spiral of Silence, where an unpopular opinion in social media is repressed only to be expressed in another way. Obviously, the more vocal proponents of Trump did not care whether the opinion climate was hostile or not, again in line with Stoycheff's findings.

2.3.2 *Somebody's watching me — Perceived surveillance*

As it was previously noted, if a person is aware of being under observation by other people, it is probable that the person will refrain from actions that would reflect negatively on them. This effect has been successfully measured in a controlled setting. Bull and Gibson-Robinson [46] studied the effect of an individual's eye-gaze on another person. They tested whether people would donate more money to charity when under the gaze of another person. The subjects were observed in situations where the donation collector looks them right in the eye when giving the donation. These were compared to situations where the collector's gaze was focused on the collecting tin instead of the donor. The gaze was found to be a potent factor in several collection situations, and that significantly more money was donated in the case of the collector looking the donor in the eye. While examining group formation and dynamics, Kurzban [47] found that men in particular were inclined to donate more of their own resources to the common good when having eye contact with others in

16 Ramin Skibba. "Pollsters struggle to explain failures of US presidential forecasts". *Nature*, 09.11.2016. Online, available at http://www.nature.com/news/pollsters-struggle-to-explain-failures-of-us-presidential-forecasts-1.20968 Accessed 21.6.2017.

the group. This same effect was not found with women, though, suggesting some alternative group dynamic.

The human gaze has an effect on people's behavior. However, the actual gazing does not have to be directly performed by a human to have this effect. Haley and Fessler [48] were among the first to observe that people behaved differently in a social setting when in the presence of images of eyes. This effect was studied during a round of the Dictator Game, a game where one player ("the dictator") gets to divide a sum of money between himself and the other player ("the recipient") in any manner he chooses, with complete anonymity. The recipient has no way to affect the decision of the dictator, who is able to take all of the money for himself without any consequences, should he choose. Haley and Fessler found that when subtle eye-like stimuli such as images were present during the decision, the dictator would act more generously towards the recipient than without the presence of such stimuli.

Burnham and Hare [49] studied decision making in people that are in the presence of a robotic eye. They found that subjects who were under surveillance contributed 29% more to the public good in an economic game, even when the surveillance was in the form a robotic eye. Ernest-Jones *et al.* [50] observed that people alter their behavior if they are in the presence of posters depicting eyes. They examined whether people were less likely to litter in a cafeteria setting when there were posters depicting eyes on the walls. They found that in a larger group, posters with eyes had less effect, but in smaller group sizes, the same posters coincided with less littering. Francey and Bergmüller [51] report similar results on the effect of eye images and littering. Nettle *et al.* [52] observed both in their experimental and meta-studies that the presence of eyes in the Dictator Game made people more likely to donate at least something (as opposed to taking everything for themselves) when eyes were present.

The results on the *watching eye effect*, as Nettle *et al.* call it, are not unambiguous, though. While researching surveillance and generosity, Sparks and Barclay [53] found that while the effect of eye images on behavior was measurable, its effect fades over time. After test subjects had prolonged exposure to eye images, there was no perceived difference in generosity with the control group with no exposure. Ekström [54] was also able to measure

the effect of eye images on behavior, but based on his observations he argues that measurements of the effect in previous studies have been biased upwards. Nettle *et al.* [52] found that while people were more likely to donate in the presence of eyes, they were not more generous in their donations.

This does not match the results previously presented by Haley and Fessler. Nettle *et al.* [52] also note that both successful and null results have been reported in replication studies. Given that a significant part of the results in psychological studies cannot be reproduced in subsequent tests [55], this kind of variation in such tests could perhaps be considered to be normal in this discipline, without attributing the discrepancy to malice or foul play by dubiously motivated researchers.

The Panopticon prison was designed explicitly to change the behavior of its inmates. Bentham was keenly aware of the effect that constant surveillance has on its targets, and he was certainly not the first person to perceive this phenomenon. This only goes to show that the reaction to perceived surveillance is a profoundly human condition that exists irrespectively of the technology used. Overt, direct surveillance provokes change in behavior of its targets towards more desirable patterns from the viewpoint of those performing the surveillance.

This is one observation that, when put into context of mass Internet surveillance, reveals a key underlying problem with current forms of network surveillance: a *de facto* chilling effect on online behavior brought about by the knowledge that all actions taken online are monitored and stored for further analysis. The actual targets of these surveillance efforts are still able to continue their operations normally as long as they know how to avert surveillance by using strong cryptography — mathematical tools used to secure information from unauthorized parties — and strict operational security procedures. More detailed discussion on cryptography can be found in Chapter 3. The people who know about surveillance but are not its intended targets per se (i.e. not terrorists or criminals) end up changing their behavior (as the knowledge that you are being monitored does have an effect on human beings). Even if one professes to be indifferent about surveillance, one cannot rule out that their behavior does not subtly change unconsciously.

2.3.3 *"I've got nothing to hide"*

A dismissal of negative effects of surveillance can be phrased as the *I have nothing to hide, so I have nothing to fear* argument. It is commonly used to play down any harmful effects of surveillance by implying that if you have done nothing wrong, you have nothing to fear. It can also be used to imply that those who object to surveillance indeed have something to hide, thus forcing the objectors to defend themselves against suspicion.

Solove [56] critiques this argument from a privacy viewpoint. He argues that we all have something to hide, whether we know it or not, and failure to realize this is due to a misunderstanding of the basic nature of privacy. He observes that those who support the *nothing to hide* argument generally tend to see privacy as negative – a necessity for hiding something harmful or illegal. Schneier observes[17] that those who espouse the *nothing to hide* argument *"accept the premise that privacy is about hiding a wrong."*

Solove is an American lawyer and thus understandably approaches the question generally from the point of view of American legislation. As the nature of privacy, and more importantly the need for it is generally universal in human beings, his argument about the nature of privacy can be reasonably generalized to other jurisdictions as well.

2.3.4 *Notes on surveillance and behavior*

According to existing literature and research, there is a strong connection between surveillance and behavior. Surveillance changes the behavior of its subjects, sometimes quite dramatically, and this observation has already been used as early as in the 18th century to create a hypothetical self-governing, behavior-altering prison. Dictatorships and oppressive regimes have used surveillance as a method for control and suppression throughout history, and there are no indications that the method would have lost its effect, at least judging by the continued use of surveillance in contemporary societies.

17 Bruce Scheneier, "The Eternal Value of Privacy" (2006) Online, available at https://www.schneier.com/essays/archives/2006/05/the_eternal_value_of.html

Surveillance in itself can be used for a multitude of purposes. In Chapter 3 we shall discuss the implementation and justification of mass surveillance in the context of the Internet and its infrastructure, but next we will examine surveillance in the wide context of an information society.

2.4 SURVEILLANCE AND THE INFORMATION SOCIETY

We have already drawn a (light-hearted but yet apt) parallel between the Panopticon and a near-omniscient supernatural being from the Lord of the Rings. When we apply modern computer systems with their advanced capacity for gathering, storing, retrieving and processing information to data analysis of surveillance data, that comparison begins to seem even plausible. A surveillance system that uses the majority of data from the Internet as its input gives us a compelling image of an omniscient surveillance apparatus capable of extracting meaningful data from even the smallest digital footprint we leave behind on the Internet.

In accordance to one of our previously used definitions in Section 1.3, an information society is one where information is the chief driver of socio-economics. Given that surveillance is all about information, we can quickly come to the conclusion that the information society provides also ample opportunities for surveillance; as the information society transforms all aspects of society, surveillance is naturally among those transformed into a new form of potential surveillance. For those who are inclined towards cynicism, the nature of the information society directly leads to the emergence of a surveillance society. This development is facilitated by data on all citizens that is plentiful and readily available. When all aspects of society converge to facilitate mass surveillance, we encounter a completely new situation where even seemingly unrelated functions and processes all contribute to the wholesale surveillance of the society.

In the following sections we will briefly examine this transformation, how it has been observed in literature, and what we can learn from this transformation process as we move further towards mass Internet surveillance.

2.4.1 *Dataveillance and the electronic Panopticon*

Consider surveillance that is not necessarily performed in real time. Instead of being a one time event — a single recorded discussion, for example — conversation information is gathered continuously and stored for processing at a later time. This type of surveillance relies on data collected by observations over time, and that data is mined for information later.

Now consider that all of this is performed using computers, also leveraging the scalability of computer-based solutions to ramp up the scale of the surveillance operation. This kind of scenario falls clearly into the definition of dataveillance, and it is clearly relevant to the discussion on surveillance in information society. Dataveillance actually is an apt description of modern data-driven surveillance efforts, where those who are under surveillance are represented primarily – or solely, as we will later observe – by data in a computer system. To describe this kind of surveillance apparatus, Lyon — among others — uses the term "electronic panopticon" [57].

Shoshana Zuboff [58] discusses and examines the *information panopticon*, examining panoptic surveillance in the context of work and monitoring personnel in the workplace. Even though information technology has developed significantly since she wrote her work in 1988, the key issues she describes are acutely relevant to current events [58, p. 322].

> *"[Information systems] can become information panopticons that, freed from the constraints of space and time, do not depend upon the physical arrangement of buildings or the laborious record keeping of industrial administration. They do not require the mutual presence of objects of observation. They do not even require the presence of an observer. Information systems can automatically and continuously record almost anything their designers want to capture, regardless of the specific intentions brought to the design process or the motives that guide data interpretation and utilization."*

Zuboff is among the first to describe the capabilities of information systems to act as surveillance tools, and makes important observations on their capabilities. One such observation related

to monitoring employees in a work environment is that while some managers saw such systems as tools to help train employees, other managers saw it to be more convenient to assign blame and justify firings [58, p. 317]. This example clearly demonstrates an early instance of unintended consequences coming from adoption of surveillance technology. Even if the intention of the creator of the surveillance was benign, others have later subverted it to other purposes.

The electronic Panopticon concept is not without its critics. For example, Bain and Taylor criticize the use of the Panopticon metaphor in the context of a workplace with strict surveillance of employees: the call centre [59]. They conclude that while some aspects of the Panopticon theory are arguably present in the call centre environment, the comparison to Jeremy Bentham's Panopticon prison is "grossly overdrawn".

As it was noted before, modern computer systems are capable of powering an Internet surveillance system that monitors practically everyone, everywhere, at every occasion, unless the target takes great pains to avoid it — if even then. The *unit cost* for surveillance has become a matter of physical – processing power or data storage – resource, not human resource. *This* is the key difference between old world mass surveillance schemes in East Germany or the Soviet Union, and modern mass Internet surveillance systems.

2.4.2 *Metadata*

A common argument for using metadata to represent actual data for surveillance purposes instead of using real communications data and context is that *metadata is only data about data*; as it contains no actual details of the content of the exchange, no privacy violation happens. This is quite easily demonstrated to not be true. Bruce Schneier noted[18] that "metadata equals surveillance", giving an example of a detective hired to eavesdrop on a target. Data about who the target interacted with, where the target was located in a particular time, it is the surveillance data in itself, not metadata derived from the actual surveillance data.

18 Bruce Schneier. "Metadata Equals Surveillance" *Scheneir on security blog*, 23.9.2013. Available online at https://www.schneier.com/blog/archives/2013/09/metadata_equals.html Online, accessed 9.5.2017.

It is possible to infer more information from metadata than it immediately meets the eye: metadata is sufficient in itself to give a strong implication on the actual data it is derived from, in addition to providing a whole new layer of information on the whereabouts and activities of the surveillance target. Schneier [7] gives an example about a Stanford study [60] on what information can be deduced from phone metadata. The researchers collected data such as when the calls were made, from whom and to where, but without any knowledge of the actual messages in the phone conversation. The researchers were able to identify people with medical conditions, owners of firearms and drug offenders, just from observing this metadata.

Zappalà [61] discusses the problems of using metadata to wage war. While metadata is on one hand considered to be just circumstantial, derived from the actual data and not indicating anything about the original data, on the other hand metadata is actually quite accurate – as evidenced by the Stanford study discussed above. Nevertheless, we encounter some gray areas when we consider what actions can be taken based on metadata. While metadata-based mass surveillance can lead to potential targeted surveillance of a person of interest in many cases, sometimes the actions taken are orders of magnitude more severe, ranging up to using lethal force. General Michael Hayden, former director of the NSA and the CIA, has stated on the record[19] that "*we kill people based on metadata.*" In this context, metadata should be considered to be as valuable and damaging to privacy as the actual content of communications.

2.4.3 *Data double*

The concept of a *data double* is discussed in the context of surveillance by Haggerty and Ericson [62]. We all have a data double, a digital representation of ourselves. It is a product of all that we read, write or interact with in the digital world. On a conceptual level, at least, our data doubles accrue matter over time, as all our actions online are "stored" in our data double. Whether this kind of data double actually physically exists de-

19 This quote can be found at 18:01 on the video *The Johns Hopkins Foreign Affairs Symposium Presents: The Price of Privacy: Re-Evaluating the NSA.* Available online at https://www.youtube.com/watch?v=kV2HDM86XgI Accessed 28.11.2016.

pends on whether some instance has taken the time and used the resources to gather all individual actions online to a single entry.

It can be quite convincingly argued that from the point of view of a system that only observes the digital world, we as people *are* our data doubles, at least as far as the system is concerned. Mass surveillance systems targeting the Internet definitely fall into this category, and can be seen as processing, analyzing, and examining our data doubles. As the operators of mass surveillance systems are commonly state level authorities, they do indeed have the power to interact with and affect us physically. This shows that the coupling between our physical selves and data doubles is quite intuitive.

2.4.4 *The Surveillant Assemblage and Rhizomatic surveillance*

For the sake of providing a more comprehensive view on the the forms of studying and modeling surveillance, we will briefly introduce some concepts that closely relate to mass surveillance and its effects. Further examination of these concepts are however out of the scope of this thesis.

The *surveillant assemblage* is a concept introduced by Haggerty and Ericson [62]. It can be thought of as a widely distributed and decentralized surveillance apparatus that incorporates information from many unrelated sources, consolidating and distilling data to form a coherent and comprehensive data double of its targets. As more and more aspects of society and life are integrated into the surveillant assemblage by providing source data to surveillance efforts, it becomes increasingly difficult to avoid leaving traces of oneself into the system. Haggerty and Ericson describe this progress as the "disappearance of disappearance"; that it is getting increasingly hard to disappear into the crowd, maintain anonymity, or hide from monitoring by social institutions. Especially, Haggerty and Ericson argue [62, pp. 619–620], it takes effort and causes you to forfeit social benefits and privileges such as voting, social security, banking, and even using the Internet in general.

The surveillant assemblage is remarkably accurate in representing the experience in the networked information society. As an example, in Finland it is very challenging to use societal ser-

vices without a computer, an Internet connection, and online banking credentials – the *de facto* online ID in Finnish society. To say that without an online presence one might as well not exist is not a great exaggeration.

Now consider if all of the data that is given to use these services would be consolidated to form a dossier on the person. It would be very accurate, and if this were to be exploited in bad faith, the damage done to the person could be irrevocable. This is why the electric Panopticon is such a frightening concept, even though the actual risk of this happening to an individual can be considered to be negligible.

Another related surveillance concept is *rhizomatic surveillance.* Based on the work by Deleuze and Guattari [63], rhizomatic surveillance is used to describe surveillance characterized by growth trough expanding use and a leveling effect on hierarchies. This means that all parts of society, even those completely unrelated to surveillance at face value, are used as a part of a larger surveillance operation — a kind of function creep in the direction of surveillance. The continuous incorporation of seemingly unrelated features and functions of society into a distributed but yet connected surveillance apparatus describes the concept of rhizomatic surveillance quite accurately. For more reading into the matter, see e. g. the work by Haggerty and Ericson [62]. Ellis, Tucker & Harper [64] provide criticism on the surveillant assemblage, arguing in turn that the combination of decentralized surveillance devices that form the surveillant assemblage is better described as an *atmosphere of surveillance*.

The existing literature on surveillance studies is large and expands constantly as further advances are made. A relatively comprehensive overview of necessary key concepts and theories pertinent to the topic at hand has been presented in this section. While the discussion on the finer points of terminology in theoretical concepts of surveillance is certainly intriguing and serves a purpose in its own field of reference, for us further elaboration and analysis is out of the scope of this thesis. We must proceed further to actual practical surveillance, how it is performed, by whom, and how the use of controversial surveillance methods is justified to the people in countries where this happens.

2.5 CONCLUSION

It is difficult to conclusively measure the perception of a populace towards mass surveillance. It is a tool mostly used by "the good guys", but mass surveillance is almost always directed towards the populace itself that, at least in democracies, ultimately is responsible for the adoption of mass surveillance. It is difficult to ignore the increasing support among the general populace for mass surveillance as a valid tool to provide security in societies. We can only hypothesize the motivations behind such support for something that is at the very least controversial.

Mass surveillance has become the go-to tool in modern surveillance systems. The reasoning for this is actually at least seemingly sound. Because of the complexity of Internet communication systems and the sheer amount of data that is transferred on these networks, combined with the difficulty of pinpointing the target of surveillance from other people, it quickly becomes more feasible to just gather all possible information and analyze it later. This approach is not without its critics, though. A report by the New America Foundation[20] claims that NSA mass surveillance is ineffective in combating terrorism, and advocates traditional investigative practices as more efficient. This view is echoed by others, such as Schneier [7], for example.

People do not in general understand either the threat posed by or the consequences of surveillance. If people know they are monitored, surveillance does incite changes in behavior, and overt actions of surveillance are felt on a truly visceral level. But if no negative consequences ever realize due to surveillance, it would be a reasonable assumption that people would "deal with it", so to speak, and continue their lives under even constant surveillance – if only the negative consequences never happen to them or anyone they know or care about. This kind of apathy in the public would be the goal for anyone trying to implement a systematic surveillance effort that would monitor truly everyone.

20 Bailey Cahall, Peter Bergen, David Sterman, and Emily Schneider. "Do NSA's Bulk Surveillance Programs Stop Terrorists?"*New America Foundation*, 2014. Online, available at https://www.newamerica.org/international-security/do-nsas-bulk-surveillance-programs-stop-terrorists/ Accessed 21.6.2017.

People require feedback for their actions to change something in their behavior. If you know that speaking about certain topics, or holding the "wrong" political views will get you in trouble, people will refrain from such adverse actions, at least publicly. The goal of mass surveillance is to notice both public and private manifestations of such unwanted behavior. Or, on general layman terms, the goal is to catch criminals and terrorists before they do any harm. This goal in itself is laudable, and this is commonly taken as justification to escalate efforts in hopes of finding more instances of targeted unwanted behavior.

In the past this kind of (ab)use of power has required surveillance and processing of gathered data. Now it is a totally different game. It is possible to practically gather and store all data produced in the world – and this is what the NSA is trying to do. If the data is not acted upon immediately, but only at a later date, people will not have the feedback of negative experiences. It is possible to gather "dirt" on all people and to use it only when it is convenient, or more likely, when it is found, as processing such large quantities of data takes a lot of time.

When we justify violations of personal rights with the goal of greater good and protection of society, extreme care must be taken. It is tradition in writings on this topic to quote – albeit disputedly – Cardinal Richelieu: "*If you give me six lines written by the hand of the most honest of men, I will find something in them which will hang him.*" We must consider very carefully to whom, if at all, we give those metaphoric six lines of text on ourselves. And if we choose to entrust our personal and private information into the hands of others, we must contemplate whether they do indeed have our best interest in their heart or not.

Surveillance and especially mass surveillance has significant potential for misuse in a very large scale, and these questions should be examined from as many points of view as possible. Laudon [65] observed that *"technology does not stand "outside" of society, but instead is a social phenomenon itself, subject to all the constraints of other social actors."* This holds especially true for solely considering the technological perspective on mass surveillance, as technology does not exist an ethical vacuum, and the issues with mass surveillance warrant careful and thorough consideration.

3

INTERNET SURVEILLANCE — ATTACKING THE INTEGRITY OF INTERNET

The enemy knows the system.

– Claude E Shannon
Also known as *Shannon's Maxim*

I'll create a GUI interface using Visual Basic, see if I can track an IP address.

– Lindsay Monroe
CSI:New York (2008)
Season 4, episode 20: "Taxi"

When a person turns on their personal computing device and opens a web browser, email client, or their favorite media streaming platform, they are immediately tracked. When they proceed on their business online, they are being tracked, profiled and categorized on several different levels. The information potentially stored by the browser lets an observer know, based only on the browsing history, a lot about its user: their job, interests, hobbies, political and social opinions, who they communicate and associate with, and so on. The same applies to the information available to the operating system on their computer or tablet, as it is responsible for, well, everything that gets executed on the device. Their ISP is similarly aware of things they do online, as it is responsible for carrying and processing the network traffic. Every online service knows things about their users – information users have chosen to divulge when choosing to use their services and agreeing to relevant terms. When such distributed fragments of data are consolidated, it can be used for purposes other than what was originally intended.

Surveillance has ceased to be a resource intensive activity, and has nowadays become more dependent on the efficiency of data acquisition [7, p. 23-28]. In the past, all personal data was spread into various different information systems, many of them not even connected to a network. This made combining personal data from these various systems to a centralized accessible database of information prohibitively expensive, if not outright impossible. Two critical developments have made this possible. First, nowadays the cost of computing has gone down: the cost of a gigaflop – one billion floating-point operations per second – in 1984 was almost 19 million USD, approximately the cost of a Cray X-MP supercomputer. In 2016, a single Sony Playstation 4 gaming console has a peak capacity of over 10 teraflops. The cost of data storage has plummeted at the same time, algorithms have been improved, and the development of communication systems has multiplied the amount of data in motion on the Internet.

When adjusted for inflation, 19 million USD in 1984 is over 54 million USD in 2016.

These changes in the global network environment have made surveillance possible on a whole new scale. As we observed in Chapter 2, surveillance affects its targets in various ways. In this chapter we examine how global Internet surveillance is implemented. First we briefly discuss the concept of infrastructure attacks against the Internet and its critical components. Next we examine the scope and extent of the NSA surveillance programs to give an overview of the scope of an advanced mass surveillance operation. We also consider the situation in the Nordic countries regarding mass surveillance. Then we proceed to examine different technological methods and operating procedures for realizing mass surveillance on a global scale, and consider how they affect the integrity of the Internet. We also briefly touch upon commercial Internet surveillance as a phenomenon. Next we proceed to examine the potential consequences of mass surveillance, and how it affect the Internet, and what effects it has already had. Finally we ponder briefly on the justification of mass surveillance, when it is actually warranted and legitimate, and whether the seemingly ever-conflicting interests of privacy and legitimate surveillance can ever be not in conflict.

3.1 ATTACKING THE INTERNET INFRASTRUCTURE

In this thesis we use the term *infrastructure attack* to refer to an attack that is directly targeted against, or takes advantage of, the basic infrastructure of the Internet. This in turn refers not only to the actual routers, cables, modems and other equipment that physically make up the Internet, but also the protocols and other immaterial properties that define the nature of Internet communication.

Because these attacks are targeted against the communication infrastructure, they are indiscriminate in how they affect their targets. In this regard, infrastructure attacks could even be argued to be comparable to weapons of mass destruction in conventional warfare. In this chapter we analyze different attack vectors and vulnerabilities that have been exploited in real life. The side effects of infrastructure attacks affect not just the intended targets, but everyone who happens to be in the (metaphorical) vicinity of a infrastructure cyber attack, resulting in for example denial of service, loss of data, loss of privacy, or financial damage.

3.1.1 *NSA and Five Eyes surveillance programs*

In the past infrastructure attacks have generally been relatively low-scale, possibly accidental, and often not motivated by malice but ignorance. As an example, some of the protocols that define the core functionality of the Internet are not particularly robust when faced with a determined attacker. For example, Border Gateway Protocol (BGP)[1] can be used to redirect traffic to unintended locations. Previous cases show both intentional[2] and

1 Internet Engineering Task Force. RFC4271: A Border Gateway Protocol 4 (BGP-4) Available online at https://tools.ietf.org/html/rfc4271 Accessed 7.2.2017.

2 RIPE Internet Coordination Centre. "YouTube Hijacking: A RIPE NCC RIS case study. Online, available at http://www.ripe.net/internet-coordination/news/industry-developments/youtube-hijacking-a-ripe-ncc-ris-case-study Accessed 21.6.2017.

accidental[3,4] redirection of a significant part of the world's Internet traffic. To give an example of a malicious infrastructure cyber attack, the Domain Name System (DNS) has previously been attacked with Distributed Denial of Service (DDoS) attacks that, had they been successful, would have crippled the Internet on a global scale.

There has been little motivation in the Internet community to fix things that are essentially not broken. While there are problems with the infrastructure, for most parts it is working as intended and such events – regardless of the intention behind them – can be described as isolated incidents. This status quo was challenged in June 2013, when the security, integrity and dependability of the Internet were cast in a controversial light due to a major data breach. The NSA and its associate intelligence agencies were revealed to have disturbingly wide signal intelligence capabilities, and that they were indeed actively used against – allegedly – both legitimate and illegitimate targets. The news reports were sourced by a former NSA contractor, Edward Snowden, who had collected a vast number of classified NSA documents and subsequently leaked[5] them to the public. The documents outline in detail the surveillance capabilities, strategies and procedures of the NSA. They also present an unprecedented view of the potential abuses that happen within the purview of the mass surveillance program. While many of these capabilities and programs were already alleged to exist, the leak provided concrete proof for the existence of a mass surveillance project of previously unseen proportions.

In this section we do not aim to present a comprehensive view of the full capabilities of the Five Eyes surveillance network. Five Eyes refers to the Signals Intelligence (SIGINT) establishments of the United States, United Kingdom and the other Commonwealth states: Canada, Australia and New Zealand [66]. First, we must recognize the nature of our source information. As all information about these programs and capabilities are classified

3 Renesys. "Con-Ed Steals the 'Net". Online, available at http://www.renesys.com/2006/01/coned-steals-the-net/ Accessed 21.6.2017.

4 Renesys. "Internet-wide catastrophe — Last year" Online, available at http://www.renesys.com/2005/12/internetwide-nearcatastrophela/ Accessed 21.6.2017.

5 Electronic Frontier Foundation. NSA Primary Sources. Online, available at https://www.eff.org/nsa-spying/nsadocs Accessed 21.6.2017.

as top secret and are in the public eye only because they have been leaked. Thus we cannot make any convincing arguments and observations on the truth behind these documents. We are rather forced to take them at face value and note, that however well-sourced and probably correct, they are not proper sources and thus we can only speculate on the truth. Second, we do not aim to make an exhaustive list, but rather focus on the main capabilities and programs related to mass Internet surveillance. By using this focus, we can better discuss the effects the programs have on the Internet, and also compare these capabilities with other actors with similar interests in mass surveillance, such as China and Russia. For an overview of the different programs and related tools, see e.g. [67, 68] for further discussion of several relevant programs.

3.1.2 *Information gathering programs*

The main information gathering program for the NSA is alleged to be the PRISM project.[6,7] It is the codename for the program under which data is siphoned from the networks of large Internet technology companies.

This was one of the first major revelations in the Snowden leaks. Several US Internet companies were implicated in providing data on their customers to the NSA. It was also implied that some companies were willing partners in giving data of their customers to the authorities, and other companies were included in the program without their knowledge. In a document dated in April 2013 describing the PRISM program, major US companies including Microsoft, Apple, Yahoo, Facebook and Google

6 Glenn Greenwald and Ewen MacAskill, *The Guardian*, 7.6.2013, "NSA Prism program taps in to user data of Apple, Google and others", http://www.theguardian.com/world/2013/jun/06/us-tech-giants-nsa-data (accessed 23.9.2015)

7 Barton Gellman and Laura Poitras, *The Washington Post*, 6.6.2013, "US, British intelligence mining data from nine US Internet companies in broad secret program", http://www.washingtonpost.com/investigations/us-intelligence-mining-data-from-nine-us-internet-compa-nies-in-broad-secret-program/2013/06/06/3a0c0da8-cebf-11e2-8845-d970ccb04497_story.html (accessed 23.9.2015)

were named as information providers to the PRISM program.[8] The response from the named companies was to generally deny any involvement.[9,10] Additional insight to the actual method on how data from Google is collected was provided in a later document, which described the process for tapping Google internal networks for data extraction when it is on the move within the internal network.[11]

The United Kingdom Government Communications Headquarters (GCHQ) has a similar information gathering program, codenamed TEMPORA.[12] In a similar manner to the PRISM program, TEMPORA intercepts communications by tapping the Internet backbone fiber optic links, allegedly with approval from the companies that operate those links. Together, PRISM and TEMPORA are responsible for a significant part of the bulk data collection efforts for the Five Eyes surveillance network. Other data collection programs are alleged to exist, such as the UPSTREAM collection,[13] but examining these programs any further is outside the scope of this thesis.

8 "NSA slides explain the PRISM data-collection program", *The Washington Post*, June 6th 2013. Online, available at http://www.washingtonpost.com/wp-srv/special/politics/prism-collection-documents/ Accessed 16.10.2017.

9 Larry Page and David Drummond, *Official Google Blog*, June 7th 2013. Online, available at http://googleblog.blogspot.fi/2013/06/what.html Accessed 16.10.2017.

10 Mark Zuckerberg, 2013. Online, available at https://www.facebook.com/zuck/posts/10100828955847631 Accessed 16.10.2017.

11 Barton Gellman and Ashkan Soltani, *The Washington Post*, October 30th 2013. "NSA infiltrates links to Yahoo, Google data centers worldwide, Snowden documents say" Online, available at http://www.washingtonpost.com/world/national-security/nsa-infiltrates-links-to-yahoo-google-data-centers-worldwide-snowden-documents-say/2013/10/30/e51d661e-4166-11e3-8b74-d89d714ca4dd_story.html Accessed 16.10.2017.

12 James Ball, *The Guardian*, 25.10.2015, "Leaked memos reveal GCHQ efforts to keep mass surveillance secret" Online, available at http://www.theguardian.com/uk-news/2013/oct/25/leaked-memos-gchq-mass-surveillance-secret-snowden Accessed 23.9.2015.

13 United States Foreign Intelligence Surveillance Court (FISC) order. Online, available at https://www.aclu.org/files/assets/fisc_opinion_10.3.2011.pdf Accessed 16.10.2017.

3.2 NETWORK SURVEILLANCE IN FINLAND AND SWEDEN

In this section we briefly outline the status of network surveillance in the Nordic countries. While the focus in this thesis is in the kind of surveillance NSA and other comparable actors are capable of, the situation in the Nordic countries is also of great concern, as they are often considered to be progressive and liberal societies in many ways, including privacy and safety of their citizens. This progressiveness does not, unfortunately, extend to network surveillance, as we shall next observe.

Network surveillance in Nordic countries is primarily conducted by national signals intelligence SIGINT establishments. In some cases there are jurisdictional divides between civilian and military organizations, and some countries have more than one agency authorized to conduct SIGINT operations. The instance responsible for SIGINT in Finland is the Finnish Defence Intelligence Agency (Finnish: Puolustusvoimien tiedustelulaitos). The Finnish Security Intelligence Service Suojelupoliisi (SuPo) does not conduct signals intelligence, but legislation to clarify the legal framework around network surveillance and to grant SIGINT capabilities also to SuPo is under preparation.[14]

The government organization responsible for SIGINT in Sweden is Försvarets Radioanstalt (FRA). Previously FRA was limited to intercepting wireless communications only, but on 1.1.2009 they gained legal authorization from the Swedish parliament to intercept wired communications that crosses Swedish borders. These broad rights were consequently restricted by an amendment to the relevant legislation, significantly restricting the operational parameters of FRA surveillance.[15]

The Internet backbone networks in the Nordic area go mostly through Sweden. They are operated by Telia Carrier, which is the

14 Suomen Eduskunta, "Verkkotiedustelu". Online, available at https://www.eduskunta.fi/FI/tietoaeduskunnasta/kirjasto/aineistot/kotimainen_oikeus/LATI/Sivut/Verkkotiedustelu.aspx Accessed 16.10.2017.

15 Sveriges Riksdag, "Lag (2008:717) om signalspaning i försvarsunderrättelseverksamhet" Online, available at http://www.riksdagen.se/sv/Dokument-Lagar/Lagar/Svenskforfattningssamling/sfs_sfs-2008-717/ Accessed 24.9.2017.

second largest ISP in the world[16] and also a Tier 1 network operator. Russia also has a significant part of its connections routed through Nordic countries, making them a prime target for SIGINT efforts for those interested in Russian network traffic.

3.2.1 *Future of network surveillance in Finland*

The Finnish authorities do not yet possess the legal capacity to perform signals intelligence gathering on computer networks, but a proposal for new legislation that would make this possible has been under discussion in the Finnish parliament. The discussion is mostly focused on who will have the authority on surveillance issues, and what will be the division of powers between the military signals intelligence and regular police data gathering and surveillance. Constitutional issues are also in the way of concentrated mass surveillance, as the inviolability of secrecy of communications is quite strongly embedded in the Finnish constitution.

This state of affairs will probably change, as the pressure to engage in mass network surveillance is noticeably increasing in the current political climate. This pressure can be seen to emanate from several sources, such as the international intelligence community and also from political maneuvers and an expectation of increased police powers bringing better results in solving domestic crimes, for example. The pressure for a new foreign intelligence act has increased in Finland in the wake of successful terrorist attacks in Europe in 2016 and 2017.[17]

3.3 COMMERCIAL INTERNET SURVEILLANCE

Internet surveillance in general is not solely in the hands of governments. Tracking, categorizing and assessing Internet users for

16 Earl Zmijewski. "A Baker's Dozen, 2016 Edition" Online, available at https://dyn.com/blog/a-bakers-dozen-2016-edition/ Accessed 24.9.2017.

17 *Kaleva*, 09.04.2017. "Pääministeriltä kova uutispommi: Juha Sipilä haluaa tiedustelulakien perustuslakimuutoksen läpi kiireellisenä" Online, available at http://www.kaleva.fi/uutiset/kotimaa/paaministerilta-kova-uutispommi-juha-sipila-haluaa-tiedustelulakien-perustuslakimuutoksen-lapi-kiireellisena/756788/ Accessed 16.10.2017.

commercial purposes is a key driving force for online commerce. Analyzing potential buyer profiles, targeted advertising and marketing are essential for all business online, and a key requirement for realizing this is the capability to identify Internet users and differentiate them into groups based on whatever relevant criteria, and even to provide personalized advertising. Internet giants such as Google and Amazon focus heavily on analyzing their users for sales purposes.

In a completely different scope but still related to network surveillance, enforcement of Intellectual Property (IP) rights can be argued to currently be one of the forms of commercial Internet surveillance. Commercial monitoring for P2P traffic is a large business for companies specializing in IP rights enforcement. Surveillance is performed by joining a P2P network and pretending to share copyrighted content with others in the peer group. All connecting peers are logged and the information is stored. This monitoring information is later sold to other companies that in turn take action towards those who they have deemed to be in violation of copyright. Recent examples of this include a case in Finland, where a person was convicted to pay a 32 000 Euro reparation fee for copyright violation.[18] This development has caused even concern in the responsible ministry.[19]

While further examination of commercial and private network surveillance would perhaps be warranted in order to form a holistic view of Internet surveillance, such actions are outside the scope of this thesis. We focus on surveillance by government actors and equivalent organizations, and while they can have some overlap (especially in the US intelligence market), examination of this research direction is left for future work.

3.4 METHODS FOR MASS SURVEILLANCE ON THE INTERNET

In this section we examine different methods that have been efficiently used to implement and realize mass surveillance on the Internet. This listing is far from exhaustive, as we do not aim to

18 Eero Mäntymaa. *"Waretukselle" korkea hintalappu: Mies tuomittiin luvattomasta elokuvien jakamisesta 32 000 euron korvauksiin* YLE Uutiset 24.11.2016. http://yle.fi/uutiset/3-9002115

19 *Helsingin Sanomat*, 20.1.2017. http://www.hs.fi/talous/art-2000005053537.html

identify each individual technique and method that can be used fro surveillance. Rather, we focus on methods that compromise the integrity of the infrastructure of the Internet and attack fundamental building blocks of online security: bulk data interception by leveraging network infrastructure, compromising device integrity in various ways, and attacking cryptography standards.

3.4.1 *Targeting various data types*

Before we move on to technical methods for data acquisition, we will first define what different forms data — the target of mass surveillance — can take. Then we move on to discuss various methods used to target mass surveillance to particular kinds of data, how successful these methods have been, and what consequences these methods have had — or can have in the future. Data can essentially be in three distinct states: *at rest, in transit, or in use.* Next we shall examine the features of the three states.

Data at rest refers to data which is stored and is not actively used or transmitted elsewhere. Files stored on a hard drive, flash drive or cloud service that are not actively used are data at rest. *Data in transit* is data that is actively being transferred from one storage location to another. This can be done over a network or even more mundane ways, such as by sneakernet.[20] *Data in use* is data that is actively being used in a computational process. It is stored in volatile memory or can be considered to be otherwise temporarily stored, such as in a temporary file, folder, or swap file.

Data stored on company databases is, by definition, data at rest. It is not possible to intercept data that is not in transit, so one important aspect of mass surveillance is to acquire access to this category of data as well. For example, social media companies generate a lot of data on their users based on their use of the service, and store this data in their systems. In many cases, this data is not ever transmitted outside the company secure networks, and therefore it is outside the scope of regular traffic interception. Against this background, it is clear that gaining access to this data is a priority to any intelligence organization willing

20 The term originates from physically moving storage medium by hand (and foot), thus the transport medium being literally a sneaker – a type of footwear.

to preform mass surveillance, as it provides huge amounts of user data and metadata that is seldom on the move.

In the context of data in transit, the norm should be that communications over untrusted networks are secured with encryption. The Internet, by definition, is an untrusted network. Therefore practically all sensitive communications are transmitted over secure connections, whether the data in question is online banking data, company internal communications, or discussion forum login information. The proliferation of encryption has made using secure connections the standard for client-server communications on the Internet. Modern encryption standards are designed to withstand attacks from practically omnipotent adversaries capable of using practically unlimited resources both in money and computation cycles. The approach of a real-world adversary such as the NSA in this case is to not play fair, but to change the rules fundamentally in their own favor.

Data in use is actively processed is a computational device. Gaining access to data in use requires at least some kind of control over the computational platform, and the process that is processing the data. This can be a difficult problem for surveillance purposes, as this effectively requires a targeted surveillance effort to gain access to a particular device and process. A clever adversary can attempt to only process incriminating evidence as data in use, never transferring it over a communication link or storing it in a non-volatile storage device. This can be an effective method for avoiding scrutiny by law enforcement agencies, but is also extremely difficult to do in practice.

Lyon [69] identifies three layers of the "surveillance iceberg" that describes mass surveillance: accessing data in transit, accessing stored data, and using spyware to compromise individual devices. We can identify similar layers to mass surveillance, and we shall next examine some of these layers and their aspects in more detail.

3.4.2 *Bulk data interception*

Security and privacy concerns have driven secure solutions for network communications to become more and more prevalent, but the default assumption of openness can be seen in the nature of existing protocols. HTTP and FTP traffic, for example, are

by default sent in the clear, without any privacy or security. In wireless networks, all receivers within range of the transmitter are able to listen in on the traffic sent by any other host; the fact that most wireless local area network interfaces on the chipset level choose by default to ignore packets not meant for them can be considered to be an act of politeness, not a feature that brings robust privacy or security to wireless networking.

Data in transit can be observed or intercepted as it is transmitted, using whatever communications system available. In our context, this is the Internet, but in principle it does not matter what kind of medium or encoding is used. These certainly have an effect on potential forms of surveillance, though. Computer networks are relatively well-organized by necessity, and this suggests that networks are also well monitored and administered. An administrator of a network is able to perform surveillance on their own network as a by-product of normal administration and troubleshooting duties. Essentially, no traffic is safe from the view of the administrator or operator of a network. This makes the service provider a single point of failure regarding security and privacy, and also a very tempting target for infiltration or some other kind of clandestine information gathering.

Anyone who has worked in network administration should be choking on their coffee right about now...

3.4.2.1 *Data acquisition from other organizations*

One of the key controversies in the NSA surveillance and spying revelations is that NSA has been targeting major Internet companies for data acquisition. A slide in a presentation that was classified as top secret describes the information gathering capabilities of NSA regarding major US Internet companies. Allegedly most major companies have been co-operating with NSA to provide access to their customer data stored on their own servers. This data has been subsequently fed to data analysis, thus making it possible to form even a more accurate and extensive coverage of all Internet users. Even people who do not use the Internet at all are not beyond being profiled if they have friends or acquaintances who do. Users can add information on their location and people they are with to social media sites, for example, and this data can contain personal information of those who do not use the Internet.

Of particular interest in this effort have been social media companies such as Facebook. This is rather understandable, as social media by its nature shows connections between people, and this is a key interest in all intelligence operations: who knows whom is a very important piece of information when attempting to create a comprehensive picture of a surveillance target. As it was noted previously with metadata, mere information of the nature of communications and participants is more than often enough to deduce very accurate information about someone, what they are doing and what they are interested in.

The legislation that governs data acquisition by government agencies in the US actually gives law enforcement significant access to any organization that has operations in the US. Targeted surveillance is possible using established information acquisition procedures. The FISC, a court that operates in a shroud of secrecy, presides on cases relating to national security, and gives relevant security organizations authorization to perform surveillance on targets and acquire data relating to that subject from other organizations. The FISC is authorized by the 1978 Foreign Intelligence Surveillance Act (FISA), and is commonly also referred to as the FISA court.

A significant part of the criticism to this process stems from the apparent lack of checks and balances in the operation of the FISC. While on paper it does have oversight, in truth it seldom rejects any application made to it[21]. The validity of the FISC oversight has been both affirmed and brought into question.

3.4.2.2 *Wiretapping major networks*

To gather massive amounts of Internet traffic, a natural target for data interception are important network junctions that handle significant amounts of traffic. The United States has a large number of important hubs for Internet traffic, as many of the major technology companies, service providers and organizations are based in the US. Also, due to the global network infrastructure, the US is a large hub of Internet traffic, and a significant part of global traffic is routed through the US. This provides ample opportunity for US intelligence to tap important traffic hubs and

21 Electronic Privacy Information Center EPIC, 2015. https://epic.org/privacy/surveillance/fisa/stats/default.html

extract huge amounts of traffic. The NSA has taken this into account and actively exploit this fact in their information gathering operations.[22]

Physically, the interception of network traffic is done with taps on fiber-optic cables in important junctions. An attempt to map such NSA splitter sites in the US is made in [67], where a crowd-sourced mapping tool, IXmaps,[23] is used to identify potential interception sites by tracking Internet traffic as it is routed through the US. Analysis of this data heavily implies that the US has the potential to use its status as a nexus and a preferred route for Internet traffic to its advantage, as several sites for traffic interception are hypothesized to exist.

3.4.2.3 *Deep Packet Inspection*

Network traffic can be examined in various different ways. *Packet switching* networks relay traffic in the form of *packets* that are comprised of a header and data payload. The packet header gives information on where the packet is coming from, where it is going, and of what type it is, among other information. The actual message being transported is contained in the data payload part of the message. Any further information about the contents and context of the message requires closer inspection of the payload. While it is possible to gather significant information from metadata, situations exist when knowledge about message context and meaning is essential. Examples include a wide range of use cases ranging from basic network administration to network surveillance.

Deep Packet Inspection (DPI) is a common name for *application layer traffic inspection*. Network communications are abstracted into different layers, each serving their own purpose and providing services to the other layers. The layer model used in Internet communications is the *TCP/IP model* or *Internet model*.[24] There are four layers in the TCP/IP model: application, transport, Internet and link layers. Each layer has their own respon-

22 *The Washington Post*, June 6th 2013. http://www.washingtonpost.com/wp-srv/special/politics/prism-collection-documents/

23 http://ixmaps.ca/

24 Internet Engineering Task Force. RFC 1122: Requirements for Internet Hosts – Communication Layers. Available online at https://tools.ietf.org/html/rfc1122 Accessed 7.2.2017.

sibilities, and the scope of the application layer is handling and managing user application data.

The goal of DPI is to gather information on the actual content of the communications instead of just metadata about the packet, such as where it came from, where it is going and when. By analyzing the actual data payload it is possible to conduct content based surveillance. Intercepting messages sent over the Internet with various messaging applications is one common use of DPI. It is usually performed with specialized network hardware, such as firewalls designed for this purpose.

3.4.3 *Compromising device integrity*

The capability to use vulnerabilities to attack individual hosts is nothing special in itself; the lowliest script kiddie uses the same principles of exploiting vulnerabilities as the elite military infiltration unit when it comes to attacking computers. Although professionals have more options available for related methods – most script kiddies are, to the benefit of all, not generally capable of infiltration, extortion, murder, and other similarly drastic measures – they too must find a vulnerability in the target system and develop the capability to exploit it. Vulnerabilities are used and leveraged to gain unauthorized access to a computational platform, which is in turn used to complete the objectives of the attacker. The motives and goals of attackers vary as much as their background, ranging from curiosity to profit and to waging active cyber war.

We can categorize attacks targeting computational platforms further into attacks where the target is an individual host, device or equivalent non-network hardware, and attacks where the target is a device that processes network traffic. These include routers and other devices that make up a major part of the key infrastructure of the Internet.

One further category would be attacks that do not discriminate between targets, and aim to gather as much computational resources as possible. Botnets – a group of computers and other network connected devices that have been compromised and can be used by the attacker to whatever purpose – are commonly assembled with trying to gather as many devices as possible to the botnet. Spam-attached malware, for example, is a common

tactic for creating and expanding a botnet. The botnet can then be used to mount DDoS attacks, for example.

The Carna botnet [70] is an example on how easy it is to gather a moderate number of hosts to a botnet on the Internet. The Carna botnet spanned 420 000 individual hosts, gathered in a relatively short period of time, mostly by exploiting default root login vulnerabilities on various Internet-connected devices. Instead of malicious purposes, the Carna botnet was used to map the IPv4 address space and to research the security of the Internet on a global scale. The results show that in 2012, hundreds of thousands of Internet-connected devices were vulnerable to attacks implementable by a simple scripting engine and some expertise.

While extremely relevant to network security in general, as tools for mass Internet surveillance, botnets are not as interesting as the other categories of compromising attacks. While there is some potential for surveillance in analyzing and gathering data processed by devices in the botnet, it is arguably more cost-effective to gather mass data from the network infrastructure directly. One could argue that gathering a botnet and then analyzing the data processed in it would be a kind of mass surveillance, but in this dissertation we shall make a clear distinction between botnets and network hardware compromise.

3.4.3.1 *Targeting individual hosts and systems*

When an attacker needs to compromise a single individual target to achieve their objective, it is considered to be a targeted attack. In surveillance context, this is very often targeted surveillance and dataveillance. If the target is a single host, the goal of the surveillance is quite certainly to gather information about a single targeted host or person, or to otherwise extract some particular information. If the purpose of the attack is not to gather data but to destroy it, these kind of operations fall under other categories such as sabotage. Compromising Internet infrastructure, however, and attacking routers and gathering all data they process, however, is clearly in the realm of mass surveillance.

Several ways exist to gain unauthorized access to a computer. The most common method is to leverage existing vulnerabilities in software known to be running on the target host, and to gain

access through this avenue. Another method is to impersonate a legitimate user of the system. This can be done by guessing the password of the real user, stealing the authentication token or some other credential information, and using this to access the target system.

Finding, exploiting, fixing and selling vulnerabilities is a business. The Common Vulnerabilities and Exposures (CVE) database[25] contains a list of existing vulnerabilities in commonly used software. These vulnerabilities have been gathered from the year 1999, and currently the number of CVE IDs is 84245. Thousands of new vulnerabilities are added to the database annually, and these are just the ones that are noticed by people who submit vulnerability reports instead of selling or exploiting found vulnerabilities.

Zero-day vulnerabilities are vulnerabilities that have not been brought to the public attention, but are actively exploited being in the wild by the discoverer. The price of a zero-day vulnerability can be hundreds of thousands USD on the black market, making them especially valued among cyber criminals and organizations that leverage them in their operations. One of the more famous cyber attack tools, the Stuxnet malware, was based on – among other vulnerabilities – exploiting a zero-day vulnerability [71].

3.4.3.2 *Targeting network hardware*

If an attacker is able to gain root privileges or otherwise compromise a computational device in a similar manner, the attacker is also able to observe whatever data is processed in the device. Gaining root access on a workstation or a server by using exploits – previously known or zero-day, it makes no difference as long as they work – provides access to whatever data is stored or processed on that particular host. Compromising the operating system of a key network router, though, gives an attacker an excellent position to gather information and survey network traffic routed by the compromised router. If it is positioned strategically in a network, it can provide access to massive amounts of data routed through it. Such methods have allegedly been used

25 Common Vulnerabilities and Exposures. Online, available at https://cve.mitre.org/index.html Accessed 15.4.2017.

by the NSA, as implicated by leaked documents[26] discussing this practice. While compromising hosts can be considered to be relatively mundane in the current network security climate, targeting network routers is not possible for most attackers, as this requires significant resources.

As a related example, in February 2015 The Intercept published an article[27] based on a GCHQ document[28] where it is insinuated that the NSA and GCHQ infiltrated Gemalto, a large supplier of secure documents and smart cards, and stole the master encryption keys for mobile phone SIM cards. Anyone in possession of these keys can decrypt all phone traffic originating and directed at phones operating with a Gemalto built SIM card. This is a clever attack that gives access to a large bulk of network communications. An unfortunate side-effect with attacks of this nature is that probably there will be several innocent bystanders in the group of compromised SIM card owners.

3.4.4 *Subverting cryptography*

One of the approaches taken by the NSA is to target the cryptography standards defining secure protocols, and also individual implementations of cryptography standards in the form of security products.[29] The details in these leaked documents are redacted or unclear as to which protocols, products or vendors have been compromised, unfortunately. Allegations have naturally been abundantly made, but proving this kind of malicious activity is difficult at best. Next we shall briefly examine a few examples of previous tampering with cryptography standards and products. This is by no means an exhaustive list of such cases, but aptly demonstrated the extent to which subversion of standards can be taken. As modern cryptography is perhaps the

26 Sean Gallagher, *Ars Technica*, 12.3.2014. "Photos of an NSA "upgrade" factory show Cisco router getting implant" http://arstechnica.com/tech-policy/2014/05/photos-of-an-nsa-upgrade-factory-show-cisco-router-getting-implant/

27 Jeremy Scahill and Josh Begley, *The Intercept*, 19.2.2015, https://firstlook.org/theintercept/2015/02/19/great-sim-heist/

28 https://firstlook.org/theintercept/document/2015/02/19/cne-access-core-mobile-networks-2/

29 *The Guardian*, September 5th 2013. http://www.theguardian.com/world/2013/sep/05/nsa-gchq-encryption-codes-security

most robust part of security, it is not surprising that intelligence agencies have concentrated significant resources for subverting cryptography solutions available to the public.

3.4.4.1 *Crypto wars*

Intentionally weakening cryptography solutions has ample precedent in the US. In what has become to be known as the *Crypto wars*,[30] the US and other governments have waged war on strong encryption, seeing it as a threat to national security interests. The chosen method for fighting strong cryptography has been censorship. After the second World War the US classified strong cryptography as a controlled export item comparable to weapons, meaning that any US company manufacturing software with cryptography elements in it must provide foreign customers a version with reduced security. This can be achieved with shorter key lengths or known vulnerabilities in the algorithm, for example.

3.4.4.2 *Data Encryption Standard*

The Data Encryption Standard (DES) [72] was a long-time *de facto* encryption standard for information processing. It was created by IBM in the early 1970's and standardized in 1977. It is a Feistel network block cipher with a key size of 54 bits – woefully short even back in the day – and some unconventional design solutions. Especially the choice of the S-boxes in the cipher have provoked doubts of deliberate weaknesses by design in DES [73]. The EFF built a DES cracker in 1998 for around $250 000. It was an ASIC designed to break DES keys, and was capable of decrypting a DES encrypted message by exhaustive keyspace search in a matter of hours [74]. If the assumptions on deliberate weaknesses in DES design are true, those with knowledge about the inner workings of the algorithm could have done comparable attacks on DES encrypted data for a fraction of the cost, way earlier than brute force attacks became feasible in 1998. DES was finally deprecated by the adoption and standardization of the Advanced Encryption Standard (AES) in 2002 [75].

30 *The Electronic Frontier Foundation*, "The Crypto Wars: Governments Working to Undermine Encryption", https://www.eff.org/document/crypto-wars-governments-working-undermine-encryption

3.4.4.3 *Key escrow*

One approach to subverting cryptography is to use *key escrow*, where a "trusted" third party keeps the used encryption keys in escrow, only providing them to law enforcement agencies when required to do so. During the previous crypto war in the 1990's, the Clipper chip was the US government proposal for encryption of digital communications. It faced widespread resistance due to legal and technical issues, and as a result the scheme was never widely implemented. It was ultimately scrapped in silence. For extensive discussion and details on Clipper and the controversy it sparked, see Froomkin's comprehensive paper on the subject [76].

Recent political discussion in the United States in 2016 has brought the use of key escrow back to the general discussion. Calls for designing systems with a "golden key" or a "skeleton key" that would allow law enforcement to decrypt communications when necessary have been suggested by multiple parties. This is comparable to key escrow, and contains all of its disadvantages with no upsides. It will remain to be seen whether backdoored encryption standards and products will be publicly announced, but as the software industry is extremely hostile to the idea it is quite probable – fortunately – that such schemes will not be publicly adopted.

3.4.4.4 *Predictable randomness – The case of Dual_EC_DRBG*

The Dual_EC_DRBG Random Number Generator (RNG) is an example of subversion and intentional weakening of cryptographic protocols. It is a RNG designed and certified by the National Institute of Standards and Technology (NIST), originally designed in the early 2000's. In 2007, concerns were raised about the security of the RNG.[31] No definite proof of foul play could be found until the Snowden archive of documents revealed more information on a concentrated effort to weaken cryptography standards globally.

It turned out that the Dual_EC_DRBG RNG was included in several commercial cryptography applications. Most notably, the

31 Dan Shumow and Niels Ferguson, 21.8.2007, "On the Possibility of a Back Door in the NIST SP800-90 Dual_EC_PRNG", http://rump2007.cr.yp.to/15-shumow.pdf

famous security company RSA used it as the default source of randomness in their BSafe product after striking a deal with the NSA.[32] This enabled the NSA to easily break the security of anyone using that particular software. Bernstein *et al.* [77] provide a good account of the Dual_EC_DBRG story.

3.4.4.5 *Attacking RC4 in TLS*

The RC4 encryption algorithm (see, cf., [73] for a description of the algorithm) is a stream cipher that is very commonly used in many applications. It was developed by Ron Rivest in 1987, and due to its inherent properties — lightweight, simple, and efficient — it has been used in nearly all relevant Internet standards; Among others, Secure Socket Layer (SSL), Transport Layer Security (TLS), BitTorrent, Kerberos, Secure Shell (SSH), and both Wired Equivalent Privacy (WEP) and Wifi Protected Access (WPA) all support the RC4 encryption algorithm.

RC4 has several known weaknesses, and some of them have already been known for the last 20 years. Among the first were the existence of a class of weak keys found in 1995 by Roos [78], and a weakness in the algorithm's key generation algorithm, found in 1997 by Golić [79]. More serious attacks have been subsequently found. There is a widely studied statistical weakness in the key generation algorithm, which causes the output to be predictable [80, 81, 82, 83, 84]. One proposed solution to avoid this was to dump at least the first 512 bytes of keystream [85]. Even this has been found to be insufficient, as it is possible to mount an attack on RC4 that takes only a matter of hours [86]. This is done by exploiting long-term biases that can be found in the keystream. In contrast to short-term biases which can be alleviated by discarding some amount of the initial keystream, long-term biases appear throughout the whole keystream and cannot thus be similarly avoided.

RC4 is also subject to a related key vulnerability [80]. This attack requires the encryption key to be split into two parts, a public key k_p and a secret key k_s, and that k_s remains the same while k_p varies from message to message. This is the design

32 Jospeh Menn, *Reuters*, 20.12.2013, "Exclusive: Secret contract tied NSA and security industry pioneer", http://www.reuters.com/article/2013/12/20/us-usa-security-rsa-idUSBRE9BJ1C220131220

approach that was taken with WEP. It was instrumental in its deprecation as insecure, as the related key vulnerability works regardless of key length or other modifiers. This combined with efficient and automated methods for exploitation all contributed to WEPs demise [87].

AlFardan *et al.* [88] have examined the use of RC4 in TLS. They found that as RC4 is vulnerable to plethora of attacks, it should be removed from active use in Internet standards as soon as possible. The attack later found by Vanhoef and Piessens [86] further increased the pressure for deprecation of RC4. The Internet Engineering Task Force (IETF) subsequently deprecated RC4 in TLS in 2015,[33] citing as the reason the new and improved attacks against RC4 implied that RC4 is unable to provide sufficient level of security to TLS.

Given that AlFardan *et al.* [88] estimate that about 50% of TLS traffic at the time was secured with RC4, storing traffic for later analysis suddenly becomes quite the fruitful approach into breaking encrypted communications. When data is stored in bulk, it can be accessed later in leisure. Should there be an identified pattern of interest, decrypting RC4 encrypted communications data is not a very difficult task for the NSA or other equivalent actor. After the encryption has been stripped away, this data can then be mined for further intelligence. This results in that communications encrypted with RC4 are perpetually at risk of being compromised at a later date.

Some protocols may allow the use of older, unsafe cipher suites in an effort in backwards compatibility. In some configurations TLS may allow fallback to SSL 3.0 for legacy systems. It is possible to exploit this option to force communications ot take place with unsafe cipher suites and subsequently breakable security. The Padding Oracle On Downgraded Legacy Encryption (POODLE) attack [89] accomplishes exactly this by forcing the use of SSL 3.0 for client-server communications. The fact that communications encrypted by RC4 can be decrypted later given sufficient but reasonable time provides a great motivation for storing massive amounts of communication data for later analysis.

33 Internet Engineering Task Force. RFC 7465: Prohibiting RC4 Cipher Suites. Available online at `https://tools.ietf.org/html/rfc7465` Accessed 7.2.2017.

3.4.4.6 *Attacking Diffie-Hellman in TLS*

Adrian *et al.* [90] describe an attack against TLS based on exploiting a common insecure version of Diffie-Hellman key exchange protocol [91]. *Logjam* is a downgrade attack that forces vulnerable servers to use a weaker version of Diffie-Hellman. This attack made it possible for the authors to compromise connections to 7% of top million web sites indexed by Alexa at the time. Logjam is based on precomputation and exploiting weak Diffie-Hellman groups used in export-grade version of the protocol.

The same principles can be applied to compromising Virtual Private Network (VPN) connections. The authors estimate that the NSA has sufficient resources to mount an attack even against 1024-bit Diffie-Hellman used in VPN connections [90, p. 7–8]. Evidence of such compromise can be found in the description of a classified NSA program TURMOIL,[34] leaked by Edward Snowden in 2013.

3.4.4.7 *Attacking the Certificate Authority system*

In cryptography, protocol descriptions often assume the existence of a neutral third party, one that has no interest in the protocol transaction itself and only acts as a third party arbiter in secure communications between two (or more) parties. The neutral arbiter has to be trusted by all participants, even though the parties might not trust each other in the first place at all. This kind of arbiter construct is both common and useful in exactly these kind of situations, where trust is either absent or hard to establish.

The Certificate Authority system is central to modern Internet security. It is a trusted third party based system, such as is described above, where a CA acts as a trusted third party in a transaction between two parties that have no prior knowledge of each other. The CA grants a proof of identity in the form of a X.509 standard[35] digital certificate, which is then used to establish a trusted connection.

34 Der Spiegel. Turmoil VPN Processing. Leaked NSA Document. Available online at http://www.spiegel.de/media/media-35526.pdf Accessed 11.6.2017.

35 Internet Engineering Task Force. RFC 6960: X.509 Internet Public Key Infrastructure Online Certificate Status Protocol - OCSP. Available online at https://tools.ietf.org/html/rfc6960 Accessed 7.2.2017.

The CA system is heavily dependent on the trustworthiness of individual CAs. Should a CA be compromised in any way, the trustworthiness of systems depending on certificates granted by that CA naturally become suspect as well. A relatively famous example of CA compromise is the case of the Dutch CA DigiNotar B.V. [92]. In 2011 the company was subject to an attack that eventually ended up toppling the company itself. The attackers were able to use the CA credentials of DigiNotar to grant themselves authentic certificates for various web services. One of the main targets was Google, as a certificate for **.google.com* was among those falsely issued. These certificates were in turn later used to facilitate a Man-In-The-Middle attack against users, a majority of these users coming from Iran.

The approach of attacking the CA system is logical if the goal is to subvert a secure communication system. Instead of attempting to break the cryptography in the communications directly, it is more cost-effective for malicious attackers to target these neutral third parties, or to establish their own "trusted" CAs. The true extent of malicious CAs in the wild is unknown, but already in 2013 a major spike in binaries signed with a certificate from malicious CAs was noted in a report by McAfee Labs.[36]

3.5 JUSTIFICATION OF SURVEILLANCE

As we have already observed in Chapter 2, surveillance in itself is a commonly used tool, albeit one with some very significant and serious potential drawbacks for both its users and targets. Nevertheless, calls for more surveillance in society for increased security are not uncommon at all. A cynic might even say that people bring the bad effects of mass surveillance upon themselves by demanding more surveillance in the hope of increased security. Potential justifications for mass surveillance are briefly examined in this section, but a detailed analysis of the dynamics of justification of surveillance is out of the scope of this thesis.

Surveillance does have a valid justification in some situations. Few can argue against just and laudable goals, such as the need to fight crime and terrorism, and such are the commonly cited

36 McAfee Labs. Malicious Signed Binaries Crush Certificate Authority Reputation. Available online at https://www.mcafee.com/us/security-awareness/articles/malicious-signed-binaries.aspx Accessed 7.2.2017.

reasons for engaging in surveillance. Another aspect of generally acceptable surveillance is monitoring known dangerous places and massive events with large crowds, for example, so that emergency response is not delayed unduly in the case of an accident or other undesirable event. It is hard to convincingly argue against using surveillance tools in those situations.

What is more interesting in this regard is that many in a given society may even consider surveillance a good thing. This is evidenced by the prevalence of people who espouse the "I have nothing to hide" -argument, previously examined in Section 2.3.3. They are confident that surveillance will target people other than themselves, and even if they should be targeted, nothing bad will happen as they are innocent and indeed "have nothing to hide, and thus nothing to fear."

A central problem in justifying surveillance is where to draw the line between the acceptable and unacceptable. In his book, Rule [93, p. 2] observes a key issue in privacy: "there is no natural line of separation between the realm of the private and personal matters of legitimate interest to others". Indeed, there exists a definite set of situations where surveillance and intrusion of privacy have generally been historically deemed to be acceptable, among them criminal investigation into more serious crimes such as murder, treason, and major financial and narcotics crimes.

The existence of valid reasons for surveillance and even in some cases, mass surveillance, makes it very difficult for a society to limit the use of these powerful tools to the bare minimum of acceptable use cases and scenarios. A comprehensive study of this problem is definitely outside the scope of this thesis. A promising solution that can help to unravel this problem is discussed later in this thesis in Chapter 7.

3.5.1 *Lawful interception*

The concept of *lawful interception* means interception of private messages by law enforcement authorities. This is done acting under the color of law, usually with a court order or a warrant, depending on the jurisdiction. Lawful interception of protected communications is an important part of law enforcement and

criminal justice as a method for obtaining intelligence on unlawful activities.

Lawful interception is generally targeted surveillance. The goal is to obtain particular information, or to follow only certain people, thus inherently narrowing the scope of surveillance to only intended targets. For example in Finland, the permit for telecommunications interception must be granted by a court, and must be very specifically targeted to a certain case, person, and topic.

One postulate, however cynical it may seem, that can be made is that when a tool for targeted surveillance (such as monitoring telecommunications or network traffic or traditional post – all valid examples in this context) becomes useful in existing mass surveillance, lawful interception tends to expand to cover the use of this tool for mass surveillance – regardless of the original intent of the law regarding mass surveillance. This can be speculated to be a kind of function creep, where something that has been found useful in smaller scale is simply taken into use in a larger context. Using data mining for identifying anomalies from Internet traffic requires storing and analyzing massive amounts of data to establish a baseline. While this can be seen as lawful interception, establishing the baseline requires processing of vast amounts of data from individuals that are not being suspected of any crime. Whether processing data from these individuals should be in turn allowed or not is a central issue in expansion of Internet surveillance, and while interesting and important, further discussion on it is also outside the scope of this thesis.

3.5.2 *Building surveillance in the society*

It is possible to create an information society that has built-in surveillance by capturing all the (meta)data that is generated from normal network communications and storing it for future use. This set of information can be both observed in real time and be accessed later at leisure, but both approaches carry with them serious threats to security and privacy of citizens in the society. When we consider surveillance from the point of view of the networked information society, finding a solution to this problem is paramount.

A massive amount of data is generated as a by-product of normal operations in all aspects of society and business. In order to provide services most companies must collect at least some data about their customers. For example, collection of usage patterns and related communication information of a smartphone, telemetry data gathered to assure that the operating system of a computer is functioning correctly, customer information on what was bought, when, and where. Various software vendors have some kind of feedback systems for collecting data on how their product is functioning and to facilitate solving potential errors. All these are a part of providing a service – one that is often explicitly wanted and paid for by the customer.

As it was noted above, managing communications infrastructure and providing communication services generates a lot of personally identifiable data on who owns which device, where it is connected to, and which devices it is used to communicate with. This data is again generated as a by-product of normal operations. There are thus very legitimate reasons for this kind of mass gathering of data, especially for improving user experience and product development in general. The important question is, when does this data gathering actually become surveillance? This naturally depends on what data is actually collected, how and for what purpose it is used, how well it is anonymized, and finally – and perhaps most importantly – in what manner and how securely it is stored. Telecommunications data has been used in the past to perform lawful interception on messages on all organized communication methods, starting from messengers and postal service.

Consider the scenario where all of the information that is generated as a by-product of normal operations of an information society can be accessed, analyzed and used in surveillance activities. Lyon argues that the use of Big Data techniques for analyzing intelligence data is changing surveillance in itself in three ways [68]: by increasing reliance on software, shifting the focus of surveillance to prediction of future rather than analyzing current and past events, and improper adaptation of techniques from other disciplines to surveillance.

Methods for controlling such an institutionalized infrastructure surveillance are rarely – if ever – technical, but rather based on softer controls such as legislation and customs. The main

problem in this case is trust and specifically, lack of trust. How far can we trust that our data that is collected for the purpose of providing a service is not used against us in the future? We shall examine trust and the related trust dynamics issues in Chapter 4.

When mass Internet surveillance and data gathering is combined with modern efficient data analysis techniques, new associations and links between people are created; events and concepts that, however improbable or trivial, can be used to show connections and intent not based in any actual reality [68]. The use of such information in a malicious manner is a serious concern for individual security and privacy. Even if this data is used in good faith to pursue justice and uphold "national security", because of the disconnect between events and intent inferred from data and reality, mistakes are bound to happen. Connections will be made between people and events that have no basis in reality. For example, correlation of location data of people to movements of a target of interest can result in a person being targeted for surveillance just because they are in the same physical location as a person under targeted surveillance.

If an illegal act is erroneously associated with a person, what would the effect on society be if those falsely attributed allegations were aggressively prosecuted to the full extent of the law. For example, terrorism is a naturally a very serious crime for someone to be charged of. In the current societal climate, however, even the mere suspicion ar illusion of association with terrorism can be dangerous. The publicly stated purpose of many mass surveillance programs is to increase national security and combat terrorism, so such crimes are routinely investigated in those programs. Therefore it would stand to reason that the erroneous accusations would be of a more serious nature. As it was noted previously in Section 2.3.1 when discussing the chilling effect, the legal system is not perfect, and even such erroneous accusations may lead to sentences. This would understandably enhance the chilling effect, as people would be more inclined to stay away from any topics of controversy in order to maintain an appearance of blamelessness.

3.5.3 *Built-in surveillance and business*

While the discussion in this section focuses on the United States, the same observations apply to any nation with a strong government and permissive legislation on mass surveillance. One of the main problems for US based tech companies is that they have no means for defending against allegations of cooperation with intelligence agencies. If the US government wants access to data in the possession of a US company, they can issue a national security letter. It compels cooperation under the threat of imprisonment for company executives, and also forces everything related to the letter to be kept secret. The recipient of a national security letter cannot thus divulge to any third party that they have in fact received one. Warrant canaries [94] — public statements on the company web page stating that no warrants have been issued to the company — have been formulated as a defense against this kind of government action, but their effectiveness has yet to been tested in an actual court case, at least in the United States.

The problem with assurances of non-compliance with authorities is that cooperation is practically mandatory. If a company in the US does not want to comply with the authorities, their options are few and far in between. For example, Lavabit – a secure email service provider used by Edward Snowden – was faced with a government request for data on one of its customers email accounts. Rather than comply with the order and break confidentiality of user data, the owner of Lavabit decided to suspend the operation of the company.[37] This was the only way to avoid prosecution and possible imprisonment. Naturally, this option is not available for technology giants such as Google, Microsoft or Apple; they must comply with any and all valid requests for data by the government. Also, as National Security Letters are accompanied by gag orders, it is a crime to even disclose the existence of such an information request, and even prominent Internet companies cannot risk consequences that would arise from doing so.

One key argument that is often raised in the defense of surveillance is that *everyone* does it. Nearly every sovereign nation has

37 Who is Lavabit? Online, accessible at https://lavabit.com/explain-lavabit.html Accessed 9.5.2017.

its own secret service, police or intelligence agency responsible for intelligence gathering and espionage. Nations collect data on other nations, and everyone acknowledges this fact. It is also the prerogative of a sovereign nation to decide whether it wants to engage in espionage.

The US is a key Internet hub with a significant amount of traffic is routed through the US, and the US is also one of the key suppliers of Internet technologies. A lot of data is also stored in data centers, in locations under US jurisdiction. The US government has previously considered any business with operations in the US under its jurisdiction and claiming access to stored data, regardless if the physical storage location is outside the US borders – a stance the US judiciary has disagreed with so far.[38]

3.6 CONCLUSION

We have examined only a limited subset of the possible ways mass Internet surveillance can be implemented in this chapter, but the most important types of harmful surveillance and data gathering methods have been discussed. It is not the purpose of this thesis to offer an exhaustive list on different methods and attack vectors — this would be practically impossible to do. Instead, we have examined a representative collection of different actors, methods and attack vectors that range from targeted attacks on single computers to mass gathering of network traffic, from traditional eavesdropping or metadata gathering to compromising Internet communication standards and basic infrastructure. All of them have one thing in common – they are making the Internet less secure on the infrastructure level.

Mass surveillance through gathering and storing information for later analysis has become the go-to tool in modern surveillance systems. The reasoning behind this development is actually quite sound: because of the complexity of Internet communication systems and the sheer amount of data that is transferred on these networks, combined with the difficulty of pinpointing

38 Orin Kerr. "The surprising implications of the Microsoft/Ireland warrant case" *The Washington Post* Online, available at https://www.washingtonpost.com/news/volokh-conspiracy/wp/2016/11/29/the-surprising-implications-of-the-microsoftireland-warrant-case/?utm_term=.60824c0472b6 Accessed 27.4.2017.

the target of surveillance from other people, it quickly becomes more feasible to just gather everything and perform the analysis later.

The actual surveillance is done on the metadata level. We must separate data gathering from data analysis, in which a massive collection of data is analyzed for information about persons or organizations of interest. This is also where the problems that are the most privacy sensitive arise: the gathering of data and its subsequent consolidation and combination with other records, and then making assumptions based on meanings derived from incomplete or incoherent data.

The surveillance potential of systematic data gathering is in the ability to compound records over time. If every single thing that a person does online is stored permanently, this information can be used at a later date – even years or decades later. Any potential harmful information could be used years later than when it was originally captured.

The role of cryptography – to provide long-term data privacy and secrecy – may not stand the test of time if the attackers are capable of subverting cryptography standards in a similar way that was previously described in this chapter. If an attacker has some inside information on weaknesses of the used cryptography solution, there exists a nontrivial probability over a timespan of years to decades that the attacker will succeed.

There are situations when surveillance is justified, even in free societies. We can easily construct examples where even the most drastic measures can be judged to be acceptable. But due to both intentional and unintentional effects, even justified mass surveillance has significant drawbacks for the average user of the Internet.

4

FORCED TRUST, MISTRUST, AND RELIANCE ON INTERNET INFRASTRUCTURE

Trust, but verify.

-Russian proverb

In this chapter we discuss trust, and especially focus on the concept of *forced trust*. We consider the implications that forced trust has in the context of Critical Governmental Information System (CGIS), and especially in the particular case of the infrastructure of an information society. We start by examining different trust definitions in applicable fields that can be found in literature. We analyze forced trust relationships in essential information society processes, and the apply this analysis to Internet surveillance.

4.1 DEFINING TRUST

The Merriam-Webster Dictionary[1] gives the following definitions for trust:

- 1 A: *"assured reliance on the character, ability, strength, or truth of someone or something"*
- 1 B: *"one in which confidence is placed"*
- 2 A: *"dependence on something future or contingent"*
- 2 B: *"reliance on future payment for property (as merchandise) delivered"*

Trust is an important aspect of both personal and societal interactions. Humans tend to trust one another, and we often have

1 Definition of trust by Merriam-Webster. http://www.merriam-webster.com/dictionary/trust

an intuition about people and their respective trustworthiness. Sometimes it is formed by stereotypes, in other instances from personal experience. We as humans are also able to differentiate other humans based on the extent they can be trusted, and with what things or issued they can be trusted with. These trust levels vary from person to person, as Schneier illustrates quite concisely [95]:

> *"I trust Alice to return a $10 loan but not a $10,000 loan, Bob to return a $10,000 loan but not to babysit an infant, Carol to babysit but not with my house key, Dave with my house key but not my intimate secrets, and Ellen with my intimate secrets but not to return a $10 loan."*

This simple yet powerful example illustrates that clearly, even what one could think of as a simple trust scenario — Do I trust this person, and what can I trust him with? — is actually a complex interconnection of different and possibly conflicting interests. Trust actually has several definitions that are wholly dependent on context of the discipline in which they are used [96]; psychology has a different view on trust as sociology or philosophy, economics has its own approach on trust, and so on.

A simple way of describing a trust relationship would be to say that *"Alice has a reasonable expectation that Bob will behave in a certain manner in a given situation, therefore Alice trusts Bob in this situation."* In a trust relationship, the *trustor* is the party who is placing trust on someone (or something) else, while the *trustee* is the target of this trust. So, in the previous example, Alice is the trustor and Bob is the trustee. A simple way of using this description of trust would be to say that *Alice is assured that Bob will behave in a certain manner in a given situation due to some of his characteristics, therefore Alice trusts Bob.* Here, Alice has no guarantee that Bob will actually do as she thinks he will, but has decided to trust Bob to do as assumed. Some definitions of trust are applicable only in situations where Alice has something to lose, i. e. the situation entails risk to the trustor [96]. The existence of some kind of risk or uncertainty is common in most definitions of trust [97].

4.1.1 *Trust in various disciplines*

We as human beings intuitively *know* what trust means, but formally defining trust is another matter altogether. Metlay [98] compares the difficulty of defining trust to the similar conundrum pertaining to defining what is pornography. As was noted above, the definition of trust varies between disciplines and context. As a consequence the true nature of trust is inherently obscured [96]. This makes it practically impossible to present all possible aspects and definitions of trust across all fields and disciplines, as trust varies according to the lens though which we observe it. As a consequence, in this work we shall focus mainly on trust definitions in the realm of information technology.

Trust is used in various forms and applications in social sciences, economics, computer science and engineering. In the past, trust and its mechanics have been discussed in computer science in the context of e. g. trusted platforms [99], electronic and Internet voting [100], trusting the preservation capabilities of digital storage [101], trust between virtual entities on the Internet [102], and trust in digital commerce [103, 104]. Economists have examined trust extensively, e. g. financial transactions in electronic commerce [105] and organizational trust [106]. Schneier describes trust as an essential part of a functional society, and goes so far as to argue that our society as a whole is completely dependent on the existence of a base level of trust [95, p. 243].

Institutional trust is a relevant trust concept in the scope of this thesis. Metlay [98] studied institutional trust in the context of nuclear waste management and concluded that institutional trust is actually quite simple to define – at least in the context he was observing. Metlay uses two dimensions for institutional trust: beliefs of institutional behavior and the apparent competence of the particular institution. When we are discussing trust on a governmental and societal level, however, this kind of institutional trust model is insufficient in describing the full extent of the trust landscape.

4.2 FORCED TRUST

An important observation in the context of Critical Governmental Information System (CGIS) and trust is the concept of *forced*

trust. It refers to the situation where the user has no choice or opportunity to affect any part of the information system, including the choice to use the system itself. This is a clear case of trust as despair [106], a type of situational trust in which the party in the position of forced trust has no other choice but to trust all aspects of the provided system.

Usually trust relationships between actors in the context of trust between users and information and communication systems are voluntary. In this case, "voluntary" means the trustor has an actual choice in whether to use and trust the system or not. CGISs do form a salient and obvious exception to this assumption, but situations of forced trust can be observed in any organization where an information system is taken into use without careful consideration. It is also true that there will always be critical voices and disagreement on what system or software to use for whatever purpose. This does not detract from the study of forced trust, but rather serves to emphasize the problematic aspects of forced trust, as it exists in the networked information society.

In other words, forced trust describes the situation in which a user is dictated to use and to trust an information system or an ICT product. As the user does not have the privilege to choose, his or her actions and attitudes are affected by the forced trust relationship. A similar situation exists on the other side of the trust relationship as well. The designer of the information system has to take into account the potential misbehavior of users – whether purposeful or not is beside the point – and implement security measures and safeguards against such events. As all information security functions necessarily increase complexity, limit usability and take resources [107], this preparation is not without its consequences.

4.2.1 *Distrust, untrust, and mistrust*

Because of the amorphous nature of trust, various definitions for subclasses of trust exist. Marsh and Dibben discuss three such trust concepts that are central to this thesis: mistrust, distrust and untrust [108]. The first, mistrust, is, simply put, misplaced trust. It describes a situation where trust is first placed and then abused in some manner, leading to a situation of trust becoming

that of mistrust. It is in effect negative trust, a negative characterization of the target of mistrust. It is possible to continue to operate in the zone of mistrust, but probably some other leverage is required than trust to facilitate further cooperation in this case.

The second is distrust. It is a subtly complex type of trust, where the expected outcome of trust is negative. A trustor with distrust expects the other party to actively act against their interests. It must be noted that it is possible to function in an environment of distrust using external (with regard to trust) control mechanisms to ensure compliance and cooperation. These include rules, regulations, legislation and contracts. Marsh and Dibben give an example on the distinction between mistrust and distrust by comparing them with misinformation and disinformation. The former is factually incorrect but possibly not maliciously, but the latter is purposefully incorrect information.

The third type of trust Marsh and Dibben discuss is untrust. It is described as a situation of positive trust, i.e. trust exists, but that trust is insufficient. In a situation of untrust, further action or assurance is required before a state of trust is achieved. Similar mechanisms that can be used for functioning with distrust can be used. The difference between untrust and distrust is that in the former there is some uncertainty towards the trustworthiness of the target of untrust, while in the latter it is clear that the target of distrust is indeed not trustworthy.

4.2.2 *Forced trust and Critical Governmental Information Systems*

In Figure 2, the trust landscape for a Critical Governmental Information System (CGIS) is illustrated in directed multigraph form. Graph vertexes represent actors related to the CGIS. The edges are directed, designated by arrow-heads. A directed edge between two vertices signifies that there is a trust relationship between those two actors, and the direction also signifies the direction of trust. The edges are also weighted with the type of trust. From Figure 2, we can clearly identify the forced trust relationships between users, government, system administrators and information system suppliers. The role of the user is particularly unfortunate in the case of user-supplier trust relationship. As the user has actually very little say into which information

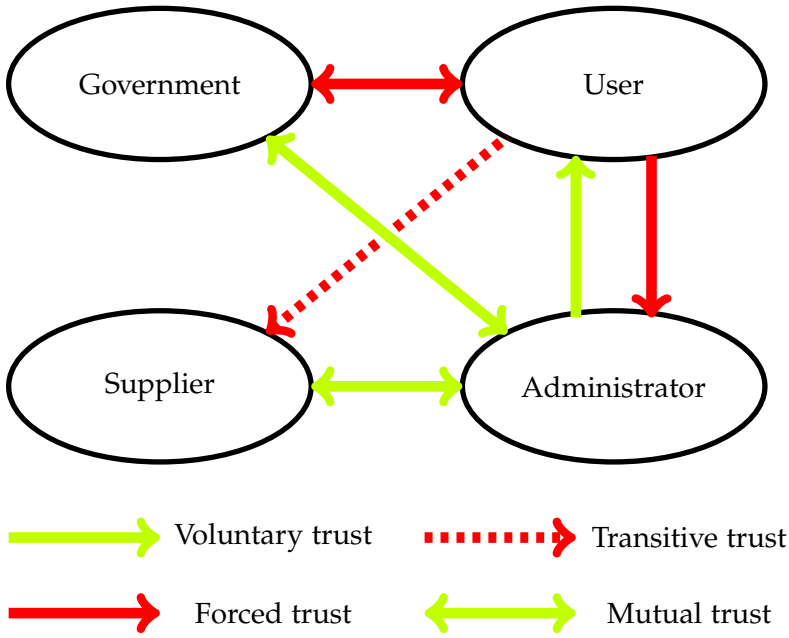

Figure 2: Trust landscape for Critical Governmental Information Systems. Adapted from [8].

system supplier will be designing and implementing the CGIS, there is no other option but to trust the supplier to keep their end of their bargain and to provide a secure and usable system. This trust relationship is indeed both transitive and forced; the user must trust the information system supplier because the government trusts the supplier, and the user is in a forced trust relationship with the government. This latter forced trust relationship is very strong. In the case of disagreement, the recourse available to users is either to vote differently and hope for a change in policy, or to move to another jurisdiction with a different government – if possible.

The government must also trust the supplier, but as an entity with significant advantages over an individual (sovereignty being the foremost), the government has a definite choice on which supplier to select. This decision should obviously be based on trust, but if distrust reigns, control mechanisms must take a major role in the relationship. Of course situations exist where the supplier has the significant upper hand. Supplier lock-in is a prime example: an information system supplier is able to make

itself irreplaceable and making it impossible or highly impractical for a government to choose another IS supplier. Natural monopolies in the software market can lead to this situation. It could be argued that from the point of view of the government this is a clear state of mistrust, but further expanding this aspect of the trust landscape is outside the scope of this thesis. We are more interested in the other cases in our example, where forced trust is due to more serious reasons than impracticality.

Users have also no choice but to trust the administrators (or "admins") of the information system. Admins are people with the capability to observe and affect the operation of the information system with the goal of maintaining proper system operation. In most cases the users of an information system do not interact with the admins, unless there is an anomaly that is visible to users and also requires administrative attention. The admins' job is to keep an information system running and functioning correctly, but they can also cause significant harm to a system or to an individual user through their actions, whether unintended or deliberate.

This trust-based issue can be mitigated by *watching the watchmen*.[2] By auditing the administrators' actions within the information system and generating a Panopticon-like control mechanism for administrators it is possible to affect the behavior of the administrators. As was observed already in Chapter 2, the Panopticon model can be used to modify behavior. It was also strongly argued against using the Panopticon model due to this exact nature – it is an interesting discussion whether using such methods for controlling the actions of a group limited in scope is more acceptable. It can be argued that the administrators are in a position of power, and thus can be subjected to more stricter means of control than regular users. This avenue of inquiry is not pursued further in this thesis, however.

A major problem with users and information systems is the propensity of users to choose bad passwords [109]. Users will go to great lengths of trouble to be lazy: for example users can develop ways to circumvent password policies to ease their own. As the password is the main authentication method for the majority of services, the persistent use of bad passwords is a major

2 "Quis custodiet ipsos custodes?" (tr. "Who will watch the watchmen?") A phrase commonly attributed to the Roman poet Juvenal.

problem for any administrator of an information system. It is also a major breach of trust in the case of a successful break-in by criminals using guessed credentials from legitimate users. Password audit tools are common IT administration tools for this reason. Passwords and insecure user tendencies are examined in more detail in Chapter 7.

In some cases, administrators can choose their customers and users, but in the case of CGISs, this is not generally possible. The administrators simply have to cope with the users behaving in an irresponsible manner – in other words, the administrators are in a forced trust relationship with the users. Especially with CGISs that are so important that the right to use them is a legally guaranteed right, it is impossible to ban users from using the system. For example, eVoting systems are in this category; the right to vote is central to democracy, and it would certainly be an interesting discussion whether you could ban people from voting based on irresponsible behavior. If even the act of killing another person does not revoke your right to vote — at least in civilized jurisdictions — what could possibly justify banning users from this system over bad passwords?

4.2.3 *Response to forced trust*

Three different actor reaction patterns to forced trust can be identified: *acceptance*, *avoidance*, and *resistance*. These patterns subsequently have several potential behavior options that can vary between actors. All possible options are not discussed further, but a selected few are presented as examples.

ACCEPTANCE. In this case, an actor chooses to accept the situation and decides to use the system. This means that even when the actor is in the position of forced trust, they still choose to use the system with forced trust. An actor can be either aware of all the ramifications posed by forced trust and still decide to use the system (i. e. "I know; I do not care" attitude). Even though an actor might find the ramifications posed by forced trust to be inconvenient, they might still end up using the provided system. This can be described as the "I'm annoyed, but not enough" attitude. These two reactions can be considered to be desirable outcomes from the point of view of the system.

AVOIDANCE. In this case an actor decides to not use the system in its full capacity, or even to avoid using it altogether. This can mean various different user behaviors, such as not giving all requested information upon request to the system, or choosing to circumvent the system in some manner. This is an undesirable outcome, as the system is not used to its full capacity, methods for parallel functionality are explored by actors, or the system is not used at all.

RESISTANCE. Finally, in this case an actor decides to actively fight against another stakeholder, or the system itself. This resistance can manifest in behavior ranging from subtly malicious such as providing false data to outright sabotage. This is obviously the least desirable outcome from the point of view of all stakeholders.

Now let us consider the trust classes described in Section 4.2.1 in a situation of forced trust in the context of information systems. In this case, it is fruitful to observe the situation when we choose not to cooperate. What are our potential choices in this situation, and how much are they choices in the first place? Examples of information systems dealing with and measuring employee performance leading into negative outcomes such as firings have been previously observed by Zuboff [58], among the first. There is a clear forced trust aspect in such systems, for example, and the user response to forced trust will depend on the type of trust.

In the case of forced untrust, the user response is hard to measure. As there is positive but insufficient trust, all responses are viable. Some may trust the system and cooperate, some may not trust the system and still cooperate because they are beneath their "bothering" threshold, and some may choose apathy or even adverse actions, but they are probably less common. Forced untrust and CGISs is perhaps the least worrisome situation, as new systems start out in this state; trust is yet to be established but the expectation is positive. Not having any trusted options can be a disadvantage, for example using a new eVoting system, with no option for paper ballots, is a situation of forced untrust.

Mistrust is more complex, and more prone to non-cooperation. Forced mistrust is a bad situation to be in; we have already trusted the other party, in this case the information system, its admins, or its suppliers, and have been betrayed in some man-

ner. Unfortunately, we are in no position to choose another option, and our expectations of future transactions are probably negatively affected. In the information system context, this can be conceptualized as using an information system that based on previous experience has delivered bad results. Avoidance may be a prevalent user response in this case, with less accepting users and an increase in actively fighting users. A user with a forced mistrust relationship to a CGIS will find their options in society reduced. To continue with the eVoting example, forced mistrust would describe a situation where an eVoting system is used even after negative results, and no other alternative is given.

Distrust is the worst type of trust discussed here, and forced distrust is a bad situation. In this case, a significant part of the users will probably actively fight against the system, making it a difficult, even hostile environment to function in. This is the situation where a user is forced to use a system that they are certain will not function properly, will cause detrimental effects to its users, or will in other ways betray any trust and expectations. To continue with the voting example, forced distrust would be using an eVoting system known to be rigged, and being provided no alternatives. In this case acts of minor disobedience and defection are expected, and full-blown sabotage is within the realm of possibility.

4.3 FORCED TRUST IN THE NETWORKED INFORMATION SOCIETY

In this section the concept of forced trust is applied in a wider scope to the networked information society. The central observation is that the citizens of a networked information society are forced to trust an inherently vulnerable infrastructure that has not been built to be secure in the first place. This forced trust can be exploited for purposes of control, surveillance and other undesirable actions.

The dilemma of forced trust in information systems in general, and CGIS's in particular, can be phrased in the context of global communication infrastructure. *Instead of questioning the trustworthiness of individual information systems it must be questioned whether the infrastructure of the networked information society itself is trustworthy.* Based on previous research and observations

in this thesis, trust in the infrastructure is definitely misplaced. Due to the nature of the networked information society, Internet users are in the unfortunate position of forced trust: participating in society in any meaningful manner requires accepting the flawed nature of the infrastructure – and then using it regardless. *The Internet is broken, but we as a society have already ceased to be able to live and function without it.*

Mass surveillance of Internet users is a special case of exploitation of this forced trust. As it was observed in Chapter 3, the United States National Security Agency (NSA) and other equivalent agencies are engaging in mass surveillance by collecting extensive amounts of data of Internet users and using it for intelligence gathering purposes. This is all clearly in the charter of an intelligence agency – it would be a peculiar intelligence agency that would *not* spy on others – and thus should be seen as, if not laudable, then merely acceptable. The problem with this lies in the problem of forced trust and deliberate exploitation of this forced trust. Internet users should be considered to be in various states of forced trust – mistrust, distrust, untrust.

The techniques and methods used to implement this data collection, however, deliberately and even carelessly exploit flaws in the basic infrastructure of the Internet, further compromising an already insecure communications network. Examples of methods and techniques used for this purpose were also discussed in Chapter 3. In reality, there is no such thing as an exploit that can only be used by the "good guys". A weakness or flaw can be exploited by anyone aware of it, and trusting the security of a system on the obliviousness of the opponent — security by obscurity -– is an information security disaster waiting to happen.

If this development continues, it will continue to make the Internet less secure for everyone, not just those who are supposed to be targeted by mass surveillance. Ironically, those who are the claimed real targets of these surveillance measures – terrorists and members of major criminal organizations – are often resourceful enough to create their own security solutions. Well-implemented strong encryption is still capable of thwarting all known cryptanalysis methods, and the information required to implement this is freely available – it is impossible to ban the concepts of algebra and number theory from common knowledge.

The Internet in itself is designed to withstand a nuclear conflict, and as a highly decentralized network it is capable of relaying information, barring a complete global destruction of infrastructure. *The problem is not resilience; the problem is trust.* If users wish to use the Internet, they must also trust it as a platform, and this is forced trust on a global scale.

4.4 BUILDING TRUST IN

How can we enhance trust in a situation of forced trust on an insecure infrastructure? Dutton *et al.* discuss *trust-enhancing technologies* such as firewalls and other technological security solutions as means to provide trust [97]. Such trust-enhancing technologies are actually technical security measures and devices that enforce conformity to predefined security protocols and policies, and thus by extension create trust by providing incentives for everyone to behave appropriately.

Schneier [95] uses various societal dilemmas, such as the *Prisoner's dilemma* (see e. g. [110]) and the *Tragedy of the commons* [111] to describe trust in society. In the Prisoner's Dilemma, two prisoners are trying to minimize their respective prison sentences after being apprehended. The best scenario in all cases is for a prisoner to betray their accomplice and get a reduced sentence, even if by remaining silent it would be possible for them to get a shorter sentence but only if both prisoners refuse to betray the other. In the Tragedy of the commons, the use of common resources in a society is examined. Even if it is beneficial for the overall society to limit the use of a common resource, it is in the interest of the individual to maximize their use of common resource, even though it will lead to depletion if everyone would act in the same manner. That is, maximizing personal gain at the expense of the society is in the best interest of the individual, even though it will lead to resource depletion as all individuals try to maximize personal gain at the expense of all others.

Schneier uses these examples to construct various scenarios of conflicting interests between different societies, and uses *societal pressures* as a control system and a facilitator for societal trust. He identifies four distinct types of societal pressure: moral, reputational, institutional, and security pressure. Moral pressure works best in small societies, while reputational pressure works

in small-to-medium sized societies. As the size of the society grows, these pressures lose effectiveness, and thus institutional pressure is needed. Finally, security pressure stems from security systems, whether technical, judicial or otherwise. All these combine to provide incentive to members of society to act in the best interest of the group, instead of trying to maximize personal gain at others' expense. In a way, Schneier argues that trust stems from various different sources in a society, and not only from technology-based solutions.

Other authors such as Nestas and Hale [100] define security devices and measures as trust building technologies. In a sense that is exactly what they are, but they do not by themselves induce trust in all actors. For a subset of users, these devices will actually induce trust, but for another subset they might cause more suspicion due to mistrust in technology in the first place. An interesting question that can be derived from this is that does one have to understand all aspects of a security measure before that particular measure can be seen to develop trust?

We can use the concept of the *certainty trough* to model this situation. Introduced by MacKenzie [112, p. 372], the certainty trough describes in an abstract manner how a persons' distance from the source of knowledge affects how certain the person is about that particular knowledge. Those who are not familiar with a technology are understandably less certain about whether it is good or not. Similarly, those who are familiar with the technology and understand all its limitations and shortcomings have a more realistic view of the capabilities of that technology, thus making them less certain about it. In between are the users, who are merely involved in using the technology but are not familiar enough with it to understand its details and limitations. Their sense of certainty is a much higher with regard to that technology than both those who are unfamiliar and those who are intimately familiar with.

This certainty trough also reflects to the level of confidence and trust one has in that particular technology. Thus those who are familiar with firewall and other network security devices may consider them to be trust building and enhancing devices, while those who are unfamiliar or intimately familiar may not consider them as trust building devices at all, or at a reduced level of effect.

The concept of signed certificates demonstrates an approach to trust in IT. It is hard to trust an unknown server on the Internet. The server can be what it claims to be, or it can be a malicious server impersonating as legitimate. There is no way to know this beforehand, so the system of digital certificates was introduced. In it, we have trusted third parties who are responsible for granting certificates to companies and entities responsible for operating the servers, who in turn use them on the servers. The certificate identifies the party responsible for the server and identifies it as legitimate. In turn, browsers trust legitimate certificate authorities automatically to facilitate a smooth and pleasant Internet experience. Usability trumps security in most use cases.

4.4.1 *Accountability in networked information society*

Dutton *et al.* [97] introduce the concept of *trust tension* to describe the juxtaposition of trust inducing and reducing elements in eGovernment. They observe that [97, p. 15]:

> [i]n e-government, it is important to recognize a possible difference between trust in the technology of the Internet - the array of equipment, people and techniques used to gain access to an e-government service - and trust in the people using and communicating through it for the provision of af a particular government service.

Their key observation is that eGovernment systems need data on citizens to function, and that the trust mechanics of these systems are conflicting. On one hand, when users refrain from providing data, this absence impedes trust because it limits accountability in society. For example, the lack of data means also that there is no audit trail to follow. On the other hand, data gathering creates trust anxieties in citizens regarding the use — and especially potential misuse — of that data. This is the juxtaposition at the heart of trust tension. Dutton *et al.* [97] establish the need for a framework for managing trust tension in society. This framework should take into account all the uncertainties stemming not only from technical but also socio-technical aspects of society.

4.5 CONCLUSION

Forced trust is a part of the networked information society. As we have observed in this chapter, there exist various trust relationships in the context of CGISs, where there is little to no alternatives in the trust relationship. This concept of forced trust can also be applied to the infrastructure of the Internet — the foundation of the networked information society. When combined with mass surveillance discussed in Chapters 2 and 3, forced trust relationships on the infrastructure of a society should not be allowed to exist. The only option is to provide alternatives, but as it is with infrastructure, it is not feasible to maintain redundant and even competitive infrastructure.

We should not encourage blind trust in eGovernment systems, but societal mistrust for eGovernment systems is not a good alternative either. On one hand, if citizens of the networked information society lose the interest to question or the capability to understand the systems they have to use when participating in government, the society is vulnerable to outside influence. On the other hand, telling people what are the actual risks in an eGovernment system makes people hostile towards the systems, due to the inability to make an informed risk analysis based on experience and fact. Both options are equally bad, as they lead to the same conclusion, even though the reasons are different. eGovernment systems are often by necessity more complex than their analogue counterparts, even if the underlying process would remain exactly the same.

So far in this thesis we have examined both the technological and societal parts of the forced trust dilemma. In Chapter 2 we examined surveillance. In Chapter 3 we in turn examined technical aspects of mass Internet surveillance, and finally in this chapter we have introduced the concept of forced trust and observed how it behaves in the context CGISs. It is also important to recognize that the problem of forced trust is not solely about trust in inanimate objects; there are people at the heart of this problem, and therefore it is significantly more complex.

Based on the research presented earlier, the conclusion that the basic structures of the networked information society are not compatible with privacy and security for the majority of its citizens is beginning to seem plausible. In the next two chapters

we shall examine two central information society processes, and observe how the fundamental weaknesses we have pointed out affect these processes, and whether they are indeed as safe as they appear.

Part II

EVOLVING PROCESSES IN THE INFORMATION SOCIETY

5

CROSSING BORDERS – MOBILITY OF PEOPLE IN INFORMATION SOCIETY

Life... is strength. That is not to be contested; it seems logical enough. You live; you affect your world.

– Jon Irenicus
Baldur's Gate II – Shadows of Amn (2000)

One of the hallmarks of an open society is free movement; people, goods and funds are allowed to travel freely. A government which restricts the movements of its populace, either foreign or domestic travel, is often a sign of a totalitarian society. Yet from another viewpoint, monitoring the movement of people and goods over state borders is a legitimate right of sovereign nations. Practically all countries monitor their borders, often concentrating on foreign nationals, high-risk individuals and potential contraband goods.

On a smaller scale, we as human beings are relatively mobile, yet also very predictable in our movements. When we consider this from a surveillance viewpoint, the predictability of our movements provides ample opportunity for malicious actors to exploit systems, processes and devices related to human movement for malicious purposes. These actors vary in motivation, resources and threat level posed to average citizens.

In this chapter[1] we focus on the security and privacy problems of biometric passports. As travel documents that are necessary for crossing borders between nations, they are an essential part of human mobility. The biometric identifier embedded passports and possibly stored in a biometric database is a relatively new

1 Parts of this chapter are based on the author's contribution to [9] and [10], and previous work on biometrics done in [113].

development, and forces us to take into consideration the interconnected aspect of information society and what it means for biometric passports. We begin by examining biometric identification in general, and then proceed to biometric passports and biometric border control. Our focus is in the abuse possibilities of biometric passport systems. Additionally, while the progress of security protocols has not halted since the publication of the above articles, the development of biometric passports and border control has been slow when compared to previous years. A more detailed examination of the details of security features introduced after 2012 are outside the scope of this thesis, barring the short discussion in Section 5.4.

The problems with biometrics have been discussed in literature for years, and yet the potential problems that have been foreseen to arise, well, have not done so. On the contrary, biometric identification is on the rise. Facial recognition algorithms have progressed significantly and are now relatively accurate: latest results claim near-parity with human capability of face recognition [114, 115]. Fingerprint recognition has gained traction in consumer electronics and is routinely used for unlocking smartphones, for example. Biometric passports are used for automatic traveler recognition on airports, and the systems have functioned without catastrophic failures. So the question is: have the underlying problems been alleviated, are we merely ignoring the problems for the sake of convenience, or is the answer something entirely different?

5.1 IDENTIFICATION AND VERIFICATION

Before we examine the properties of biometrics, we must first examine one of the central problems for which biometrics is applied: *authentication*. From the point of view of an information system, a user should be allowed to use only the services that the user is permitted to use. Regardless of the mechanism how these permissions are decided, authentication aims to preserve the confidentiality and integrity of a computing system by only allowing authorized access. In order to differentiate between users, they must be identified in some manner. Biometrics can be used to give a solution to the two following problems [116].

1. Proving that you are who you say who you are
2. Proving that you are *not* who you say you are *not*

The first problem is *verification*, in which your identifier – whatever the identifier is depends on the system used – is matched against a single sample known to be yours *a priori*, making this a *one-to-one* comparison. If the result of the matching is positive, you have proved to be who you said you were, and your identity is verified to be your own. An easy practical example of this is when a convenience store clerk checks your ID when purchasing restricted items.

The second problem is *identification*. It is a *one-to-many* comparison where the purpose is to verify your identity to be yours and only yours. The goal is to prove that there are no duplicates in the database for your entry, meaning that there is only one of you. If the comparison yields only one match, your identity is unique. The definition of the question can also be thought of as you stating that you do not belong to a certain group, and then this is verified by comparing your biometric to those in the group. If it yields no match, they you are not who you say you are not, and have passed the verification test. A practical example of this would be a terrorist no-fly list or a list of *persona non grata* [117].

Identification and verification are distinct problems. Verification only requires a 1-to-1 comparison between two identifiers, but identification is a 1-to-N comparison, which becomes resource intensive as N grows. While not a factor when N is small, if we must compare an identifier to a database containing hundreds of thousands or millions of samples, the speed and accuracy of the matching process becomes highly relevant to the overall system functionality.

Identification and verification can also have different definitions outside the realm of biometric identification, but such alternate definitions are not discussed further in this thesis.

5.2 BIOMETRICS FUNDAMENTALS

When we attempt to tell one person from another, we instinctively recognize certain elements such as height, body structure, voice and facial features from other people and are able to use

them for this task. Humans instinctively use biometric identifiers for recognizing other individuals. Biometrics is the use of physiological and behavioral identifiers for identifying a person. The use of biometrics and biometric identifiers in recognizing people goes back at least to ancient Egypt [118]. The foundations for modern biometric recognition were built in the 19th century, when the use of biometrics in forensic and criminal investigations began to emerge (see, for example, [119]). Biometrics were mainly used in law enforcement context for the greater part of the 20th century, until the emergence of civilian applications in the mid-to-late 20th century [118].

For a behavioral or physiological characteristic to act as a biometric identifier, it should fulfill the following conditions [120]. An identifier should be *universal*, meaning that every person should have the characteristic, and just not some subset of people. The identifier should be *unique* in itself, so that the characteristic on each individual is different. The characteristic should also be *permanent*, thus remaining relatively invariant over time. Finally, the identifier should be *collectible*, i. e. the characteristic should be measurable quantitatively. Practical applications of biometrics impose three other conditions to a possible biometric [120]. The characteristic should have good *performance* regarding identification accuracy, resource requirements and environmental factors affecting accuracy. The use of the characteristic should be *acceptable* to the people who use it. Finally, a system based on the characteristic should be hard to *circumvent*, either by accident or by defectors.

Based on their performance on the previously introduced criteria, different biometrics have different application domains and uses. Some of the most prominent biometrics are categorized according to these criteria in Table 1 [121].

By examining the different properties of biometrics, we can see that not all biometrics are suitable for every task. More powerful biometrics such as the iris and DNA have low acceptability ratings for users, which can stem for example from inherent protectiveness to body integrity [118]. Another possibility is the prominent use of that particular biometric in law enforcement purposes [122]. While fingerprints are generally associated with criminal investigations, it can be argued that if a person is famil-

Table 1: Biometric identifier characteristics. Scale: H=High; M=Moderate; L=Low. [121]

	Fingerprint	Iris	DNA	Face	Signature	Voice	Gait
Universality	M	H	H	H	L	M	M
Uniqueness	H	H	H	L	L	L	L
Permanency	H	H	H	M	L	L	H
Collectability	M	M	L	H	H	M	H
Performance	H	H	H	L	L	L	L
Acceptability	M	L	L	H	H	H	H
Circumvention	M	L	L	H	H	H	M

iar with the use of fingerprints in non-threatening situations, as a result they can be more accepting of their use in general.

Two distinct groups of biometrics can be discerned. If a biometric is fixedly related to the physiology of the person, it is considered to be a *physiological* identifier. If the biometric is a representation of a learned property, it is considered to be a *behavioral* identifier. While physiological identifiers — height, body proportions, fingerprint, iris, DNA etc. — are straightforward, behavioral identifiers include such characteristics as signature, gait, keyboard typing patterns, and speech. While physiological identifiers are consistently stronger, behavioral identifiers can be used in recognition to a certain extent. An intuitive way to recognize a person from a distance is, for example, the way a person walks.

5.2.1 *Biometric recognition*

Biometric recognition is based on sampling the target biometric identifier and storing it for future reference. The sampling process naturally varies by the identifier — e.g. photographing for face, audio recording for voice, physical sample of saliva for DNA — but the end result is a sample of the original biometric. If necessary, multiple samples of the same biometric identifier can be taken and combined, but the end result is a representative of the original biometric identifier, referred to as a *template*. It contains all essential information about the biometric identifier, with the important property that it can be compared, automat-

ically or manually, to other templates, yet again derived from samples of a biometric identifier.

Biometric identification is not exact in nature: it is impossible to do a bit-wise comparison of two samples of a biometric identifier and get a match. The sampling process introduces errors and variations to the samples due to differences in user interaction, sensor accuracy and environmental factors. For example, fingerprints can be smudged or the finger itself can be damaged in some way, and the person giving the fingerprint sample is very unlikely to be able to use the exact same pressure and positioning when giving the original fingerprint. Facial images can vary between expressions, angle, lighting and makeup. Iris images can be slightly off-position, or there can be occlusion from the eyelid in parts of the iris. These behavioral and environmental issues practically guarantee that biometric samples have high *noise*. In this context, noise refers to any external factors affecting the biometric sampling process. Even the most precise sensors have their limits in accuracy, and when this is combined with behavioral and environmental factors, they cause systematic errors in any biometric sampling process [123].

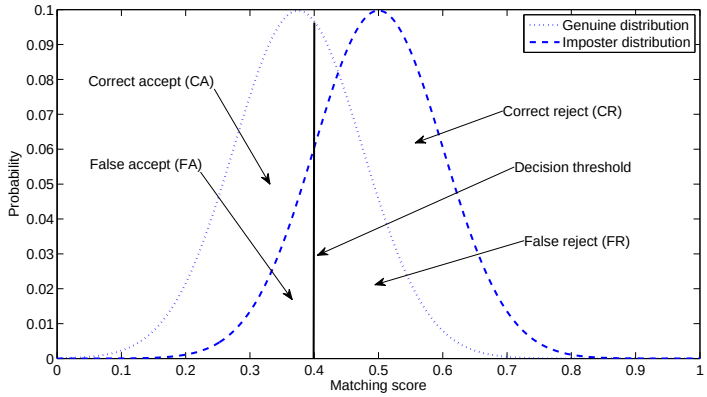

Figure 3: Example probability distributions of genuine and imposter templates in template matching.

Biometric identifiers are reduced to a biometric template, which contains the essential information of the biometric, and which is also comparable to other templates. Because of the differences in biometric samples, templates constructed from two different samples are naturally different. The challenge of biometric recog-

nition is in ascertaining if the two templates come from the same original biometric identifier, i.e. from the same person. The two templates are compared for similarity, and a decision of match/no match is made according to this comparison. Because this comparison is not exact, the result of this matching is a probability of a match. The threshold for a successful match defines the probabilities of matching errors in the process. An example is given in Figure 3, where the comparison metric between templates is the *Hamming distance*, and the threshold is set at 0.4. Hamming distance in this case is the distance between two binary codes, or specifically the amount of disagreeing bits between the two binary codes divided by the total amount of bits compared.

Because of the statistical nature of biometric identification, the possibility for errors in matching exists. The types of errors possible in statistical decision making and by extension, biometric matching, are *False Accept* (FA) and *False Reject* (FR), corresponding to Type 1 and Type 2 errors, respectively. *False Accept Rate* (FAR) and *False Rejection Rate* (FRR) define the efficiency of the decision making scheme. In addition to these errors, the correct results are defined as *Correct Accept* (CA) and *Correct Reject* (CR).

The probability of making an error in biometric matching can be formally defined. Let the two probability distributions in Figure 3 be $P_{Im}(x)$ and $P_{Ge}(x)$, for imposters and genuine users, respectively, and the decision threshold is $\tau \in x$. Now the areas under these distributions can be defined as [124]:

$$P(FA) = \int_0^\tau P_{Im}(x)dx \tag{1}$$

$$P(CR) = \int_\tau^1 P_{Im}(x)dx \tag{2}$$

$$P(CA) = \int_0^\tau P_{Ge}(x)dx \tag{3}$$

$$P(FR) = \int_\tau^1 P_{Ge}(x)dx \tag{4}$$

For these equations, the following equalities and inequalities are true [124]:

$$P(CA) + P(FR) = 1 \tag{5}$$
$$P(CR) + P(FA) = 1 \tag{6}$$
$$P(CA) > P(FA) \tag{7}$$
$$P(CR) > P(FR) \tag{8}$$

5.2.2 *Problems of biometrics*

The use of biometrics for identification and identity verification purposes is not without problems. Biometric identifiers are, by their definition, unique to the subject, and permanent. This identifier uniqueness presents a serious vulnerability in the case of widespread use of biometrics. If the biometric identifier is ever compromised, i. e. leaked into the wrong hands, the identifier cannot be reliably used for identification. Because biometric features are permanent, they do not change over time, and so are considered compromised for the rest of the life of the subject. While in science fiction it is possible to change body parts for the purpose of avoiding biometric recognition,[2] in the real world this is not (yet) possible. Advances in medicine can perhaps make this a reality, but so far, we will have to abide by the status quo of only having one physical body, mostly without changeable spare parts.

The salient problems of biometrics can be divided into technological and societal categories. Technological problems arise from the ramifications of technology, while societal problems stem from various different sources, most notably from the people operating, administering, designing and implementing biometric systems. People make mistakes.

5.2.2.1 *Technological problems*

Taking the nature of biometrics into consideration, the authentication problem with a biometric identifier becomes interesting.

2 An example from popular culture can be found in the film adaptation of *Minority Report (2002)*. In the movie, John Anderton, portrayed by Tom Cruise, has to undergo eye transplant surgery to avoid iris-based automatic biometric recognition in public places.

While biometrics have an advantage versus the two other alternatives – something you have and something you know – one critical flaw remains. A key card can be revoked and a new one can be assigned in case the previous token went missing. A password or a PIN code can be changed at will if it becomes compromised. This is simply not possible with biometric identifiers. This is one of the strongest arguments against large-scale biometric identification; if biometrics is in large scale use, biometric spoofing will become more than lucrative – almost a necessity for any criminal. Revocable biometrics [125] have been proposed as a solution to this problem. Using the biometric identifier as a base, distinct revocable biometric identifiers are derived from the original by distortion or other means. If such a derivative identifier becomes compromised, it can be revoked and a new derivative identifier can be issued, with a different distortion function. By choosing the method of distortion properly, it is hard to discern between different possible identifiers, even if the distortion function is known to the attacker. Other methods later proposed for obscuring templates include fuzzy vaults [126, 127]. Further analysis of methods for template scrambling are however out of the scope of this thesis.

The use of such distortion functions for fingerprint templates may present one additional problem. If the distortion function is known to the attacker, they might be able to calculate candidates for the original template. For example, ISO standardized fingerprint templates are reversible [128] to natural fingerprints, so great care must be taken to assure security and privacy of templates. Template generation should be a one-way function, so that the original identifier cannot be accessed from the template. In most cases this is a moot point, as the template is effectively the identifier itself, as far as automatic recognition systems are considered. But in the case of fingerprints, this can pose a problem. While a human expert will probably not be fooled with a reconstructed fingerprint [128], automatic systems are more susceptible. While generating fingerprints form a desired template and plating them into a crime scene is perhaps not an efficient forensic countermeasure, it can cast undue suspicion on the owner of the fingerprint.

It is also possible for a person to not possess a certain biometric at all. Fingerprints can, and do, erode from manual wear

and tear, contact with chemicals, and injuries. It is possible to not have fingers at all, for example in the case of amputees. There are several scenarios where it is perfectly plausible that the person is unable to use a biometric identification system because of a missing biometric. It depends on the administrator of the system how such outliers are managed. For example if a Finnish citizen applying for a passport does not have index fingers, other fingers are used instead and if the person has no fingers at all, then the passport is granted without embedded fingerprints[3]. While this is the current situation, one cannot know into which direction these policies will evolve. A worst-case scenario could be that a person without a necessary biometric would effectively become a second-grade citizen in the society, because of the inability to use said biometric for recognition.

The accuracy of a biometric recognition system depends on several factors. Identifiers have different properties which translate to different recognition accuracies, as is evident from the characteristics in Table 1. The inherent accuracy of an identifier varies, and environmental factors add uncertainty to the recognition process. The occasional error in recognition can happen, but when the volume of identification grows, the statistical error grows accordingly. This results in very accurate systems still giving erroneous results for a small portion of the users.

One can argue that the fraction of false identifications or rejections is so small that there is no noticeable overall effect to the users, but this problem cannot be isolated as a technological issue. The result of a false identification or rejection is uncomfortable to the user, and because of the trust placed on biometric systems, this will attract undue suspicion towards the user. Additionally, the response to such an incident is a societal issue; depending on the policies in place it can vary from a discreet retry of identification to immediate incarceration. This is a problem technology alone cannot mitigate, and technology actually facilitates this problem extensively. This will be discussed in the following section.

3 Suomen poliisi. Lisätietoa passin hakemisesta. Online, available at https://www.poliisi.fi/passi/lisatietoa_passin_hakemisesta Accessed 25.9.2017.

5.2.2.2 *Societal problems*

As a society, we usually trust technological solutions, especially if we do not fully understand the particular functionality behind the systems. The certainty trough [112], previously discussed in Chapter 4, can also be applied to the use of biometrics, similarly to trust-enhancing security measures. Those who are very unfamiliar with biometric recognition systems may have reservations on their functionality and accuracy, because they do not know how the systems operate in the first place. Those somewhat familiar with biometric systems – limited use experience or limited marketing level knowledge of systems – tend to be more optimistic on the security, privacy and effectiveness on biometric recognition systems. Finally those with intimate knowledge of the underlying problems tend to demonstrate more suspicion and reservations towards biometric systems and their use.

It should be noted that biometric identification systems are, in a broad sense, exactly those trust-enhancing and generating security measures discussed in Section 4.4. Biometrics are used in access control and identification, both functions essential to a secure information system, or in a larger scale, a secure networked information society.

Most people tend to fall into the two first groups – unfamiliar and somewhat familiar. This has implications on how these systems are perceived, and more critically, how people who – for any reason – end up getting the metaphorical short end of the stick. Perception and assumption of guilt is critical in this regard. If someone gets caught by a facial recognition system designed to identify shoplifters, most people witnessing the event would directly assume that the person was guilty of *something*, even without any evidence to the contrary, as security technology solutions are seen as functional and reliable. A false identification, whether positive or negative, can cause substantial harm to an individual in the form of increased scrutiny, stress, and other possible societal disadvantages.

The next question is where biometric identification devices are used. Because of the mental image of high security, they are found in high security processes and facilities requiring authentication, such as borders, vaults, restricted areas. The other approach to biometric recognition is of convenience. Because your

biometric is the only key you need to possess, users of the system don't have to carry keys with them, or remember pass codes. For example, fingerprint scanners as access control devices to gyms and equivalent facilities are more and more common. These kind of systems value convenience over security, so usually the only authentication factor is the biometric, no tokens or PINs. This may increase convenience and even security, but the underlying issues are still present.

If we disregard false reject situations when attempting to access a gym locker of failure to unlock your smartphone on the first attempt, the other side of biometric recognition offers more serious scenarios. A false reject error when crossing national borders can mean, depending on the nation, minor inconvenience at best and incarceration at worst. In general, the higher the security level of the protected target or process, the higher the response to an error. An unfortunate example can be found in the case of an individual who was unduly attached to the investigation of the Madrid bombings in 2004 as a result of a fingerprint match performed by the FBI [129]. An interesting problem in biometric identification arises when the expert who performs the matching of biometric identifiers has contextual information about the situation, suspect or motives related to the investigation. This has been found to have an effect on expert performance [130]. This in turn leads to the situation where a false match or nonmatch can happen due to external circumstances, and the subject of the identification has very limited possibilities for recourse and compensation on any adverse effects resulting from this error. At worst, it can mean incarceration as an enemy combatant or equivalent label, effectively stripping the person of any rights whatsoever.

Because of the nature of biometrics, the existence of a perfectly accurate biometric recognition system is science fiction at best. While this fact is often overlooked in marketing and policy positions, false positives and negatives are always a possibility – however small – in biometric recognition.

5.3 BIOMETRIC PASSPORTS

One of the most visible manifestations of biometrics to the people is their use in passports. The biometric passport standard [131]

developed by the ICAO – a United Nations specialized agency – describes the specifications that biometric passport designs must comply with. The biometric passport is a standard passport with an embedded Radio Frequency Identification (RFID) chip, which in turn contains remotely readable information on the passport holder. The RFID chip is a standard ISO14443-A or ISO14443-B contactless chip. The standard allows for storing fingerprint and iris biometrics on the chip [131], and the chip also contains the information on the passport information page. A biometric passport functions as a regular passport, but the data contained on the RFID can be remotely read by a border official, and the biometric data can be used for automated biometric matching of the passport holder.

Biometric passports have been in use for a bit over a decade, and the standards have been under constant development as new security threats have been identified and mitigated. While the biometric passport is standardized and the specifications are publicly available, the implementations of biometric passports by individual countries are consistently kept secret. This has resulted in a concentrated reverse-engineering effort by security researchers [132]. These efforts have revealed several vulnerabilities in current biometric passports.

5.3.1 *Data structure of ICAO passports*

The data on the chip is structured into a Logical Data Structure (LDS), which all passports must implement as is without exceptions. The LDS is subsequently divided into different Data Groups (DG). The structure of the LDS is presented in Figure 4. Some future but not yet implemented features of the LDS have been left out of the figure for the sake of clarity.

As it can be noted from Figure 4, the ICAO passport standard allows for storing multiple biometrics of the passport holder, namely face (DG2), fingerprint (DG3), iris (DG4), and signature (DG7) [133]. Iris images are not yet actively stored on any passports, but the capability is built in the standard for future use.

Detail(s) recorded in MRZ	DG1	Document type		
		Issuing State or organization		
		Name (of holder)		
		Document number		
		Check digit – document number		
		Nationality		
		Date of Birth		
		Check digit – date of birth		
		Sex		
		Date of Expiry or valid until date		
		Check digit – DOE/VUD		
		Optional data		
		Check digit – Optional data field		
		Composite check digit		
Encoded identification feature(s)	Global interchange feature		DG2	Encoded Face
	Additional features		DG3	Encoded Finger(s)
			DG4	Encoded Eye(s)
Displayed identification feature(s)	DG5	Displayed portrait		
	DG6	Reserved for future use		
	DG7	Displayed signature or usual mark		
Encoded security feature(s)	DG8	Data feature(s)		
	DG9	Structure feature(s)		
	DG10	Substance feature(s)		
	DG11	Additional personal detail(s)		
	DG12	Additional document detail(s)		
	DG13	Optional detail(s)		
	DG14	Security options		
	DG15	Active Authentication public key info		
	DG16	Person(s) to notify		

Figure 4: Structure of the LDS on an ICAO biometric passport, according to [133], Figure 1.

5.3.2 *Security measures in the ICAO biometric passport*

The data integrity on the passport chip is verified using Passive Authentication (PA) [134, p. 21–22], which computes the hash digests of the data stored on the passport and compares it to the Security Object (SOD) stored separately on the passport. Any tampering with the data on the passport will result in PA failure.

Notably, PA does not protect against passport cloning attacks, because the clone passport would have identical data structure and content, and thus would pass PA without any problems.

Basic Access Control (BAC) [134, p. 7–10] is the primary defense measure against unauthorized access of a biometric passport. It prevents access to the chip without the owner's consent by requiring RFID reader to possess the Machine-Readable Zone (MRZ) of the passport, printed on the passport information page. Encryption keys are derived from the MRZ and used to establish secure communications between the reader and the passport. This is a mandatory security feature in all biometric passports.

Active Authentication (AA) [134, p. 23–26] protects the passport chip against substitution, i. e.against having a cloned chip in a passport, as PA is unable to differentiate between a cloned chip and the original. It is based on a non-readable private key on the passport, and is implemented on a number of passports [135, p. 19], but is not a mandatory security feature.

Extended Access Control (EAC) [136] is an enhancement to the previously presented passport security solutions. It consists of various protocols for authenticating both the passport chip and the reader, and has its own certificates for readers, thus mandating the existence of a public key infrastructure. Two versions of EAC exist: EACv1 and EACv2 both provide terminal and chip authentication, but v2 provides improved security and privacy protections by using Password Authenticated Connection Establishment (PACE) [134] as a replacement to BAC.

5.3.3 *Problems with biometric passports*

With the emergence of biometric passport technologies, the problems and weaknesses related or derived from them have also become more salient. The use of biometric passports presents several different problems to security and privacy of passengers all around the world. Some problems happen because of inadequate technical protections on biometric data, some happen because of unforeseen consequences of having such a system in place, and some problems manifest because of foreseeable abuses of the system, without equivalent checks and balances to keep the system from being misused.

The technical protections on biometric passports are inadequate on several levels [10]. The possible attack scenarios against ICAO biometric passports are various, and in the following sections these attacks are examined in detail.

5.3.3.1 *Eavesdropping attacks*

The passport has a wireless interface through which information is read from the embedded RFID chip. These wireless communications can be eavesdropped by an attacker, or the attacker can initiate a connection to the chip on their own. The range to which the communications can be recorded or initiated by an adversary varies. Readers provide the RFID power through induction, and are thus limited in range. While commercial readers can interact with RFID chips to the distance of a couple of centimeters, with a specialized antenna for the reader this can be increased to 25 meters or more [137]. This can be used by an attacker to interact with the chip from further away. Intercepting the communications of a reader/passport conversation is also a matter of physical proximity. As eavesdropping communications does not require supplying power to the chip, the range of listening to an exchange between a passport and a reader can be extended with suitable equipment up to 500 meters [137].

Because of this unavoidable vulnerability, all communications between the passport and the reader are encrypted. This process is known as Secure Messaging (SM). First the data contained in the MRZ – the two lines of text on the information page of the passport – is read using an optical reader. From this data, the document number, date of birth and date of expiry (with associated check digits) are concatenated, forming the data entity *MRZ_information*. The key seed k_s used for deriving session keys for communications between the passport and reader are the 16 most significant bytes of the SHA-1 digest of the MRZ_information field. The encryption algorithm chosen for encrypting communications between passport and reader is 3DES in CBC mode. [131] While for example AES [138] would have been a more robust choice, when the ICAO standard was under development, AES was still a relatively recent algorithm.

Encryption does not prevent an adversary from eavesdropping and storing the communications for later analysis. Because

the MRZ is used to derive the encryption keys for communications, and there is no random element included in the key derivation process, the security of SM depends solely on the MRZ. The key seed k_s is 128 bits, but only contains information which is on the MRZ. It has been shown to have insufficient entropy due to the structure of the underlying data in the MRZ [139]. While online attacks against a passport are not feasible, bar extended contact with the passport such as on an intercontinental flight [135], offline attacks can be mounted against captured communications at the attackers leisure. The more is known about the passport holder (age, nationality, etc.), the easier the offline attack becomes.

5.3.3.2 *Identification and traceability attacks*

ICAO passports use standard RFID collision avoidance techniques for the situation where there are more than one RFID chips within the range of a reader. The passport generates a random ID number used for identification of individual chips. A reader can differentiate between passports by these ID numbers to avoid collisions between communications between chips. While ICAO passports should always provide a new random ID for each time a reader powers up the chip, some passports give a constant ID number [135]. This enables an adversary to track the movements of a single passport with a known constant ID; whenever a passport with a known ID comes within range, it can be identified. Other passports provide a random ID but with the prefix byte '08' [135]. This allows an adversary to identify a biometric passport by looking at the first byte of the collision avoidance ID. On the extreme side of attacks made possible by identifying passports on the fly are possible terrorist applications, such as RFID triggered bombs, designed to detonate when a certain number of passports of a certain nationality are detected in reader range [140].

The response of a passport to an unauthorized read request can also be recorded and fingerprinted [141]. While the ICAO standard defines a certain behavior for passports subject to unauthorized read attempts, passports tend to answer differently to malformed queries, or queries done in incorrect phases of the protocol. This is speculated to be because of different manufac-

turers having individual implementations of the passport standard, all differing in actual behavior. These implementations are also usually classified as industrial secrets, which leaves researchers the only option of reverse-engineering them to determine their behavior in error situations.

The RFID chip itself can also be identified by its physical layer properties. RFID chips of the same model and manufacturer can be differentiated on the modulation shape and spectral features of a chip responding to correctly formed and out of specification signals [142]. While this method requires careful experimental setup and care, it can still be used for fingerprinting RFID chips with a reasonable accuracy.

5.3.3.3 *Location privacy and predictability of human mobility*

The location of a person at a given time can be considered trivial information, but access to even a limited set of location information for a person makes it possible to make relatively accurate predictions on how that person will behave. Song *et al.* [143] have analyzed the predictability of human mobility. People are very predictable in their daily movements. While the day-by-day movement data of a person may appear to be random to a casual observer, given sufficient data points for analysis the potential predictability of the movement of a person is high. It was also found that there are no marked differences in predictability between people constrained in a 10 km radius and people who travel more than 100 km per day.

These observations strongly confirm the old adage of people being creatures of habit. This makes targeted surveillance of a person so much easier: it is enough to know only a few data points of a persons daily itinerary and it is possible to predict their movements with relatively good accuracy. Also, deviations from this established pattern are easy to notice, quickly raising red flags to potential observers.

A common example of location data that is automatically generated is cellular network control data. This refers to both data that is automatically generated and necessary in order to provide the cellular service, and to purposefully generated cellphone location data. The former is necessary information that the network needs in order to connect calls and relay data to the phone.

The latter is used to provide people with location-based services, mostly in mobile applications.

Location data is often anonymized to protect the privacy of users, while still providing enough information to generate the necessary services. This anonymization is, unfortunately, not effective in hiding individual users in large datasets. Zang and Bolot [144] examine the effectiveness of anonymization in providing privacy protection to users. They found that the steps necessary to effectively anonymize large-scale datasets for the purpose of privacy protection of individuals significantly impacts the utility of said data to service providers. This clearly makes such anonymization efforts unappealing to operators.

5.3.3.4 *Conclusions on vulnerabilities*

While the security features of biometric passports have been evolving with each new form of attack, it can be noted that the threats to biometric passports and the data they contain still exist. Even though EACv2 can be considered to be secure, the fact that its implementation is voluntary and up to bilateral agreements between countries does not make it an efficient solution. The ICAO standard mandates that only implementing PA is mandatory, thus making other security solutions optional. The fundamental flaw that all passports must accept BAC authentication regardless of any more advanced authentication systems implemented on the passport gives the attackers a known, good attack vector with the possibility of gaining at least some of the data on the passport [137]. Traceability and identification attacks are also unmitigated, and provide a way to track passports and their holders.

It can be argued that using these techniques is expensive and not worth the effort for an organization or an individual, and that potential use scenarios are limited. In this case, these security concerns are perhaps over-emphasized for most potential actors. Entities that could realistically leverage biometric passport system vulnerabilities include sovereign states or "Three-Letter Agencies" — first tier intelligence organizations such as CIA, FSB, Mossad and MI6. On one hand, a player of this magnitude would have access to practically any necessary resources, methods and manpower required to breach the security of bio-

metric passport systems. On the other hand, it can be argued that, being an extension of the government of a sovereign country, they could just simply ask for the information stored on the passport from other government agencies and be done with it. Still, the weaknesses in the passport standard provide an avenue of exploitation which leaves little to no trace of possible tampering or misuse, and confers plausible deniability of any attack ever taking place.

5.3.4 *Biometric databases*

In Finland, the biometric data gathered from passport applicants is also stored in a national biometric database.[4] A comprehensive database of citizens' biometrics can be used for noninvasive purposes, such as identifying the deceased in case other methods fail. The criticism against national biometric databases is usually based on concerns about privacy, leakage of sensitive data, function creep, and malicious use of biometric data [10]. The purpose of these databases has been the subject of controversy in, among others, Finland and Germany, where completely opposite views have been formed on biometric databases. In Germany, the existence of such a registry was seen to be in conflict with the German constitution and the right for self-governance of personal information [145]. In Finland, the existence of a national biometric registry was not seen as a problem, as long as the collected data is not used for other purposes than it was originally gathered for.[5]

Privacy concerns about biometric databases are centered on the nature of biometric data, and the ability of a person to decide whether to give such data to someone else or not. The inclusion of biometric data in passports makes the scenario of international travel practically impossible without such a document, thus making it mandatory to give your biometric identifiers to the government for storage and use. The user, who in the case

4 Suomen Eduskunta. Passilaki 6 a §. Online, available at http://www.finlex.fi/fi/laki/ajantasa/2006/20060671 Accessed 19.6.2017.

5 Eduskunnan perustuslakivaliokunta, Perustuslakivaliokunnan lausunto PeVL14/2009 vp. Online, available at https://www.eduskunta.fi/FI/Vaski/sivut/trip.aspx?triptype=ValtiopaivaAsiakirjat&docid=pevl+14/2009 Accessed 25.9.2017.

of eGovernment systems is the citizen of the networked information society, actually has very few options. Either the citizen must trust the institution in question to handle the data responsibly, and use the system, or opt out. The latter case usually means that the citizen cannot participate in the activity — such as travel, in this case — at all. Currently in Finland it is not possible to have a document which serves as an official proof of identity — a passport or an identity card — without giving fingerprints to the national fingerprint registry. This can be worrying, at the least. While a driving license is not an official proof of identity in Finland, it acts as a *de facto* proof of identity in most aspects of Finnish society. It is not always accepted in some key situations when dealing with authorities, as its primary function is to show that the holder has a right to drive a vehicle. Also, not everyone has a driving license in the first place due to excessive cost or personal choice. This leaves no alternatives for those not able to afford or unwilling to get a driving license or a passport.

Biometric data is considered to be sensitive personal information. Biometric identifiers are relatively immutable, so they are valid for the lifetime of the user and cannot be changed if an identifier becomes compromised. Biometric identifiers can implicate presence of or predisposition to certain genetic diseases [122]. Even the lack of biometric data can reveal information. If a person has no recorded fingerprints for a normally recorded finger or fingers, it can be assumed that the person has a medical skin condition or is missing the fingers and/or related limbs altogether. Most importantly, some identifiers provide a method for identifying a person without their consent or knowledge, violating their privacy.

As an example of function creep, the government can use the collected passport data to other purposes that was originally intended. This has already been seen in Finland [10], where the use of the national biometric database has been widened from its original narrow scope. Look at the other effects of this kind of behavior. In Germany, the storing of biometric identifiers into a national database was explicitly forbidden by the constitutional court [145]. The storage of such private and confidential information was deemed to be unconstitutional under the German right of managing your own data and how it is used.

The European Court of Human Rights (ECHR) has taken a stand against retention of fingerprints in national databases of a person not convicted of a crime [146]. In the judgment, it was unanimously found that retention of a persons fingerprints after being found not guilty was a breach of Article 8 of the European Convention of Human Rights,[6] which guarantees an individual the right to respect for private and family life. This gives a strong signal that the use of biometric data collected for a distinct purpose should only be used in that particular purpose. This, at least, lays some of the concern shared by security and privacy researchers to rest, but it yet remains to be seen how the member states will react to this ruling in practice.

5.3.5 *Biometric border control*

With the integration of biometrics into passports, the automated use of biometrics in border control has subsequently increased. While the goals of biometrics are to make traveling documents harder to forge and verifying identities of travelers faster and more reliable, the use of biometric identifiers in border control has also other dimensions.

When a border control official reads the chip on a passport, all of the information on the passport chip is displayed to the official. This is obvious, as the information needs to be checked by the official – that's what the officials are there for. Before electronic systems, the border guards would examine each passport closely to attempt to ascertain their authenticity, and that it belonged to the passenger presenting it.

The experiences of automatic biometric recognition on national borders have not always been favorable in the past. In the UK, a system named Iris Recognition Immigration System (IRIS) was introduced in 2005[7] and decommissioned in 2013.[8] The addition of biometrics in passports and learning from the failed experiments in the past have made automatic border control a reality in contemporary international travel. By comparing the stored

6 Council of Europe, European Convention on Human Rights. http://conventions.coe.int/Treaty/Commun/ListeTraites.asp?MA=3&CM=7&CL=ENG

7 BBC News http://www.bbc.com/news/uk-england-17058448

8 Findlaw.co.uk http://www.findlaw.co.uk/law/immigration_emigration/other_immigration_law_topics/iris/30854.html

identifiers in the passport with a new sample taken with an autonomous system, it is possible to relatively reliably identify a person with a fully automatic process. Such biometric identification systems are already in use at airports all around the world. So far they would appear to be working as intended, or at least this is the impression given by the lack of reported issues with the latest automatic passport control systems.

Computer systems make it possible to store, index, and search large amounts of data. This has already been discussed in the context of surveillance in Chapter 2. This differentiates biometric border control from the past, when it was not possible to effectively cross-check passport records. Now, having a registry of people who have crossed national borders in a particular country is not massively problematic in itself. The real issue with security and privacy arises when we consider how much we trust the *country* and its authorities that are doing the reading, storing, processing and managing the information?

A general model of a biometric passport system and border control is described in Figure 5. It describes the central actors in a biometric border control scenario in a networked information society. Key data and how that data flows from actor to another is also illustrated. From the point of view of the citizen, the forced trust relationship that exists between the citizen and the government providing the passport becomes even more complex, as governments are forced to share information with other governments when their citizens cross international borders. In this case the forced trust on the passport and its security features is also extended to the foreign government reading the passport during border crossings. The foreign government does not have any of the responsibilities the travelers' own government has regarding safe and secure storing of their travel information. International agreements on read access to biometric passports may indeed take this into account, but it is quite difficult to remain uncynical regarding such agreements, especially when governments can be so asymmetrically matched in power and influence. Gathering as much information as possible on foreign citizens is in the interest of all sovereign governments.

If a foreign government does not store the data according to robust security standards, there exists the possibility to gain access to information on citizens by attacking a completely different go-

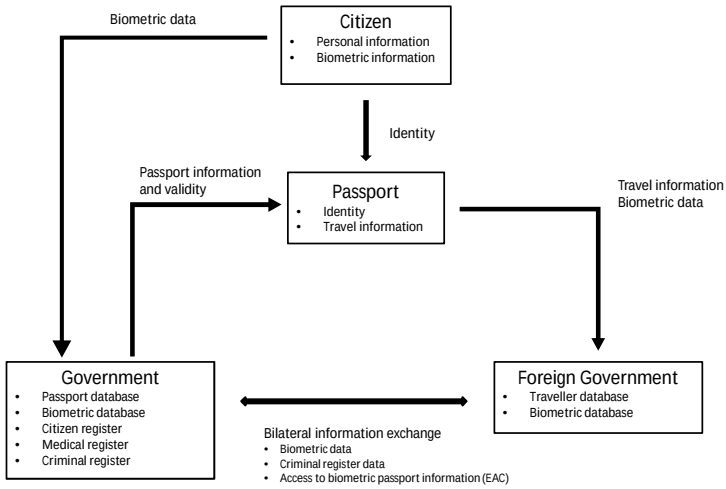

Figure 5: Description of the information flows with biometric passport systems.

vernment information system that is originally storing the data. In the case of Germany, where biometric data is not stored centrally at all, this could lead to biometric data of such citizens being compromised through a less privacy conscious government. Other potential malicious parties include spies, criminals, terrorists, and corrupt officials. These adversaries that are interested in malicious actions against a biometric passport system are shown in Figure 6.

5.3.6 *Biometric passports and forced trust*

Crossing national borders is a good example of an important societal process in the modern world. It requires copious amounts of trust from all parties involved: the traveler, the border control agent, and in a wider scope, their respective nations. Free movement between countries can be seen as a major trust issue. The Nordic agreement on travel documents and the free mobility within the Schengen area, for example, facilitate mobility across borders. There are complicated processes involved behind the scenes to make sure that this kind of freedom can be granted to people. The initial assumption is that nations trust each other and their citizens, and expect reciprocity from other nations and citizens – this assumption of trust is necessary for the system to

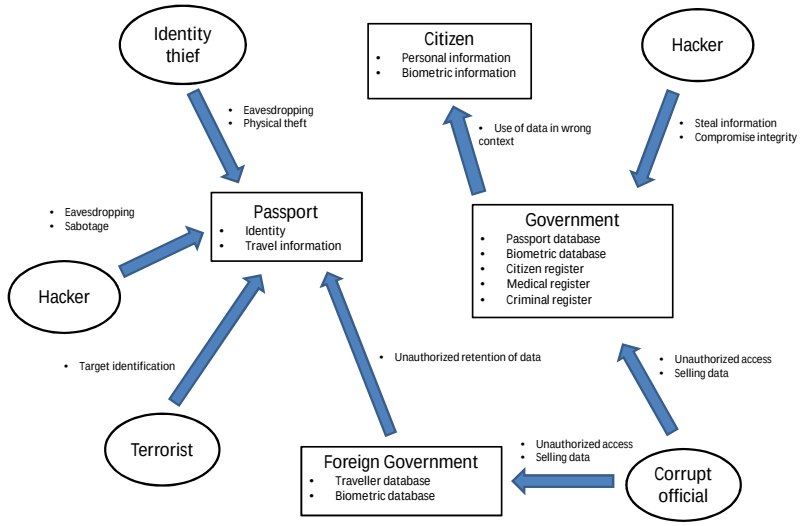

Figure 6: Adversaries in a biometric passport system.

work, as was discussed by Schenier [95]. Some countries do not trust their citizens, however, and limit their capability to travel freely.

The concept of forced trust is pertinent to the discussion on biometric passports. Owners of biometric passports are in a forced trust situation with their government, the passport manufacturer, and any governments that are able to access the passport information, whether it is through regular travel or information exchange. Biometric passport systems are clearly a part of critical infrastructure due both to the critical function they provide and the importance of the information that they store. This must be taken into account when designing these systems, and care must be taken to examine all potential trust relationships in a critical manner.

It can be argued that choice exists in this scenario. Indeed it has always been the case that traveling to certain countries will make entering others next to impossible. For example, traveling to Iran after visiting Israel is practically impossible, and vice versa. But this choice has only to do with entrance to a country. In the case of biometric passport systems, there is more on

the line than just location privacy – something that can be debated whether it should apply on national borders or not. Now when you consider the forced trust relationships illustrated in earlier in Chapter 4, Figure 2, in the case of biometric passport systems, there is not just one government entity present. With biometric passports, biometric identifiers and their storage must also be considered. Whenever the holder of a biometric passport wishes to travel to another country, they are immediately forced to trust *the information systems of the destination country* in addition to those of their own.

5.4 BIOMETRIC PASSPORT SECURITY – 5 YEARS LATER

It can safely be said that the potential security threats that have been discussed here in this chapter have not been realized in the last 5 years after they were previously discussed by Heimo, Hakkala & Kimppa [10]. Biometric passports have become commonplace, and automatic passport checking at airports is quite mundane, and there have not been – at least publicly – any incidents with biometric data stolen or misused in any manner, whether stored in biometric passport databases or not.

Biometric passports and their security has progressed from the time when the analysis in this chapter was done in 2012. According to notes in the ICAO standard, the more insecure security features such as BAC are finally being phased out starting from 2018 [134]. This alone would warrant a new examination of the security situation for biometric passports. Avoine *et al.* [147] have made an extensive survey of vulnerabilities in existing biometric passports, detailing many of the vulnerabilities previously discussed in this chapter. The security situation of biometric passports in 2017 and beyond, especially taking into account the forced trust aspect and societal issues discussed in this thesis, is a potential avenue for new research in the future.

The security environment is currently one where biometric identification data is not a very valuable commodity. It is used in some authentication applications, but widespread popularity has still evaded biometric recognition. Smartphones and other personal computing devices have become perhaps one of the main target applications for biometric recognition, by virtue of fingerprint or face recognition in unlocking the device. The con-

venience of using biometrics in this manner makes it a killer application in this area, but does not make biometric data sufficiently valuable in itself to make it a target for criminals to start breaking into biometric passports. Indeed, the threat to biometric passport data comes more directly from nation state level actors than from individual hackers and criminals. So far there exists no way to effectively monetize fingerprint data, thus removing it from the scope of mundane criminals. The underlying issues with biometric data have not been removed or alleviated, however. The same risks that were identified five years ago are still relevant in discussion on security and privacy of biometric passport users.

5.4.1 *New threats to biometric recognition systems*

In addition to previously identified risks, new ways to misuse biometric information have surfaced. With the adaption of fingerprint recognition in smartphones and other mobile devices, fingerprints now have, for the first time, monetary value. This extends beyond merely making it possible to steal smart devices by acquiring the fingerprint of its owner and replaying it back to the device. The current main threat is the use of mobile device fingerprint identification systems as strong authentication in online banking.

In early 2016, Nordea, a bank operating in the Nordic countries, begun offering its customers access to their online banking with Apple's TouchID. It is the fingerprint recognition system for Apple's iPhone.[9] After a while, competitors such as OP have adopted the same method for accessing online banking systems with fingerprint recognition.[10]

9 Nordea Oy "Mobiilipankki". Online, available at http://www.nordea.fi/henkiloasiakkaat/paivittaiset-raha-asiat/internet-mobiili-ja-puhelinpalvelut/mobiilipankki.html Accessed 26.1.2016.

10 OP Osuuskunta. "Uutta OP-mobiilissa: verkkoviestit ja tunnistautuminen sormenjäljen avulla". Online, available at https://www.op.fi/op/henkiloasiakkaat/tilit-ja-maksut/uutta-op-mobiilissa--verkkoviestit-ja-tunnistautuminen-sormenjaljen-avulla?cid=151858901&srcpl=3 Accessed 7.5.2017.

Apple's TouchID system for the iPhone can be fooled with a fake fingerprint.[11] If TouchID can be used as an authentication mechanism for financial transactions, and fooling the sensor with a relatively simple fake fingerprint is possible, this causes a scenario that has not existed before; an easy way for attackers to monetize biometric data.

While using the fingerprint reader undoubtedly provides a better user experience and more streamlined access to online banking, one clear threat scenario arises from this development. The biometric recognition system works even when the user is incapacitated. If an attacker is able to gain access to the victim's phone when he is incapacitated (one obvious example would be due to excessive consumption of alcohol or other narcotic substances), it is possible to access their banking details without the knowledge of the victim. Biometric recognition does usually not require coherence or awareness from the target. Some systems do measure liveness signals, making sure that the input is from a living human being, instead of an inanimate replica, but there is no such check for awareness. As was previously noted in this chapter, biometric recognition can be covert, designed to work without consent of the target. This can be seen as a vulnerability in this kind of biometric recognition systems.

Iris recognition (see e. g. [148, 149, 150] is also used in modern smartphones for user authentication. For example, Samsung Galaxy S8, published in 2017, has the option to use iris recognition for unlocking the phone. This system has already been shown to be vulnerable. An attack demonstrated by the Chaos Computing Club[12] is capable of fooling the recognition system by using a photograph of the user, a printer with sufficient print quality, and a contact lens.

Vulnerabilities in emerging technologies and implementations are expected. The issue with using biometric authentication in a wide scope is that it leads to biometric identifiers becoming a commodity to be stolen, sold and misused. Up until now the use

11 Chaos Computing Club. "Chaos Computer Club breaks Apple TouchID" Online, available at https://www.ccc.de/en/updates/2013/ccc-breaks-apple-touchid Accessed 26.4.2017.

12 Chaos Computing Club. "Chaos Computer Clubs breaks iris recognition system of the Samsung Galaxy S8". Online, available at https://www.ccc.de/en/updates/2017/iriden Accessed 6.6.2017.

of biometrics has been marginal from the point of view of an average user. With the advent of biometric recognition in commonplace devices and systems, the threat scenarios that researchers have been discussion for years are closer to reality than ever. Further research on the effect of commoditization of biometric identifiers is clearly needed. This is left for future work in the field.

5.5 CONCLUSION

Biometric passport systems are a part of critical infrastructure of the networked information society. The process that they facilitate — crossing borders — is equally important to society, as mobility of people is necessary for a society to function. When examined from the point of view of mass surveillance and forced trust, biometric passport systems bring with themselves multiple issues that pose a credible yet unrealized threat to the security and privacy of citizens in the networked information society.

When the biometric identifiers of an individual are placed on the passport chip, they are readable to the outside world by definition. The security measures in place can only protect the information so far. One of the widely accepted axioms of information security is that if an attacker gains physical access to your system, the game is over. The passport is a small item, easily lost and easily pilfered (and later returned, if the situation warrants it) by an attacker. Therefore it is questionable whether using biometrics in this manner is sustainable or even recommendable from the privacy point of view.

The addition of biometric databases to the equation makes it even harder to trust a government with an individuals biometric data. The original intent of such a database has been to prevent a person from gaining multiple identities by performing a match of the applicants fingerprints against the database. Using biometrics for this purpose of catching duplicates is easy and requires significantly less work than manual checking. The problem is that biometric databases can be used for other purposes than originally intended, and that governments are more than willing to extend their investigative powers to such data.

To keep these threats from becoming reality, extreme care must be taken in increasing the use of biometric data in CGISs, as the

inherent problems of biometrics are yet unsolved, and combining them with issues of forced trust on an unsafe communications infrastructure provides ample opportunity for problems to manifest and threaten the security and stability of society.

6

SOCIETAL INTERACTION AND DECISION MAKING

An Outside Context Problem was the sort of thing most civilisations encountered just once, and which they tended to encounter rather in the same way a sentence encountered a full stop.

Iain M. Banks, *Excession (1996)*

They screwed with democracy. I'm not going to lose much sleep when democracy screws them back.

–Deputy Field Prefect Bancal
Alastair Reynolds, *The Prefect (2007)*

In a networked information society where surveillance is commonplace and the infrastructure itself is not trustworthy, societal interaction in a secure and trusted manner is challenging to implement. In this chapter we specifically focus on collective decision making in the form of voting. We analyze the potential hazards of using online voting systems for collective decision making in the networked information society, taking into account the problems with the underlying infrastructure and issues with forced trust we have previously identified in this thesis.

Old threat models have become obsolete in the wake of the revealed capabilities of agencies performing network surveillance, as discussed in Chapter 3. The adversaries have more capabilities than what has previously thought to be possible. This gives rise to a new set of requirements and possible restrictions on the use of online voting in the networked information society. These issues should be considered when arranging elections, from small organizations to national level.

6.1 VOTING

One of the fundamental processes in any society is collective decision making. Whether the decision is about who will be elected the leader of a group, what are to be the rules governing a society or what stance to collectively take on a particular matter, a common method for facilitating this process is voting in an election. This makes the process of arranging an election an important task to get right. The legitimacy of a decision by vote relies upon all parties to the process accepting the majority decision to be binding, and for that to happen the participants must be able to trust the process. Concisely, without trust in the election process, there are no legitimate elections, and societal problems tend to follow the lack of trustworthy democratic decision making.

Because of this pivotal role, elections have been throughout history the subject of much controversy, attention and contention. The goal of candidates is to win the election, and sometimes the pull of power has been so strong as to invite foul play in some form or another from candidates. Of course, the entire election can be rigged from the start, which would imply that the election officials responsible for arranging the elections are corrupt. Real-life examples of such blatant election fraud are unfortunately quite abundant. As long as there is perceived power and benefit to be won in an election, voting fraud is a potential event that can happen. The finer nuances of election fraud are out of the scope of this thesis, we merely acknowledge that election fraud indeed does happen and the motivation to do so for whatever purposes exists whenever there is a sufficient amount of power to be shared — or seized.

In this context it is natural that different techniques, technologies and processes for voting have evolved to mitigate or even eliminate the possibility of foul play, and thus bolstering the legitimacy of elections. The right to participate in government through free and secret elections is a basic human right, defined in Article 21 of the Universal Declaration of Human Rights.[1]

1 The United Nations, The Universal Declaration of Human Rights, https://www.un.org/en/documents/udhr/index.shtml#a21

6.2 VOTING SYSTEMS

The implementation of an election can be done in various ways. It can be a show of hands in a crowd, a verbal casting of votes, paper ballots, or some kind of electronic election system. In this thesis we focus on electronic voting systems that are implemented over the Internet, i. e. Internet voting. For discussion on other of electronic voting systems and their issues, see e. g. the work by Mercuri [151].

In contrast to traditional pen-and-paper ballot voting, electronic voting uses electronic systems in different capacities during an election. The systems can be differentiated between those which use automated systems for vote tallying and those where the entire process is entirely electronic, from vote casting to tallying. Examples of automatic vote tallying systems are punch cards and optical scan ballot systems [152]. Direct Recording Electric (DRE) voting machines [153] and different Internet voting protocols[2] are examples of systems with fully electronic vote casting and tallying.

Numerous accounts of problems, irregularities and issues associated with electronic voting machines have been reported in literature. To give some examples, Mercuri [151] discusses the security issues of electronic voting systems in depth. Kohno *et al.* [154] assess issues with DRE voting systems used in US elections. The system audit for the 2008 Finnish municipality elections [155] found serious weaknesses in the new electronic voting system used in the elections.

In this chapter we focus on voting systems and protocols designed to be used over the Internet. While any voting system with fully electronic vote casting and tallying could be relevant to this discussion, the use of Internet for voting is most interesting by far due to the forced trust issues on the Internet infrastructure and mass Internet surveillance.

With Internet voting, the voter cannot in any way verify the system used. The average citizen has very little experience with the kind of complex cryptography that election protocols require, and thus are forced to trust the election system provider in providing a robust voting system and also trust the officials

2 See for example Helios. https://vote.heliosvoting.org/

completely to fulfill their duties in providing an impartial election.

This trust relationship has been examined in literature. Nestås and Hole [100] claim that simply having a secure Internet voting system is not enough but that all stakeholders must also have trust in the system. This trust does not emerge spontaneously but must be built and maintained. This is more easily achieved — if at all possible in all cases — in organizations of small to medium size than in for example national elections. One can assume that an organization is centered on a certain theme or issue, or to represent its members in a collective manner, thus making it easier for people to trust each other not to try any foul play. This does not, however, take into account external threats to the election. And in the case of governmental elections, this trust relationship becomes one of forced trust on an infrastructure and system that the users are incapable of understanding, and for which they have no control over.

6.2.1 *Voting system models*

To analyze the effects of infrastructural weaknesses and forced trust on voting systems, a model of a generic voting scheme is required. Numerous models have been presented in literature, but the basic principles of voting are relatively universal.

Sampigethaya and Poovendran [156] have presented a generic model for voting, in which four distinct entities are identified: *Voter, Authority, Candidate and Adversary*. A *voter* is an entity that is eligible to vote in an election. An *authority* is the party responsible for arranging the election. An election has a set of *candidates*, to which voters give their votes. An *adversary* is a malicious entity that attempts to manipulate the election. Adversaries are further divided into *internal* and *external* adversaries. An external adversary may attempt to affect the results of an election by coercing or bribing voters and also may attempt to compromise the privacy of voters. An internal adversary may similarly attempt to compromise the privacy of voters, and additionally may aim to corrupt or manipulate the voting tally.

The model includes a generic voting scheme consisting of five phases, shown in Figure 7 [156]. The first step in elections is the announcement. Additionally this contains agreement of used

protocols, eligible voters, etc. The next step is registration , which contains identification of voters by the election authorities, and potential pre-voting for candidates already at this stage. The next step is the actual voting event, where ballots are cast. The ballots, whether paper or electronic, may contain additional information related to the election in addition to the vote. The last step is the tallying of votes, where the results are tabulated, and the final result of the election is announced. These steps are performed in this order. The fifth step is verification, which can be taken at any time between registration and tallying of votes.

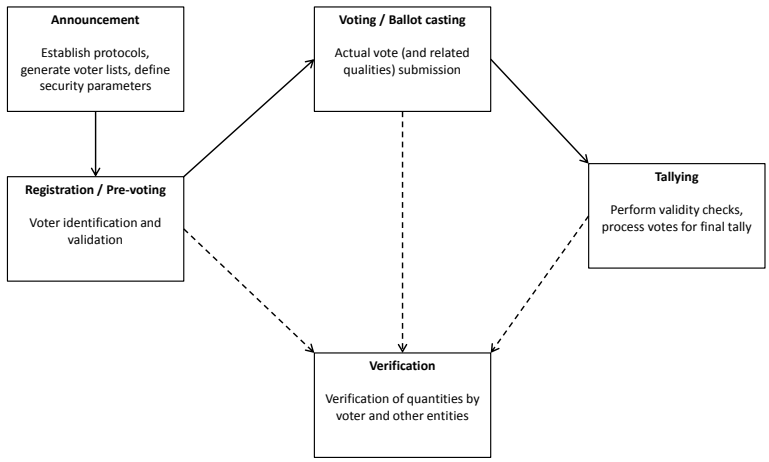

Figure 7: Description of a generic voting scheme. Adapted from [156].

6.2.2 *Voting security frameworks and taxonomies*

Various frameworks and taxonomies for analyzing, classifying, and comparing voting schemes have been presented in literature. A short survey of previous work is detailed here, and summarized in Table 2.

Gritzalis [157] discusses the principles and requirements for a secure e-voting system. He proposes a set of requirements for e-voting systems that are further divided into functional and

non-functional requirements, but does not present an adversarial model.

Lambrinoudakis *et al.* [158] likewise provide security requirements for electronic voting systems and analyze a set of proposed protocols according to their definitions, but they do not provide an adversarial model.

Sampigethaya and Poovendran [156] have defined a framework and taxonomy for evaluating electronic voting security. In their work they analyze more than 25 previously presented schemes and find that none fulfill all the security requirements outlined in their framework. They do not, however, provide a clear adversary model.

Çetinkaya [159] presents an improved framework based on the work of Sampigethaya & Poovendran [156], but similarly without a clear adversary model.

Langer *et al.* have examined [160, 161] different threat models, security requirements and adversaries for electronic voting systems in their taxonomy for electronic voting systems. Their work also includes an extensive adversary model. Langer [162] has further refined her taxonomy with an improved adversary model.

Yumeng *et al.* [163] review general requirements and implementation options for secure electronic voting systems, but do not provide an adversary model.

Neumann and Volkamer [164] develop a taxonomy that focuses on Internet voting systems. They derive technical and functional requirements for a safe Internet voting system. They also provide an adversary model that considers computationally unbound adversaries, but without the capability for short-term cryptography compromise.

The threat model Springall *et al.* [165, pp. 709–710] use to analyze the Estonian Internet voting system takes state level adversaries into account. They note that cyber warfare has become reality, and that the election infrastructure of a nation would be a tempting target for such operations. They also cite an example of such an attack [167], where attackers with reported ties to Russia succeeded in crippling Ukraine's voting infrastructure before the presidential election in 2014. Springall *et al.*do not present a clear adversary model, but allow advanced but realistic capabilities for adversaries when attempting to compromise an election

Table 2: Adversary models in Internet voting. Legend: AM= Adversarial Model; CS=Computational Security; CC=Cryptography Corruption capability; STCC= Short-Term Cryptography Corruption capability.

Paper	AM	CS	CC	STCC
Gritzalis, 2002 [157]	No	No	No	No
Lambrinoudakis *et al.*, 2003 [158]	No	No	No	No
Sampigethaya & Poovendran, 2006 [156]	No	Yes	No	No
Çetinkaya, 2007 [159]	No	Yes	No	No
Langer *et al.*, 2010 [160]	Yes	Yes	No	No
Langer *et al.*, 2010 [161]	Yes	Yes	No	No
Langer, 2010 [162]	Yes	Yes	Yes	No
Yumeng *et al.*, 2012 [163]	No	No	No	No
Neumann & Volkamer, 2014 [164]	Yes	Yes	No	No
Springall *et al.*, 2014 [165]	(Yes)	(Yes)	(Yes)	No
Neumann *et al.*, 2016 [166]	Yes	Yes	No	No

system. This includes detailed information on the inner workings of the voting system and sufficient reverse engineering capabilities to make any reverse engineering efforts practical in a short amount of time. They also assume the capability to send targeted malware to voters' computers as client-side attacks, and the possibility to compromise voting client before it is delivered to the voters.

Neumann *et al.* [166] extend and improve the work by Neumann and Volkamer [164], but still do not consider adversaries that are capable of compromising cryptography protocols used in Internet voting platforms, either over a long or short time period.

6.2.3 *More accurate adversarial capabilities*

Existing electronic voting frameworks discussed above generally assume — with few exceptions — that the lowest layer of communication infrastructure in itself is inviolable or at least not practically exploitable in any wider sense. This is indicated by

having adversarial capabilities limited to controlling the communications channel, intercepting messages or identifying communicating parties. The Dolev-Yao adversary model [168] which assumes a perfect public-key system (unbreakable one-way functions, tamper-proof public directory, no leak of secret keys) is often used as a basis for this kind of models. Indeed, this adversary model was the basis of all presented models in Table 2, with the exception of the model presented by Langer [162].

Langer [162, p. 48] presents the following adversary capability regarding cryptography:

IIIA The adversary is able to break any cryptography which provides only computational security.

This provides the adversary capability to compromise computationally secure cryptographic algorithms. This consistently interpreted in literature to mean than this break happens in a distant future, and that currently computationally secure algorithms are not compromised. When we consider the capabilities described in Chapter 3 regarding surveillance and cyber warfare, the above assumptions can no longer be considered sufficient. Thus when we consider using frameworks for classifying and analyzing new Internet voting solutions, we must choose from those which are capable of modeling realistic adversaries. Langer [162, p. 48] is among the few who have seriously considered the potential adversarial capability to compromise existing cryptography. Both new and improved versions of old frameworks should also take this realistic adversarial capability into consideration.

Even those frameworks that explicitly consider the possibility of an adversary simply breaking cryptography assume that this will happen over a long time period with a computationally unrestricted adversary. Based on previous observations in this thesis, in election scenarios of sufficient magnitude — national elections, for example — it is reasonable to assume the existence of an adversary with the short-term capability to compromise cryptographically protected assets in the election. VPN, for example, can be used to provide end-to-end security over an insecure Internet connection. It is possible for a sufficiently advanced adversary to compromise the security assumptions of VPN [90], as was previously discussed in Chapter 3.

6.2.4 *Power, the chilling effect, and elections*

Trust is an essential part of a functioning society, and this is exemplified by elections. There exists a collective trust threshold that must be passed for the majority of a society to consider any elections to be legitimate. There is no actual threshold value of trust – in itself difficult to define – for everyone; each member of society has their own perception of what is a fair and legitimate election. This threshold is seldom zero: there will always be dissenters who will not trust elections to be fair, regardless of who wins. Further discussion on the reasons for this is outside the scope of this thesis, we shall only observe that it is extremely unlikely to get 100% approval for elections.

Shifting election platforms towards Internet voting creates an unnecessary layer of forced trust: the Internet voting system can be designed according to best practices and known security assumptions and parameters, but the trust on the platform is still forced from the point of view of the citizen. Now, when new adversary capabilities are identified that are able to compromise key security assumptions of a secure communications channel, this trust is further cast into doubt.

The chilling effect was previously discussed in Chapter 2 in the context of mass surveillance. Voting is by nature anonymous to make certain that all opinions are heard and there is no possible way to pressure voters to vote in a certain way. If voters are uncertain whether their votes are truly anonymous, it can lead to reduction in the legitimacy and trustworthiness of the election, and also reduction in the likelihood of an individual to vote. The political views of a person are commonly agreed to be sensitive personal information, as persecution due to political opinion is unfortunately commonplace.

In order to achieve the capabilities described above, an adversary must be an exceptionally powerful entity — a nation state level actor for example. Indeed, allegations[3] have been made of intelligence agencies and hackers attempting to affect elections,

3 US Department of Homeland Security. "Joint Statement from the Department Of Homeland Security and Office of the Director of National Intelligence on Election Security" Online, available at https://www.dhs.gov/news/2016/10/07/joint-statement-department-homeland-security-and-office-director-national Accessed 13.6.2017.

and the scale of these alleged operations is vast. Electronic and Internet voting systems are targeted by adversaries in elections where the potential gain is worth expending vast intelligence resources.

It is tempting to assume that only elections of significant scale are subject to attention and influence from third parties, or that elections in general are not used for gathering information on those who vote. This is not necessarily true, as elections in organizations and associations may also be targeted for the purpose of gathering information on voting behavior of interesting subjects. This can be true for more frivolous voting instances as well. Azerbaijani security officials have reportedly interrogated[4] citizens who voted for the Armenian candidate in the 2009 Eurovision song contest, citing concerns of national security.

The Eurovision Song Contest is deadly serious to some people, though.

The Azerbaijan case testifies against the idea that in certain marginal cases where there is not much at stake, such as the win of Eurovision, Internet voting poses little to no risk with added convenience. What can be inferred from voting behavior in even such frivolous cases can be extrapolated to mean something completely different when combined with other information gathered from an individual, further complementing their data double as discussed in Chapter 2. Thus using any kind of Internet voting system, even in frivolous cases, *can* leak information on those who use the system.

6.3 CONCLUSIONS

Internet voting is an essential part of societal interaction that is implemented on an insecure platform. Issues with infrastructural security, platform security, reliability of algorithms, and potential attempts to affect the end-result of elections must all be taken into account when making decisions on whether to use Internet voting systems or not. In the context of network surveillance and the Networked Information Society, as the significance of the election grows so does the probability of potential outside influence, and some of the factions that attempt to exert their in-

4 Sean Michaels *The Guardian*, 18.8.2009. "Azerbaijan authorities interrogate music fans in Eurovision probe", Online, available at http://www.theguardian.com/music/2009/aug/18/azerbaijan-authorities-interrogate-music-fans Accessed 13.6.2017.

fluence may have the capabilities described in Chapter 3 at their disposal.

Assumptions of information-theoretic privacy are not in danger in light of the revealed capabilities of Internet surveillance agencies. But if the integrity of elections rely on the assumption of computational security and privacy, however, this may lead into the situation where this trust is simply misplaced. Future work in this field will show what is the extent of potential damage to Internet voting systems from the weakening of fundamental building blocks of the networked information society.

The discussion in this chapter is technologically not in-depth, rather observing the potential capabilities of adversaries in modern Internet voting systems. Future work should be directed into examining whether the advanced attacks and methods for Internet surveillance described in Chapter 3 are feasible in real-life scenarios. The potential is undeniable, but as the security and accuracy of elections is not perfect with any existing system, security compromises are always made in the end.

Part III

SECURITY ENGINEERING SOLUTIONS FOR FUTURE INTERNET

7

AUTHENTICATION, IDENTIFICATION AND TRUST BUILDING

Humans ought to have their lives, happiness, and autonomy protected. And this principle of justice —- the protection of fundamental human rights — should guide us in shaping ethical policies for using computer technology.

– James H. Moor (1999)

So, the combination is 12345? That's the stupidest combination I've ever heard in my life! That's a kind of thing an idiot would have on his luggage!

– Dark Helmet
Spaceballs (1987)

In this chapter research-based suggestions for improving security and privacy of Internet users all around the world are presented. The suggested measures take both concrete and abstract forms. We begin by applying a data ownership and control model to mass surveillance records, with the goal to provide a justifiable and ethically sound way to gather, store and process surveillance records. Next, we present empirical research both on human-chosen passwords in online services and password behavior of university students. We especially focus on the linguistic properties of passwords that contain natural language, and especially how people use natural language in human-chosen passwords.

7.1 OWNERSHIP OF DATA AND MASS SURVEILLANCE

In this section we examine the problems that arise when data ownership gets into conflict with mass surveillance. As a solu-

tion, we propose applying a concept borrowed from discussion on ethical management of patient information to the problem of protecting the rights and privacy of citizens in the networked information society, while making it possible to conduct lawful surveillance.

Ownership of data has been extensively discussed in literature from the business side. Data is vital to most, if not all, businesses. Referring to previous work by Davenport *et al.* [169] and Strassmann [170], Redman crystallizes issues of data ownership in business quite succinctly [171]: *"Indeed, the politics of data ownership are among the most brutal in many enterprises."* Data can be so critical to businesses that the whole existence of a company can depend on who owns a certain piece of intellectual property.

We argue that it is not unreasonable to extend this characterization to data ownership relationships between organizations and customers, or governments and citizens. A common quip about the revenue model of "free" services is that "if you are not paying anything for your service, you are not the customer, you are the product being sold". As a real world example, questions of data ownership and the use of user-generated data for profit have been at the heart of the discussion concerning Facebook since its founding in 2004. The questions of data privacy, data ownership and what can be done with collected data are central questions for the information society. Intellectual property rights are paramount to a society based on information, but in addition all other collected data, whether originating from industrial or business processes, or people and their behavior, is equally valuable. The Whatsapp instant messaging application has over one billion users,[1] and a significant part of the value of the company is derived from this mass of users. Information about users is seen as a commodity, and it is used to generate business in an information society.

Now consider the information that is gathered with systematic mass surveillance. This contains data about all aspects of the lives of its targets, from their personal communications and online behavior to location information and personal connections, and even personal health information. The latter can be argued to be clearly within the realm of acceptable privacy. If patient

1 Whatsapp blog, "One billion". 1.2.2016 Online, available at https://blog.whatsapp.com/616/One-billion

records are considered to be such critical information as to warrant extreme protections, isn't it also justifiable that other information of comparable compromising potential should also be afforded the same protection? There is no silver bullet solution – at least not an obvious one – for the problems that arise from these issues. Next we will examine one potential framework that could be used as a starting point in solving this problem in an ethically justifiable manner.

7.1.1 *Datenherrschaft - Mastery over data*

Mass surveillance generates massive amounts of data. From the point of view of targets of surveillance, that data is derived form their persons and activities, and contains significant amounts of personally identifiable information. Here we have an opportunity to examine the ethical dimensions of mass surveillance, and approach it through ownership of data. We use the concept of *Datenherrschaft* together with mass surveillance in order to explore possible ethically sound foundations for surveillance.

The definition for Datenherrschaft is given by Kainu & Koskinen [172, p. 54] as:

> "the legal right to decide the uses of, and continuing existence of, in a database or another compilation, collection or other container or form of data, over a entry, data point or points or any other expression or form of information that an entity has, regardless of whether they possess said information, with the assumption that sufficient access to justice is implemented for a citizen to have this power upheld in a court of law."

The difference between Datenherrshcaft – mastery over data – and ownership of property rights can be condensed to four key differences. First, Datenherrshcaft is *nontransferable* [11, p. 102]. While ownership of data can be transferred to another, the mastery over data is permanently bestowed. Koskinen also notes that the mastery over information *can only be given to the person from whom the information is derived from*.

Second, Datenherrschaft is *not a compensation or a reward* [11, p. 102]. Koskinen observes that while work done by someone

is seen as a justification for granting that someone immaterial property rights to the result of that work, in some contexts this is not applicable. In the context of healthcare, the compensation comes in the form of salary, so there is no need to compensate the person who compiles data about the health of a patient into patient information with ownership of that immaterial property.

Third, Datenherrschaft is *bestowed by default to the person from whom the data or information is derived* [11, p. 103]. It is commonplace to pass on intellectual property rights to parties that have done no actual intellectual work in creating the protected information. For example, in Finland it is possible to transfer intellectual property rights of an invention made by an employee to the employer. In contrast, Datenherrschaft is inherently a part of the originator of the data [172], and thus cannot be transferred to another entity.

The fourth and final difference is that while intellectual property rights protect works based on an artistic process and effort, Datenherrschaft can rather be used to *protect work based on facts and derived through scientific methods, instead of an artistic or creative process* [11, p. 103].

7.1.2 *Ethical data-driven surveillance*

Datenherrschaft has been previously used by Koskinen [11] to examine and outline ethically sound and justified approaches to the ownership of patient information. They contain significant amounts of data on a person and their private matters. Patient information should thus be seen as sensitive, and that it needs explicit protection from misuse and unnecessary observation by others.

While questions of ownership in regard to patient records and health information have already been discussed in literature, surveillance has not been examined in similar manner from the perspective of Datenherrschaft. This work is the first of its kind, opening the discussion on the applicability of Datenherrschaft to surveillance data.

Concepts such as the Surveillant assemblage and Rhizomatic surveillance, previously discussed in Chapter 2, give us tools to observe how omnipresent surveillance works and how all aspects of society are being harnessed as sources of data for surveil-

lance. Datenherrschaft gives us a framework through which we can observe surveillance in an ethically justifiable manner, and gives the potential to create practices that are better in protecting the individual citizen of the networked information society.

Koskinen, Kainu & Kimppa [173] illustrate the problems with patient information ownership. They observe that traditionally the ownership of intangible things such as information is governed through intellectual property rights. Patient information, however, is not suitable for being governed by the same legal framework. Even though patient information is immaterial, it is still irrevocably bound to the source of the original information – the patient.

Now, let us consider the data that is used as the basis of mass surveillance. This, similarly to patient information, is immaterial information that is bound to the originating source – the person. This data is also used to make surveillance related decisions that can have an effect on the person on various levels, even including their health. This parallel that can be drawn between patient information and data that is used to conduct surveillance serves as the basis of the argument that Datenherrschaft can also be used as a solution in protecting citizens of the networked information society from undue violations of privacy.

In an ideal world, the information that is gathered with the express purpose of tracking people should be governed by the principles of Datenherrschaft. The person from whom the information is directly derived should have some kind of control over that information, even though it is used by other parties. Indeed, the concept of the data double (see Section 2.4.3) gives some credibility to the demand for application of Datenherrschaft. If information that is, in the virtual world, considered to represent a person (or even to *be* that person), it is only natural that the person should have a degree of control over the data double.

If we examine the vastly different types of information that can be used for surveillance as described in Chapter 2, we immediately encounter the problem of defining the information to be protected by Datenherrschaft. Particularly, we should be able to identify the different kinds of data that are collected of us. In his research, Koskinen examined a particular set of information – patient records – but if we are to extend the scope of Datenherrschaft to other kinds of data, we must be able to accurately

identify the information to which the principle of mastery over data is to be applied.

7.1.3 *Datenherrschaft and metadata*

When we consider metadata from the point of view of Datenherrschaft, we immediately encounter some questions. At what point does data derived from data cease to be under the mastery of the original source of the data? Should we also own and have mastery over data about data from us? This is a central issue with applying Datenherrschaft on personal data used for surveillance.

One clear requirement for managing and combining metadata in any larger form or shape is that it cannot be identifiable or traceable to a single individual, a reasonably sized group of individuals, or a certain demographic (imagine the outcry if you were to single out black people from a mass of metadata based on some characteristic) and finally that this information cannot be reasonably deduced from the mass of data. Anonymization of data, in other words, is a clear prerequisite when any larger sets of personally identifiable data are processed.

Can we infer from this societal requirement – it is societal, as it stems from legislation and the perceived need for such – that we should really also have some mastery over surveillance metadata about us? We clearly continue to exert influence, albeit non-personalized, on how data about us is handled. An obvious counter-argument is that surveillance is by necessity a covert operation and that to give the subject or subjects of surveillance a notification let alone some actual power over this surveillance process would be ridiculous and counter-productive.

We come into apparent conflict with Datenherrschaft, which argues that the individual should have the right to decide how their data is used. This conflict can be resolved with an inherent property of Datenherrschaft – mastery over ones data confers not only rights, but also responsibilities. There exist justifiable situations where the Datenherrschaft of a person over their own data can be overridden. Koskinen, Kainu & Kimppa [173] outline that prioritizing life, health and liberty over possessions is justifiable, and thus Datenherrschaft can be superseded in these cases.

Koskinen [11] clearly states that Datenherrschaft exists in a legal framework where it is possible for the rights of other people to supersede and override an individual's mastery over their own data. Similarly in the case of constitutional republics, the constitution allows citizens certain constitutional rights. These are then in turn limited in some justified cases, and the guideline for doing this usually is to do any limitations in as non-intrusive manner as possible. These situations should be codified separately and with the principle of least interference.

This would mean that a person cannot exercise their Datenherrschaft to stop the use of their data in lawful use. Investigations that are conducted according to legal guidelines certainly fulfill this criterion, but now we face another problem: what should be considered as lawful use, i. e. how should the law be written? This is clearly out of the scope of this thesis, but the observations made in Chapter 2 give us a starting point when considering the negative effects of surveillance. Balancing these with the rights of the citizen is a difficult task, but Datenherrschaft has the potential to be the underlying guideline that facilitates ethically acceptable and justifiable choices.

As we have discussed in Chapter 2, surveillance is performed by both public and private instances: companies such as Facebook and Google have made their entire business dependent on practices that approach (or practically are) surveillance. Should Datenherrschaft be considered in a similar manner as a constitutional right that can be limited only in certain, legally justified cases? Should it cover all information that is derived from our persons, or should it have some, more precise limitations? In the case of surveillance, could this Datenherrschaft be used to provide ground rules for surveillance and tracking in general? Answering these questions is left for future work.

7.1.4 *Final remarks and future work on Datenherrschaft*

A key observation made by Koskinen [11] about Datenherrschaft is that it seems to be intuitively suitable for also other kinds of private information rather than just patient records. The idea of applying Datenherrschaft to other kinds of data is thus tempting. Surveillance, as was noted earlier in this thesis, has its own inherent problems with privacy and freedom of speech, and thus ap-

proaches for ethical solutions to managing surveillance and the vast data associated with it in a morally and ethically justifiable manner are quite desirable. Datenherrschaft has the potential to be the ethically sustainable solution to this difficult problem and the initial step that is desperately needed in managing person-derived data. Further research is required before this potential can be realized.

Future work in this area should include a comprehensive analysis on the applicability of Datenherrschaft to the case of lawful interception and surveillance. The nature of personal and privacy-sensitive information should be analyzed in a similar manner as Koskinen [11] has done to patient and health information, through the ethical theories of Locke, Kant, Heidegger, and Rawls. This could shed some light on whether the perceived suitability of Datenherrschaft to this problem is merely an illusion, or if there is a real and robust foundation on which to build the future of the networked information society. The Internet does not forget, but with information that is used to track and monitor us, we should have mastery over what is remembered.

7.2 BAD USER PASSWORDS — TURNING THE TABLES ON TRUST

In this section we discuss password security and how it relates to trust. To begin, consider forced trust relationships discussed in Chapter 4 from the point of view of an administrator of an information system. In previous discussion it was pointed out that users have a very unfortunate position when considering forced trust in systems and infrastructure, but the trust relationship between administrators and users is also two-sided. Administrators are also forced to trust the users to behave in a proper and secure manner. If a user does something to jeopardize the safe and correct operation of the system, it is the administrator's duty to intervene and remove use privileges from that user. For example, a user could use poorly chosen passwords, making password guessing attacks against the system feasible.

If the system is deemed to be of such societal importance that using it is a right instead of a privilege, that makes the situation a lot more interesting. In such cases suspension of user rights would be temporary at best, and the administrators would have no option of banning a user with poor security performance his-

tory. That is, the administrators are forced to trust the users to behave in a secure manner and are thus in a position of forced trust, as was observed in Chapter 4.

Systematic attacks on websites and applications performed by malicious users, ranging from hobbyists to organized criminal gangs, place an increasing amount of weight on password security. This is asymmetric to the ability of an average user to create and maintain effective passwords, often leading to bad password practices such as using simple dictionary words and reusing passwords across information systems. In the current networked environment where the average user can have accounts in dozens of Internet services, all of them requesting the user to generate a new password, the temptation to reuse passwords is understandably hard to resist.

In order to illustrate the seriousness of the forced trust scenario regarding passwords, we focus research on two aspects of password security. First we examine the prevalence of natural language in passwords through an analysis of Finnish web user passwords that have been leaked to the Internet. There have been a number of significant password leaks from Finnish web sites in the last 10 years, and we use some of these leaked password sets in our analysis. We use Natural Language Processing (NLP) tools that have been developed for the Finnish language in this analysis. The goal is to gain insight into what kind of natural language users tend to use in their passwords. This will improve general understanding of password security, and also provides an opportunity to build better systems that are capable of withstanding attacks and taking into account the human factor in passwords.

Second, we study university students' password behavior with an anonymous online survey. We gathered information on the use, creation and retention policies of students participating on a course on information security and society. The course is targeted for freshman and sophomore computer engineering and computer science students. Analysis of the results of this survey is provided later in this section. Together these two empirical analyses provide an interesting starting point for the discussion on forced trust and how far you can trust the user of an information system, even when (and *especially* when) you have no other options available.

7.2.1 *Passwords in authentication ecosystems*

Users are asked to create a new password almost every time when we start using a new service. We all know what kind of an experience that is: trying to come up with a new password for a service you may not use ever again. Research on software ecosystems has shown that the majority of applications have a really low retention rate on customers. People try out software or a service, don't find it useful, and forget it – and also leave behind their freshly created account with – this is quite common, as we shall observe later in this Section – a reused password from other services.

A significant factor regarding password security is that all passwords are not created equal – or treated equal, for that matter. Users should treat their passwords to systems that contain classified company assets differently from their passwords to an online quiz service, for example. Based on the results of the survey, however, this kind of password reuse is often a very possible scenario. The worst case is that the user reuses their passwords to critical systems in other, less secure ones. This allows attackers capable of breaching a weakly protected website to use extracted passwords to breach more critical systems.

When we consider this, it is not a coincidence that people tend to use the same password over several different services. Although there are attempts to consolidate passwords under one master password or credential -– whether it is your Facebook account, a Single Sign-On (SSO) system, or a password manager – more often than not you are faced with creating a new password when you start using a new service or application.

When taking into account these realities of user password behavior, in a good scenario the choice of a password is affected by the application or service it is used to protect. Users should to choose throwaway passwords for services they use rarely, or when they perceive the service having a low security impact on their personal information. Conversely, users should choose better passwords for high security applications, but this conclusion is not supported by the fact that passwords leaked from low-security services have been used to gain access to more critical ones, meaning that users use the same passwords for different services. This gives more impact to password leaks from web

services, some of which use only borderline secure methods for storing user passwords and protecting them from attackers. In fact, one of the larger analysis datasets was reportedly stored in plain text on the breached server.[2]

7.2.1.1 *Attacking password based systems*

Password based systems can be attacked in a multitude of ways. A detailed examination of the state of the art of password cracking is outside the scope of this thesis, but we shall briefly discuss the basics relating to password attacks as a starting point.

The simplest approach is the *naïve brute-force attack*, in which the attacker guesses passwords at random until finding the correct one. This kind of attack is rarely successful, unless there are some very weak passwords in the system. A more successful approach is to use a dictionary of common user passwords, and go through them in order of popularity. Using such a word list as the basis of the attack relies on the assumption that users choose weak and often used passwords. This has been shown to be quite true in several studies; people are absolutely horrible in choosing passwords, and this is reflected in that by using a plain dictionary attack it is possible to consistently compromise roughly 25 % of passwords in a given set [109]. Even more advanced attacks exist, such as those based on different probabilistic techniques [174] and guessing strategies [175].

To protect against dictionary attacks, system administrators have adopted password generation policies. They enforce a set of rules on password generation, disqualifying passwords considered to be weak. A typical example of such a policy would be to disallow passwords of length under 8 characters, or to require the use of letters, numbers and symbols in the password. Other policies might check the password against a dictionary of known weak passwords, or even actively enforce password security by constantly attempting to crack user passwords using tools such as John the Ripper[3] or Hashcat.[4] If a user password is cracked in audit, the user is forced to change the password.

2 Iltalehti. Älypään tietoturva oli surkealla tasolla. Online, available at http://www.iltalehti.fi/digi/2010032311346474_du.shtml Accessed 23.9.2017.

3 http://www.openwall.com/john/

4 https://hashcat.net/hashcat/

Multi-word passphrases are one potential option for replacing passwords. Bonneau and Shutova [176] have previously studied the linguistic properties of multi-word passphrases. They were granted access to the Amazon PayPhrase system, which allows Amazon users to associate one or more multi-word passphrases to an Amazon account. The passphrase in this case is the sole authentication token, so there are no associated account names required for authentication. In this study, the database of user-chosen passphrases was used as an oracle, capable of answering whether a particular passphrase had been used or not. This naturally limits the scope of results, as they did not have access to the full database. They found that if users choose grammatically correct passphrases naively conforming to natural language, those passphrases are not as effective in mitigating guessing attacks as randomly chosen passphrases, but they were better than traditional passwords.

Bonneau and Shutova [176] also made some observations on the linguistic properties of passphrases. When examining passphrases with two words, the most common type of natural language phrase used was nominal modifier - noun (example: operation room) with almost 10 % share. When examining bigrams from the Google n-gram corpus [177], the most common type to succeed in matching to a passphrase were adjective-noun with 13,3 % share. These results imply that when people use natural language in passwords, they choose common and grammatically correct options.

Another common method for users to create passwords is to use keyboard patterns. Schweitzer *et al.* [178] have examined user behavior when faced with password creation policies that do not allow natural language, for example. In this case, users resort to the keyboard as a source of seemingly random passwords, but they are really just patterns on the keyboard. Schweitzer *et al.* outline a technique for visualization of password patterns and provide heuristics that allow this kind of seemingly random passwords to be attacked.

7.2.1.2 *Challenges with password research*

Passwords are a hard subject to study. A researcher unwilling to break the law in most jurisdictions has to rely upon leaked pass-

word materials found on the Internet, left behind by crackers that have compromised real services and stolen real credentials. This is morally in the gray area, but can be justified because the data is already out there. The damage is done, so why not make use of it and use it to build better and more secure systems? Those with malicious intentions do not care about these questions at all, so the point is moot in their case. Also, it is understandably hard to access password material that is in production use, or even old, obsolete password files, because usually IT operations people are keenly aware of their duties and responsibilities. Some researchers have been able to secure cooperation from large organizations, and with significant safeguards in place, have been given access to live password material. Most researchers are nevertheless forced to rely only upon sets of passwords leaked to the Internet. This is the case in this research as well, and all reservations – the question of authenticity being prime – that come with using leaked password material apply to the results here as well.

7.3 ESTIMATING THE STRENGTH OF FINNISH WEB USER PASSWORDS

In this section the security properties of passwords chosen by Finnish web users are examined. Nine separate data sets for this study have been selected, labeled DS1-DS9. All of them are password data sets that have been leaked to the Internet by hackers. All datasets are allegedly from Finnish websites and online forums, a claim that is strongly supported by a simple browsing of the password material in the datasets. While the origin of the data sets cannot be confirmed, the authenticity of the source material is believable based on the public discussion around the password leaks at the time.

The basic information about the source password sets is shown in Table 3. Most of the data sets used here have been leaked to the public by attackers that compromised a vulnerable web service and were able to acquire the whole password file in plaintext. Data sets DS1, DS3, and DS5 are, as far as we are able to

Table 3: Source data sets for password analysis.

Set	Number of passwords	Origin
DS1	55 487	Cracked
DS2	127 508	Leaked
DS3	14 606	Cracked
DS4	67 100	Leaked
DS5	6 332	Cracked
DS6	11 976	Leaked
DS7	11 945	Leaked
DS8	23 917	Leaked
DS9	16 286	Leaked

ascertain,[5] the result of password cracking, and thus may not accurately represent the full password set for that particular web service. The effects of this difference between the datasets on the results is discussed in Section 7.3.1.2.

The data sets have a combined total of 335 157 entries. This gives us a moderate number of passwords to analyze and as far as we are aware, an analysis of Finnish web user passwords in similar scope has not been performed before.

7.3.1 *Password patterns*

A password can be described as a pattern of letters, numbers and special characters [179]. In this thesis we shall use the following notation for password patterns. Letters, whether upper- or lowercase, are denoted as L, numbers as N, and special characters are denoted with S. Each symbol in a password is examined and the corresponding symbol is replaced with L, D, or S. Consequent symbols are abbreviated with a number in front, denoting the number of consequent symbols. For example, the password "password1234!" corresponds to the pattern 8L4DS. The patterns are extracted with a Python script that parses the

5 It is explicitly stated within the original dump file of DS3, for example, that the passwords have been cracked from hashes.

passwords from a text file, extracts the patterns, and writes them into a new text file.

7.3.1.1 *Results*

The patterns for password sets DS2, DS3, DS4 and DS7 have previously been analyzed by Laine [180]. We add five more datasets and analyze the patterns found in the datasets. The results are presented in Tables 4, 5, and 6. Depending on the dataset, the ten most common patterns found in the datasets account for 48.1%–63.71% of all passwords. Based on this analysis, focusing on the most common password patterns would seem to be a good attack strategy against Finnish web user passwords.

The most common patterns (6-8 letters) would also support the conclusion, that the most common passwords are natural language dictionary words. It is also quite likely that they are in nominative case, as simple dictionary attacks are at least somewhat effective against base cases of words even in the case of Finnish passwords. This was observed by Dell'Amico *et al.* [109], who found that a lowercase Finnish dictionary was eventually able to compromise 20.24% of a password set of Finnish origin. This is examined further in Section 7.3.2, where the linguistic properties of passwords are analyzed. Treemaps are used to visualize the password pattern data for two datasets, DS2 in Figure 8 and DS9 in Figure 9.

7.3.1.2 *The effects of cracked password datasets*

Three of the examined datasets are considered a result of password cracking. The significance of this, when considering the results, is that we are assessing how the composition of passwords affects the cracking difficulty of a password set. We should therefore not use data which is the result cracking to assess how easy or hard cracking is, as this would lead to us making conclusions based on biased data. In a dataset acquired through cracking it is possible, even likely, that the attackers were able to crack only a subset of the whole password set. Those passwords that were cracked are therefore weak, as resistance to cracking *is* a key characteristic of a good password. Therefore, conclusions on the quality of passwords should not be made based on the cracked datasets.

Table 4: Ten most common patterns found in datasets DS1-DS3. Legend: Pat=Pattern; #: Number of patterns found; %: Percentage from total.

DS1			DS2			DS3		
Pat	#	%	Pat	#	%	Pat	#	%
6L	6659	12	6L	16924	13,28	6L	1930	13,21
8L	4852	8,74	7L	11067	8,68	7L	1562	10,69
7L	4602	8,29	8L	10908	8,56	8L	1455	9,96
5L	3344	6,03	5L	9807	7,69	5L	980	6,71
9L	2412	4,35	5L2D	5962	4,68	6L2D	727	4,98
10L	1989	3,58	9L	5441	4,27	5L2D	721	4,94
5L2D	1664	3	6L2D	4497	3,53	9L	570	3,9
6L2D	1419	2,56	6D	4277	3,36	4L2D	567	3,88
6D	1369	2,47	4L	4064	3,19	4L4D	404	2,77
11L	1210	2,18	10L	3961	3,11	4L	389	2,66

Table 5: Ten most common patterns found in datasets DS4-DS6. Legend: Pat=Pattern; #: Number of patterns found; %: Percentage from total.

DS4			DS5			DS6		
Pat	#	%	Pat	#	%	Pat	#	%
6L	8354	12,45	6L	685	10,82	6L	1406	12,17
8L	7263	10,82	8L	538	8,5	8L	1086	9,4
7L	6039	9	7L	531	8,39	7L	996	8,62
5L	3670	5,47	5L2D	418	6,6	5L	525	4,54
6L2D	3358	5	6L2D	321	5,07	9L	517	4,47
9L	3254	4,85	9L	282	4,45	6L2D	515	4,46
5L2D	3229	4,81	10L	238	3,76	5L2D	472	4,08
10L	2473	3,69	6D	217	3,43	10L	419	3,63
6D	2021	3,01	7L2D	156	2,46	6D	378	3,27
4L2D	1795	2,68	4L2D	127	2,01	4L2D	282	2,44

Table 6: Ten most common patterns found in datasets DS7-DS9. Legend: Pat=Pattern; #: Number of patterns found; %: Percentage from total.

DS7			DS8			DS9		
Pat	#	%	Pat	#	%	Pat	#	%
6L	1817	15,21	6L	3102	12,97	6L	1304	8,01
8L	1079	9,03	8L	2120	8,86	7L3D	1275	7,83
7L	1044	8,74	7L	2044	8,55	8L	1218	7,48
5L	694	5,81	5L	1258	5,26	7L	991	6,08
6D	558	4,67	9L	983	4,11	9L	966	5,93
9L	446	3,73	6D	954	3,99	6L2D	450	2,76
6L2D	406	3,4	5L2D	841	3,52	5L2D	441	2,71
5L2D	383	3,21	6L2D	813	3,4	5L	431	2,65
4L	360	3,01	4L2D	603	2,52	10L	418	2,57
4D	326	2,73	3LD3L3D	566	2,37	6D	340	2,09

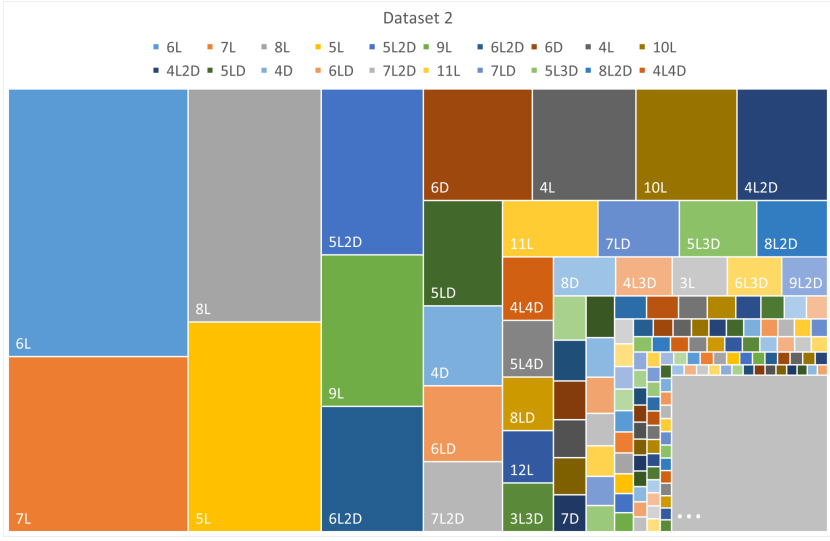

Figure 8: Treemap of patterns found in DS2.

However, in the above analysis we have observed password patterns, and they are consistent throughout all nine sets, regardless of the origin. This implies that the cracked password sets do

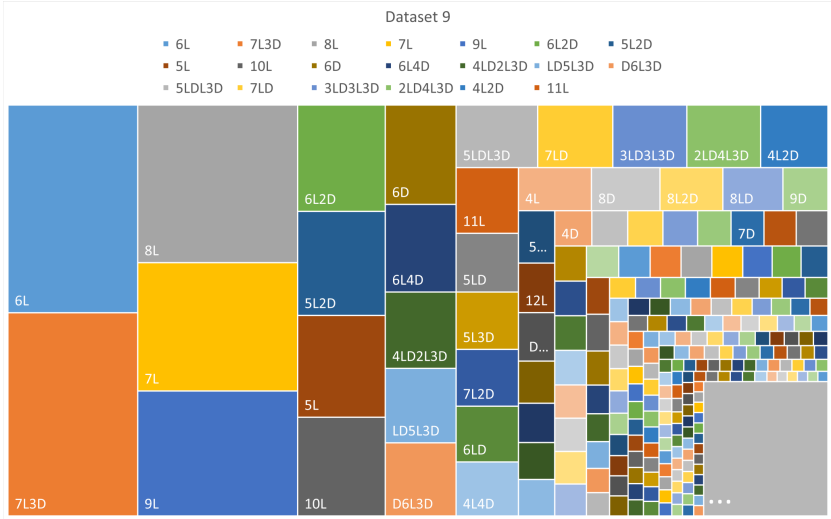

Figure 9: Treemap of patterns found in DS9.

not significantly differ in composition to whole sets. Indeed, if the three cracked password sets are omitted from the analysis, the patterns found in the remaining sets still suggest that using dictionary words is a good strategy for an attack on Finnish web passwords.

Strong passwords are not easily found via cracking, and it is entirely possible, that the cracked datasets contain only the fraction of passwords the attackers were able to compromise. We cannot ascertain *post hoc* whether the fraction of uncracked passwords was of significant size or not when considering the source data used by the attackers. What we *can* ascertain is that the pattern distributions are remarkably similar to datasets known to be leaked to the public in their entirety.[6]

Finally, we are working with data that has, to reiterate, originally been obtained illegally by hackers and published online for various unknown motives. In this context, the conclusions drawn from the source material are quite reasonable. A more robust analysis is outside the scope of what can be derived from source data of this nature.

6 For example, the poor security that made the leak of DS2 possible was publicly discussed in the national news at the time.

7.3.2 *Linguistic properties of Finnish passwords*

In addition to pattern analysis, in this study we perform a deeper assessment of the linguistic properties of Finnish web user passwords. We have chosen Finnish passwords as the target of this analysis due to multiple reasons. First, they are readily available on the Internet. Second, the Finnish language is uncommon, especially in use by others than fluent speakers — estimated to be at around 5.4 million worldwide[7] — making it quite probable that if the password contains Finnish language, the user is a native speaker. Third, Finnish is a relatively complicated language notorious for its difficult grammar and plethora of various cases and inflections due to its agglutinative nature. This gives us the opportunity to observe how users choose natural language in passwords in a language that has potential for complex structures.

7.3.2.1 *Analysis method*

For identifying natural language in passwords we use the Open Source Finnish Morphology (OMorFi) morphological analyzer tool.[8] It is built on Helsinki Finite-State Transducer Technology (HFST).[9] In this study we use the command line tool for OMorFi. It takes a string as input and returns a morphological analysis of that string if it is Finnish, and a notification if it is not.

7.3.2.2 *Preprocessing*

The raw password data is first processed to remove possible privacy sensitive information from the original password dump file. For example, user names and email addresses are removed from the files to be analyzed, leaving only the raw password data in a text file. These passwords are then parsed from the file and processed individually. For each password, all possible substrings of the password are checked for natural language, and all analy-

7 Lewis, M. Paul, Gary F. Simons, and Charles D. Fennig (eds.). 2015. Ethnologue: Languages of the World, Eighteenth edition. Dallas, Texas: SIL International. Online version: http://www.ethnologue.com/18/.

8 http://www.ling.helsinki.fi/kieliteknologia/tutkimus/omor/

9 http://www.ling.helsinki.fi/kieliteknologia/tutkimus/hfst/

sis results for a password and its substrings are stored as text in the output file, one line per password.

7.3.2.3 *Sets with cracked passwords*

The cracked datasets DS1, DS3, and DS5 do not differ significantly from the other datasets in their composition of natural language. If the cracked data sets are excluded from the analysis, the results do not change in any meaningful way. The analysis results for these three data sets are included here for the sake of completeness, but to be certain that the conclusions are unaffected by potentially biased data, they are omitted from the conclusions and discussion.

Passwords composed of dictionary words are easy to crack in a brute-force attack and therefore it is likely, that all of the dictionary word passwords in the original data were compromised. Upon visual examination, some of the passwords are strong passwords that should be resistant to brute-force cracking attacks. One possible explanation for this, in addition to pure luck on the part of the attacker, is, that the original passwords were inadequately protected in the password file, making it possible to compromise the source password data in entirety. It is impossible to ascertain whether this is the case with these datasets but given the composition of the data, it is not an unreasonable assumption.

7.3.2.4 *Results*

The main result of the analysis of the leaked password sets with OMorFi is that Finnish web users generally tend to use natural language in their passwords. In all leaked (i. e. not cracked) datasets the amount of passwords that consisted of a single natural language word or had one embedded among other characters was relatively similar, ranging from 44.5% to 54.6%. The breakdown of password compositions is shown in Table 7.

The fraction of passwords that were not identified even in part as a name or natural language was on average 25.2%. This is the part that contains "good" passwords that conform to many password creation guidelines[10] in current use, but also contains

10 Previous guidelines emphasized complexity, passwords such as *u4Vsj83%5=)*. Newer guidelines emphasize length over perceived complexity. This is a good

Table 7: Results of password analysis with Omorfi.

Set	Total	Single word	Contains word	Name	Other
DS1	*55 487*	*7 865*	*22 764*	*10 310*	*14 548*
DS2	127 508	23 498	46 137	28 963	28 910
DS3	*14 606*	*1 018*	*5 980*	*3 622*	*3 986*
DS4	67 100	11 403	23 588	17 810	14 299
DS5	*6 332*	*940*	*2 570*	*1 314*	*1 508*
DS6	11 976	1 859	4 303	3 238	2 576
DS7	11 945	1 652	4 240	2 467	3 586
DS8	23 917	2 929	8 703	3 653	8 632
DS9	16 286	1 728	5 526	1 865	7 167
Total	335 157	52 892	123 811	73 242	85 212
Total (leaked)	258 732	43 069	92 497	57 996	65 170
Total (cracked)	*76 425*	*9 823*	*31 314*	*15 246*	*20 042*

passwords such as *!qaz"wsx* or *12345*. Therefore the 25.2% fraction of good passwords by this criteria can only be considered as an upper bound.

When we further examine passwords that were composed of a single Finnish language word, we observe that they are predominantly nouns in nominative case. The results of this analysis are shown in Table 9, from where we can see that on average 95% of nouns — the clearly dominant word type — are in nominative case. This result further highlights the bad state of web user password security. Even with a relatively complex language such as Finnish, with plethora of choices in grammatically correct cases, users choose basic forms of natural language words. The result is in line with the findings of Dell'Amico *et al.* [109], discussed earlier in Section 7.3.1.1, who were able to compromise 20.24% of a dataset with a dictionary attack. The fraction of passwords made of single words out of all passwords varies between the datasets, as is shown in Table 8.

development, and will lead to a more secure authentication ecosystem in general.

Table 8: Average fraction of password types compared between leaked and cracked datasets.

	Leaked	Cracked
Single word	16.6%	12.9%
Contains word	35.8%	41.0%
Name	22.4%	19.9%
Other	25.2%	26.2%

The result also reveals that the relative hardness of the Finnish language does not offer any protection against password guessing attacks; all an attacker needs is indeed a good dictionary of Finnish words, as the overwhelming majority of passwords are dictionary words. If Finnish web users want to increase the security of their natural language passwords, they should consider using words in other, even esoteric cases.

To summarize the findings, there certainly are strong passwords among the 75% that contain natural language. But given that almost all single words are in the nominative case, it is possible, even unlikely, that this fraction is any different for passwords with embedded natural language. The complexity gained in this manner is therefore not as significant as it should be. The 25% of passwords that do not contain any natural language are not analyzed for composition, and therefore can contain both strong and very weak passwords. Based on this analysis, Finnish web users do not generally use good passwords in their services, and while natural language is used, the potential it provides for password complexity is left unrealized.

7.4 PASSWORD BEHAVIOR OF UNDERGRADUATE STUDENTS

A study on password behavior was conducted on undergraduate students, with the goal of gaining insight on password behavior of students. The study was conducted as a part of a course on information security and information society, aimed towards first and second year university students. The general background of students on the course was in IT, either in computer science or computer engineering. The study was repeated the next year during the same course.

Table 9: Word classes for passwords comprised of a single Finnish word.

Set	Nouns in total	Nouns in nominative	Verbs	Adjectives	Other
DS1	7 316	7 031	218	163	168
DS2	21 960	21 201	562	531	445
DS3	933	859	40	37	8
DS4	10 727	10 235	242	315	119
DS5	903	866	16	15	6
DS6	1 737	1 651	58	47	17
DS7	1 550	1 465	39	37	16
DS8	2 759	2 616	75	49	46
DS9	1 623	1 556	42	46	17
Total	49 508	47 480	1 292	1 240	842

7.4.1 *Research methodology*

During the first year 34 students answered the study. The repeated study got answers from 31 students, bringing the total population in the study to $N = 65$. The study was conducted as an online study, with measures in place to secure the anonymity of participants so that a set on answers and a particular student could not be connected.

The students were asked to characterize their password behavior by answering the following questions (translated from Finnish).

Q1: How many username-password pairs do you have?

- Answer options: 1-3; 4-10; 11-25; more than 25 (choose one).

Q2: Do you reuse passwords in different services?

- Answer option: open question.

Q2_A Do you reuse passwords in non-important accounts only?

- Derived question category based on open answers to Q2. Answer options: yes; no.

Q3: In how many services do you use your most commonly used password?

- Answer option: open question.

Q4: Do your passwords contain the following?

- Answer options: lowercase letters; uppercase letters; numbers; special characters; Scandinavian characters (choose any).

Q5: Do your passwords contain natural language?

- Answer options: no; yes, single words; yes, single words with letter substitutions; yes, several unrelated words; yes, several words with grammatically correct context; other (choose any).

Q6: Which of the following methods do you use to remember passwords?

- Answer options: memorization; writing down on paper; use password manager; save to browser on computer; save to browser on mobile device; other (choose any).

Q7: Do you use a password generator?

- Answer options: yes; no (choose one).

Q8: Do you use two-factor authentication?

- Answer options: yes; no (choose one).

Q9: Do your devices automatically lock themselves?

- Answer options: yes; no (choose one).

Q10: Which portion of your passwords fulfill the requirements for a strong password?

- Answer options: open question.

Basic data manipulation and plotting of results was first performed with Microsoft Excel, versions 2013 and 2016. More detailed statistical analysis was performed with R version 3.4.0. Parallel set diagrams were created using the Parallel Sets tool.[11]

11 Robert Kosara, "Parallel Sets". Online, available at https://eagereyes.org/parallel-sets Accessed 11.5.2017.

7.4.2 *Analysis results*

The results of the analysis are presented next. First we shall examine the results at a more general level, followed by more detailed statistical analysis.

7.4.2.1 *Number of passwords*

Answers that were free-form or multiple choice were quantized to discrete comparable values when necessary to compare the results. The results to Q1 are visualized in Figure 10, from which we can observe that 20% of the respondents have more than 25 accounts with passwords to remember. The majority of students have between 4 and 25 accounts with passwords.

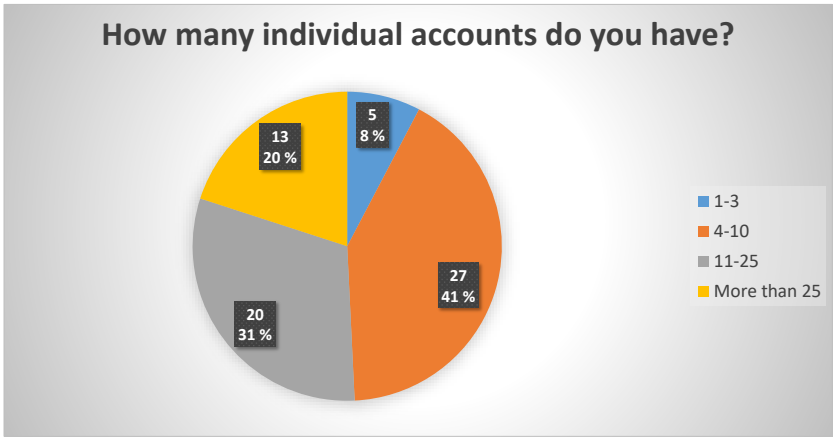

Figure 10: Number of individual passwords among the responses.

7.4.2.2 *Reuse of passwords*

For Q2, the open question was adjusted to the following scale: no; in some (1-4); in several (5-20); nearly all. An additional information category was also extracted for those who reported reuse only for non-important services, which amounted to 17 respondents (27%). The results are shown in Figure 11. One fifth did not reuse their passwords at all, and 41% only reused a password (or several passwords) in a small number of services.

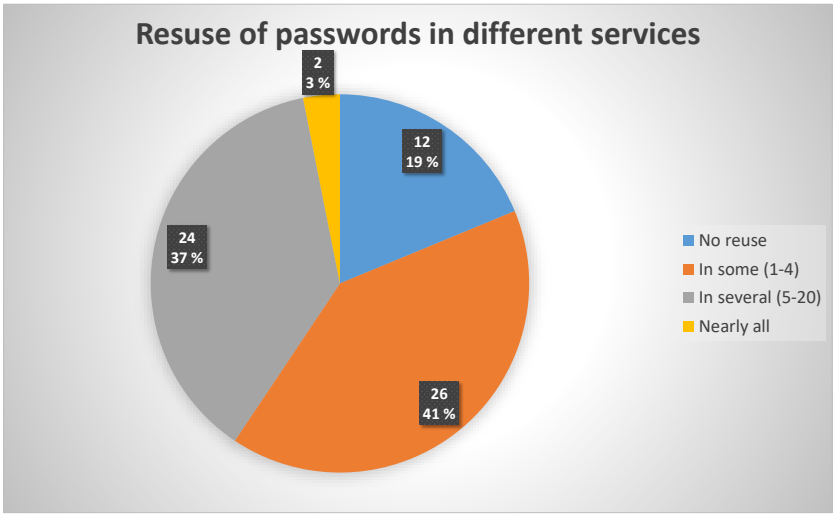

Figure 11: Password reuse frequency in among the responses.

7.4.2.3 *Number of services with same password*

A common mistake for an user to make is to reuse a single password in several different services at the same time, thus making all of them vulnerable in case of a data breach. The aim of Q3 was to measure in how many services the most common password is used. 45% of respondents used their most commonly reused password in 3-5 different services. 12% used the same password in eleven or more services, providing attackers with ample attack surface in case of a password leak. The results are shown in Figure 12.

7.4.2.4 *Password composition*

The composition of the students' passwords was asked in Q4. The results are shown in Figure 13, where the answers are gathered into groups based on which types of characters the students reported using in their passwords. The vast majority use regular alphanumeric characters in their passwords, but around one fifth of the answers had a more exotic combination of characters.

How these more special characters are divided among the respondents is shown in the parallel set diagram in Figure 14. From here we can see that student behavior in choosing what characters they use in passwords is clearly dominated by the

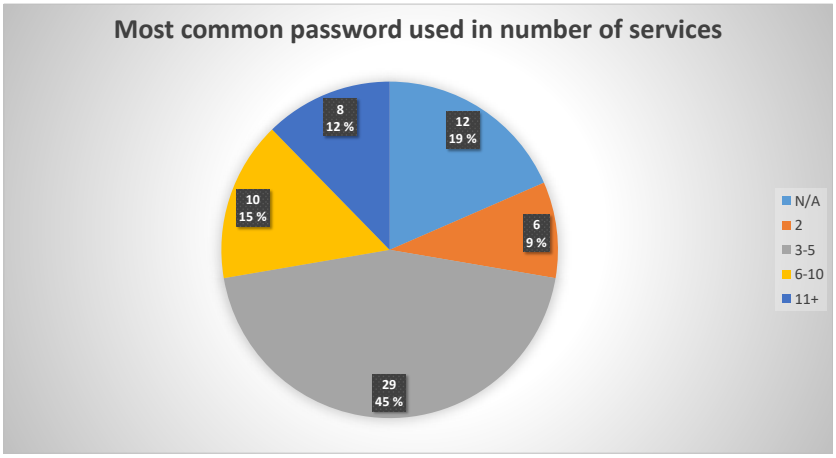

Figure 12: Number of services that the most commonly reused password is used in.

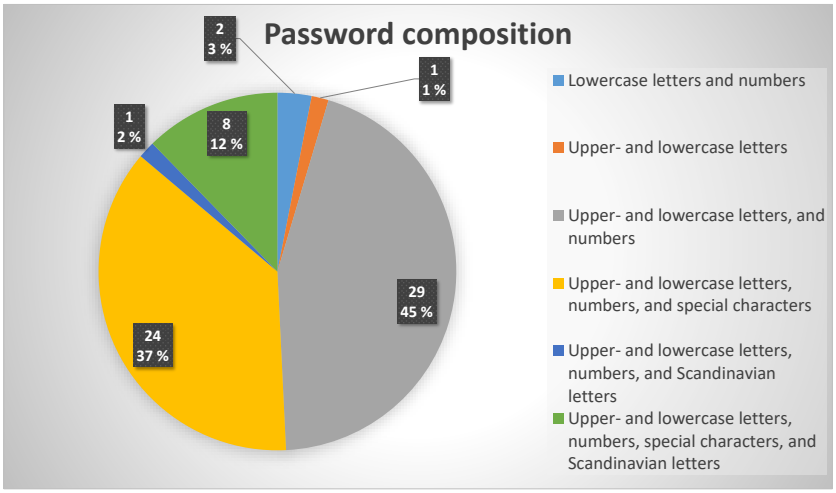

Figure 13: Composition of passwords among the responses.

more common options, and that those who have more exotic combinations, e. g.Scandinavian letters but no special characters, are clear outliers.

7.4.2.5 *Natural language in passwords*

The use of natural language in passwords is the subject of Q5. The students were asked to provide details on how they incorporate natural language to their passwords, if they do so. Out of 65

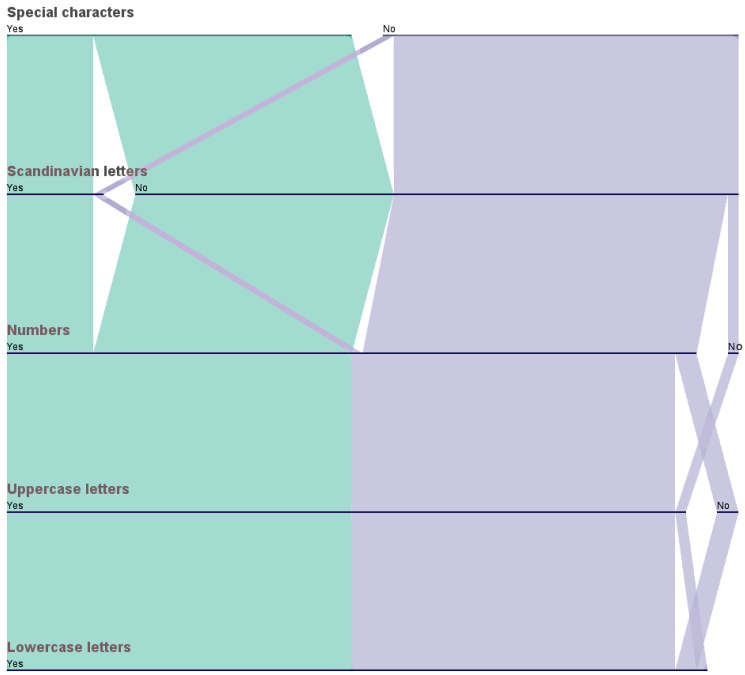

Figure 14: Parallel set diagram for answers on password composition.

responses, 26 reported that they do not use natural language in their passwords at all, while the rest had varying combinations of natural language incorporated in their passwords. These are illustrated in Figure 15. The most common methods of incorporating natural language to passwords were to either use single words with letter substitutions (i. e. i -> 1, a -> 4 etc.) to mask them, or to use several unrelated words with no context as the password. This method can also be (a bit humorously) called the "Correct Horse Battery Staple" method, popularized by the xkcd webcomic.[12]

7.4.2.6 *Password storage*

In Q6, students were asked how they store their passwords. The most common method was to memorize the passwords (60%), followed by storing passwords in the browser password safe on a computer (29%). 18% of students also reported writing pass-

12 Randall Munroe, "Password", *xkcd*. Online, available at https://xkcd.com/936/ Accessed 11.5.2017.

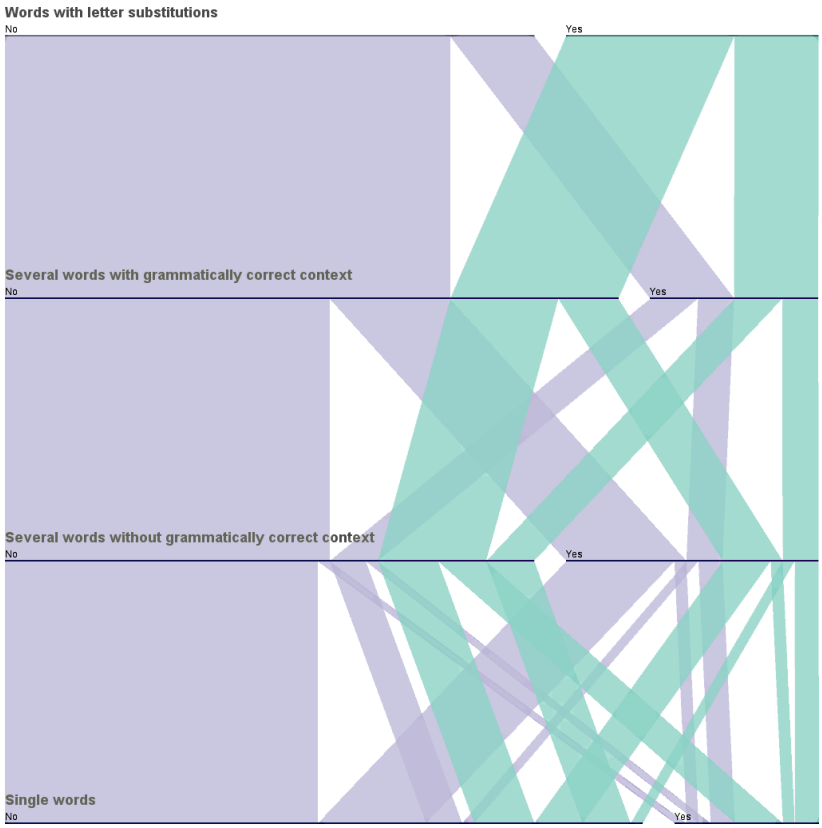

Figure 15: How natural language is incorporated into passwords among the responses.

words down on paper, and only 13% of students saved their passwords on their mobile devices.

There would seem to be a clearly identifiable group of security conscious students who do not save their passwords to any other systems except password managers or their own memory. For the most, memorization was the most common method for storing passwords.

7.4.2.7 *Password management and two-factor authentication*

A summary of the answers to questions Q7-Q9 is shown in Table 10. A majority of respondents do not (14%) use a password manager to store passwords, and using two-factor authentication is less common (38%) than using it (62%). Respondents were split (49% versus 51%) on whether their devices automatically lock

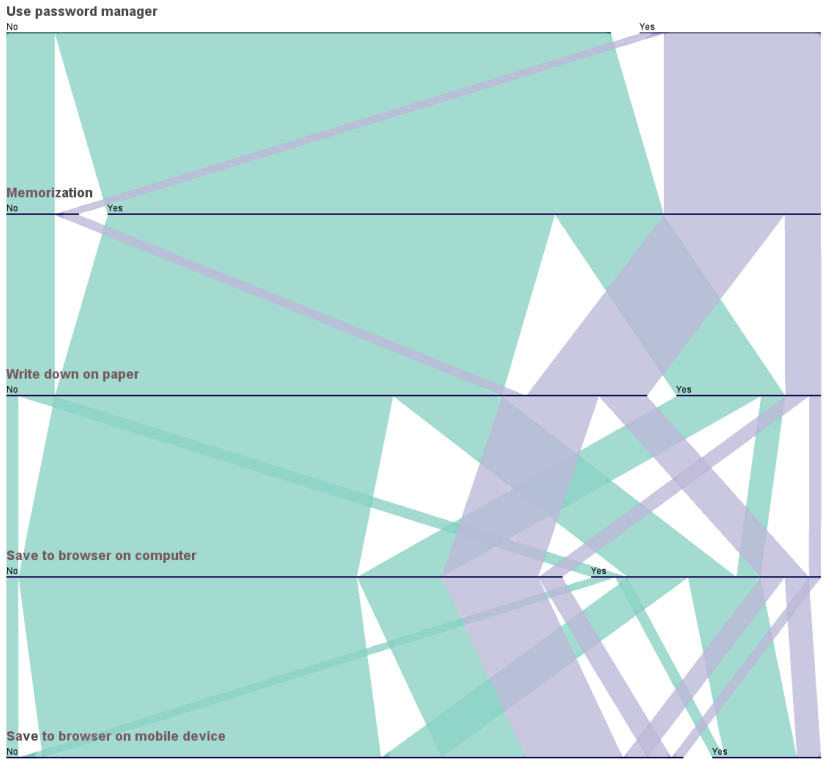

Figure 16: How students keep track of their passwords.

themselves while being unattended or not. The results are further visualized in Figure 17.

The number of two-factor authentication users at first seems to be relatively high, especially in relation to how rare applications that support it are. It may be that the question statement is flawed, as most students who answered probably have used two-factor authentication without even realizing it with one major application: online banking. All Finnish banks offer online banking services, and all banks require some kind of two-factor authentication to use online services. Taking this into account, the percentage of those who use two-factor authentication should be nearly 100%. It is possible, even likely, that most students did not understand this while answering the questionnaire, thus limiting the applicability of this result. It is also possible that due to the nature of the target demographic – university CS and CE students – it is more likely for the students to use some services

Table 10: Answers to questions Q7-Q9 that measure password security related behavior in students.

	Password generator	2-factor authentication	Auto-lock devices
Yes	9 (14%)	25 (38%)	32 (49%)
No	56 (86%)	40 (62%)	33 (51%)

that currently support two-factor authentication, such as Google 2-Step Verification[13] or Blizzard Authenticator.[14]

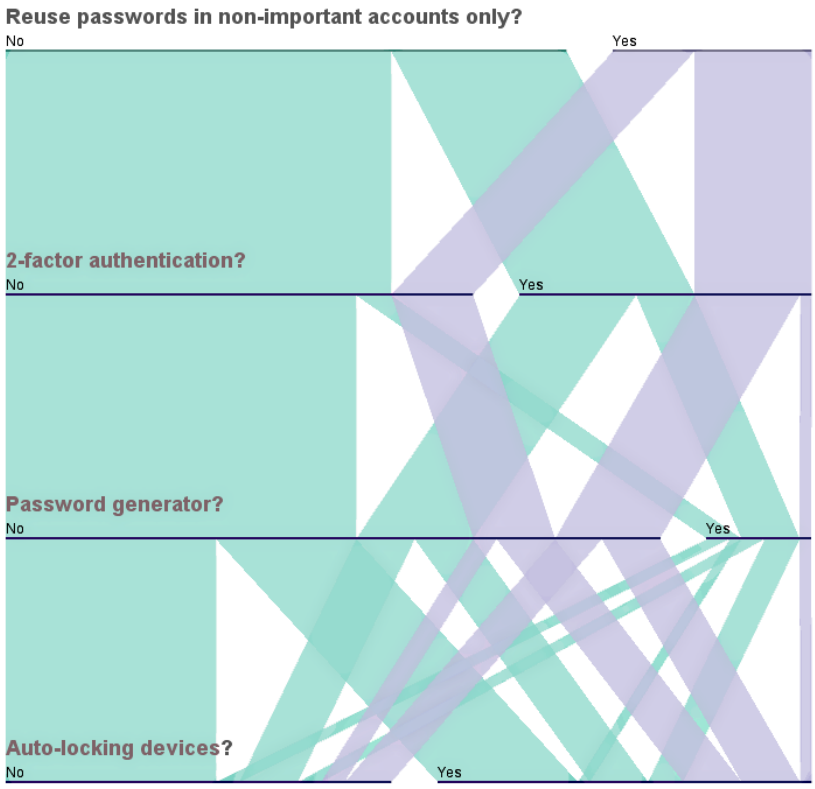

Figure 17: Visualization of password security related behavior of students.

13 https://www.google.com/landing/2step/
14 https://eu.battle.net/support/en/article/24520

7.4.2.8 *Password strength estimation*

Finally in Q10 the students were asked to assess their own password strength by estimating how many of their own passwords fulfill the requirements for a strong password. Almost 80% of students estimate that at least most of their passwords are strong passwords. The results are shown in Figure 18.

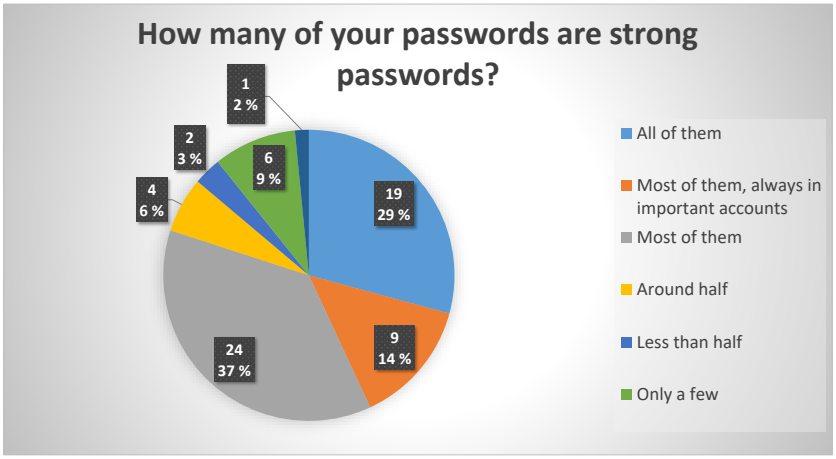

Figure 18: Student estimation of the strength of personal passwords.

7.4.3 *Statistical analysis*

The study was not designed with statistical analysis in mind and there were no *a priori* hypotheses formulated for any of the results. The goal of the questionnaire was to give a rough view of password behavior of students on various areas. Due to this lack of design consideration, the analysis on data from this questionnaire is by necessity *post hoc* and the results should be treated as such — pointing out potential avenues of further research, but not necessarily giving strong evidence. The sample size is also quite small, and thus future studies would benefit from a more extensive data gathering period.

Due to the lack of purposeful design, some questions in the questionnaire were added on as interesting side notes and some questions have redundancy with others. The goal of the study was to examine the use habits related to passwords and particularly password reuse. To accomplish this, several questions had

to be omitted from the final multiple comparison analysis. Questions relating to natural language in passwords were left out due to both being poorly designed and out of scope for password use behavior.

The following questions were eventually chosen for the multiple comparison test: Q1, Q2, Q2_A, Q3, Q4 (Only answers with special or Scandinavian characters), Q6, Q7, Q8, Q9. The decision to omit parts of Q4 were due to poor design of the original question. Nearly everyone uses alphanumeric characters in passwords, but the special cases of Scandinavian letters (å,ä,ö) and other special characters are more interesting.

The chosen question categories analyzed for statistical independence with the Pearson's chi-square test (χ^2). The null hypothesis H_0 for all comparisons is that the data sets are statistically independent.

A total of 91 tests were performed on the reduced question set. The initial p-value threshold for the χ^2 test was set as $\alpha = 0.03$ due to the nature of the source data. Out of the 91 tests, 12 were found to have a sufficiently low α, indicating that the null hypothesis should be discarded for these tests. As we are performing a *post hoc* study, these p-values are adjusted using the Holm-Bonferroni correction method for controlling the family-wise error rate in multiple comparisons. Here we use the more commonly used threshold of $\alpha = 0.05$. The results are shown in Table 11. Only the first three tests are described in more detail, as they are the most interesting ones in this analysis.

The results of this test can be characterized as inconclusive. Based on the χ^2 analysis, we can identify three cases where the null hypothesis H_0 of data independence could be rejected, but as Holm-Bonferroni correction takes into account the cumulative sum of p-values and H_0 is rejected only for those tests with cumulative p-value under 0.05, the only H_0 with the potential to be rejected is for test with ID 3.

This result implies that for the test group there is a linear relationship between the number of accounts with recycled passwords and with how many accounts the most reused password is used with. This indicates that when users reuse passwords, they tend to reuse one particular password throughout all the systems where they indeed do use a reused password. Even though the answers to Q2_A indicate that many understand the

Table 11: Results of the chi-squared test with Holm-Bonferroni adjusted p-values for $\alpha < 0.05$.

Test ID	p value	Adjusted p value	Test
3	0.0004997501	0.04547726	Q2 - Q3
11	0.0004997501	0.04547726	Q6: browser - mobile browser
12	0.0004997501	0.04547726	Q6: pwd manager - Q7
5	0.0039980010	0.35182409	-
7	0.0094952524	0.82608696	-
1	0.0104947526	0.90254873	-
4	0.0104947526	0.90254873	-
8	0.0114942529	0.96551724	-
2	0.0124937531	1.00000000	-
9	0.0134932534	1.00000000	-
10	0.0179910045	1.00000000	-
6	0.0299850075	1.00000000	-

risks of recycling passwords, in some situations users judge the risk of reusing passwords to be insignificant enough. In these cases, there would seem to exist a single password that gets chosen. In turn this leads to users having several online services accessible with a single password. These results are illustrated in Figure 19.

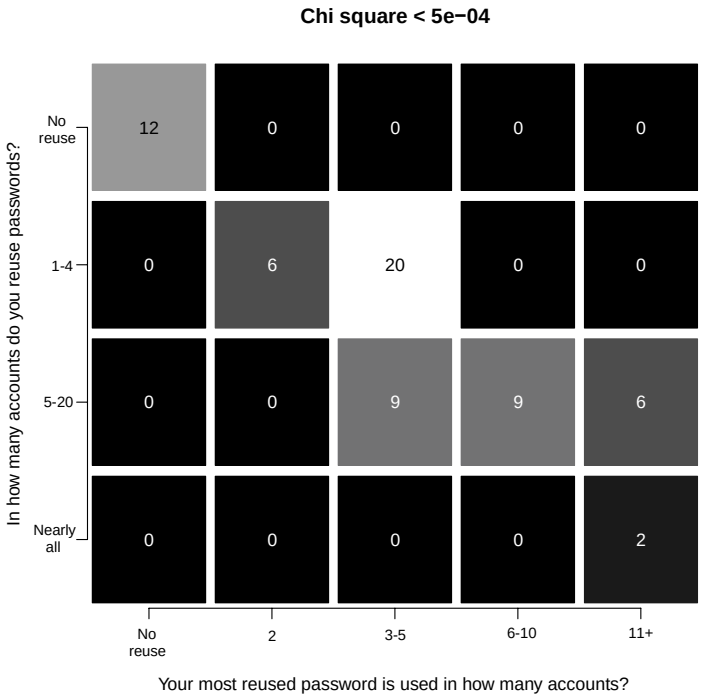

Figure 19: Heatmap depicting the relationship between number of accounts with reused passwords and how common the most reused password is.

Other potential dependences can be found in the data, but the data is not conclusive enough to reject H_0 for these relationships with this source data. Such is the case with the use of password managers and the use of password generators. Students for the most part do not use either password managers or password generators, but those who do use password managers also tend to use password generators. This can perhaps be due to password managers often having a built in password generator, and those who use a manager tend to use all its capabilities. The heatmap of the relationship between password manager and password generator use is shown in Figure 20.

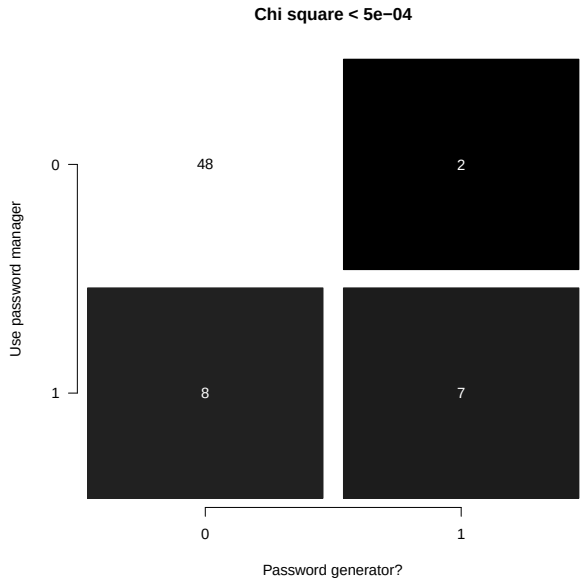

Figure 20: Heatmap depicting the use of password generators against the use pf password managers.

The second dependency that is implied in the data is between users storing their passwords in browsers on computers versus mobile devices. Storing passwords in a browser is a common method to save time and improve usability and the browsing experience. In this study, the majority did not store passwords in their browsers on either platform. Those who do save passwords on computers are not guaranteed to use the same feature on a mobile platform, though. No one used mobile device browsers exclusively for storing their passwords. The results are illustrated in the heatmap in Figure 21.

7.4.4 *Discussion on the limitations of the study*

The demographic for this study were university students, studying computer science and computer engineering. While age was not asked in the questionnaire, most freshman and sophomore students are in the 19-23 age bracket. The results are thus not representative of the whole population, and thus cannot be generalized to all users of computer systems. Further studies could at

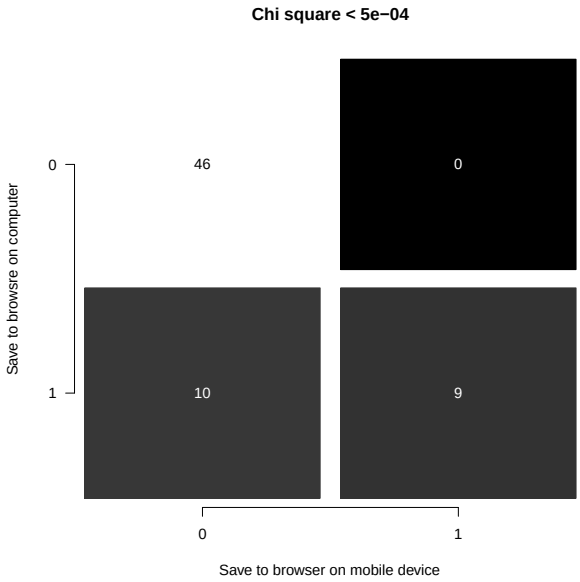

Figure 21: Heatmap depicting how students store passwords in browsers depending on platform.

first concentrate on measuring comparable statistics from other age groups, other university disciplines, and also from other education levels, and proceed from the results towards a more comprehensive understanding of password use patterns and behavior.

Future work includes a complete redesign of the questionnaire so, that the relationships this study hints at can be either confirmed or rejected. This will require more data points and a better design of experiment. Another potential way to improve on the results is to widen the scope of research beyond university IT students, and to see whether other populations behave differently. IT students are quite probably not representative of the general populace regarding information security behavior. All in all, better design of experiment and a larger set of answers are required before any hard conclusions can be made.

7.5 CONCLUSION

In this chapter we have examined methods for building trust in the networked information society. First we discussed extremely pertinent question of ownership of data and how it relates to its use in mass surveillance. By borrowing the concept of Datenherrschaft — mastery over data — from IT ethics and applying it the this problem, we have highlighted a potential solution for building a better foundation for the trust that the networked information society requires to thrive.

Next we have examined the security of password-based systems — one of the main authentication mechanisms in use in the networked information society. Even though the password has its disadvantages as an authentication tool, its popularity in practical use has not yet diminished. This also makes it one of the central trust building security mechanisms, and thus improving our understanding on how people use and choose password-based systems also serves to further increase the overall trust in society. As the administrators of information systems are often forced to trust the users not to choose insecure passwords — especially with CGISs — improved understanding on how users choose and use their passwords makes it possible to build better systems and to educate the users on how to act in a secure manner.

In our analysis we found that Finnish web users tend to use natural language in their passwords. By using NLP tools to analyze the linguistic properties of these passwords, it was found that even though the Finnish language gives unique opportunities to use complex natural language in passwords, users tend to choose dictionary words in nominative case. The results on the survey of password creation and use habits of university students were inconclusive, but give some implications on potential dependencies and direct us towards asking the right questions in order to gain more understanding in the matter.

In the next chapter we shall examine security education and how it relates to the security and privacy of citizens in the networked information society. Teaching good security practices and common issues is a central method for improving overall security, safety and trust in the information society.

8

INFORMATION SECURITY EDUCATION

NEO. *"I know Kung Fu."*
MORPHEUS. *"Show me."*

– *The Matrix* (1999)

Education is an essential part of building a more secure information society in the future. People are at the heart of a society, and how people behave with regard to security significantly shapes how an information society works and behaves. In this chapter we examine education as a tool for increasing security awareness. Educating users on security and privacy issues is especially critical for the long-term security of the Internet. Users that are ignorant can be deceived, guided and coerced into doing things no educated user would do, and also are able to ask the right questions and demand security and privacy for themselves. Users that do not understand the technicalities and underlying principles and concepts of either security nor privacy cannot even begin to behave in a secure manner, let alone demand security and privacy from the services and technologies they use.

A central conclusion drawn from the issues presented earlier in this thesis is that the general situation for security in information society can perhaps best be improved through education. While aiming education efforts to security experts is laudable and will provide better security professionals, this is insufficient in the big picture. To actually force a paradigm shift in security thinking throughout our society, we must endeavor to provide good and efficient security education to everyone — from laymen to security professionals and researchers.

Before we can even begin to accomplish this, we must first have a clear definition of what is necessary security knowledge

for the average citizen: security issues that everyone should know and be aware of. Moreover, we must have a continuum of security knowledge and understanding levels, ranging from lowest to highest, and with sufficient granularity to make each separate level justified in itself.

The following three sections are derived from existing publications. They are mainly based on the authors' contribution to [14], [12], and [13], respectively. Section 8.4 contains new and unpublished original research.

In Section 8.1 we define three different levels of security knowledge, targeted for higher education engineering students. These are further divided into three different learning profiles, separated by the engineering discipline of students. We must also be able to effectively educate our students in matters of network security. Industry cooperation and its benefits in security education is examined in Section 8.2, where we examine what aspects of network security education can benefit from having a strong industry partner. We especially focus on arranging practical lab exercises and how this can be enhanced by such cooperation. In Section 8.3 we examine efficient virtualization of network security education. Finally in Section 8.4 we assess university students' attitudes on various information security concepts using content analysis on student study journals on an information security course.

8.1 KEY INFORMATION SECURITY EXPERTISE AREAS FOR ENGINEERING STUDENTS

As we as a society move further into a true information society, it is just not the security professionals who have to be cognizant of information security issues. Students of other engineering disciplines, other professionals, and also the average layman also need to be aware of information security issues. What everyone needs to learn is very much dependent on their chosen study area and professional profile, but even now everyone needs to learn at least a limited subset of information security knowledge, and this development will continue.

In this section a framework is defined for efficient teaching of information security to higher education engineering students. The goal is to eventually extend the use of this framework fur-

ther into other disciplines and levels of education. The objective is to educate engineers with information security focus, but it can be applied to other disciplines to provide students with a strong grasp of essential information security concepts, and the ability to adapt and use this knowledge in practice. Therefore it becomes vital to correctly define what key areas of expertise in information security are for all key demographics, and having a clear vision on what to teach to and how to approach each group.

8.1.1 *Information security and the challenges of information society*

Information security is, by a broad definition, the discipline that is concerned with securing and protecting information systems against internal and external threats. Actual definitions range from "[...] protecting information and information systems from unauthorized access, use, disclosure, disruption, modification, or destruction [...]"[1] to describing information security as a parasitic order machine [107]. A universally accepted general definition seems to be elusive for information security, and provokes discussion [181].

When we consider information security as a field of study, we can immediately acknowledge its complexity and wide scope. To give a short introduction on different areas in information security, some closely related topics include i.a. theoretical computer science [182], cryptographic protocols [73], semiconductor physics [183], computer network security [184], software engineering [185], policy management [186, 187] and social engineering [188, 189]. Just the number of different aspects information security has suggests that teaching all of these subject matters to every engineering student in higher education is practically impossible, let alone impractical.

Our society has transformed into a networked information society, in which important society functions rely on computers and the Internet. To be able to function effectively and safely in this modern society one must have internalized, among other topic areas, basic concepts of information security, with the asso-

1 44 United States Code §3542. https://www.law.cornell.edu/uscode/text/44/3542 Last accessed 1.4.2015

ciated knowledge and behavior models. On a layman level some common ground rules can be easily defined, such as "do not respond to phishing emails" or "never give your password over the phone to someone". It is vital for the security of the individual to understand what personal information is and how it should be protected by businesses and the government. Other important areas of knowledge are what it means to have a secure connection to a web server, what are the realistic capabilities of possible malicious attackers with regard to different services, what attacks can target you in person, or what can be gathered from your data and metadata, and so on.

Therefore we must be able to solve a difficult equation: we need to simplify complex issues sufficiently to teach them to a wide audience, but also at the same time we must provide cutting edge education for those who are interested in learning the state of the art in information security. Combine this with limited resources, and we must begin to choose our education profiles very carefully, and target our efforts to the correct audiences. We can easily argue that everyone should have at least a basic grasp of information security concepts, but here we will mainly focus on higher education engineering students in IT related fields, as they are on the forefront of information society as consumers, creators, maintainers and protectors.

While it is out of the scope of this thesis to explore this any further, we must additionally remark that the question of information security core competence is also a relevant question for all other fields of education, and by extension, all of society. Many basic functions in our society have been transformed: online banking, online shopping, social networks, and even voting in elections can be done online in some instances. These are all basic functions that everyone needs access to, and everyone needs to be able to use, and also perceive the possible risks and threats that these new services have. This means that everyone needs to have a certain understanding of the underlying system, and this brings us to information security and information security education, and the importance of recognizing that what is now considered to be in the realm of IT professionals will eventually creep into other aspects of life, and eventually be a part of everyone's basic knowledge [190].

8.1.2 *Learning profiles and different demographics for security knowledge*

Definition of different student demographics is a key step in the process of improving information security education by targeting the correct information to the correct demographic group. This targets education precisely to the students that need it. Here we focus on engineering students, with a particular emphasis on IT and IT security engineering students. Existing literature is far from homogeneous on this matter. Studies focusing on information security awareness, education and training often target company employees, separating them based on titles, responsibility areas or employees and contractors [191], or diversify the responders based on age and level of education [192], or focus on university staff and policy makers instead of students, and with the goal of mapping out vulnerabilities in the security environment rather than assessing student learning development [193]. Previous studies are often made with a more general student population and targeting less specific security knowledge [194], or are directed towards student ethics rather than in-depth student security knowledge [195, 196].

As we have no obvious usable prior demographic division to use as a model, and we are targeting a specific subgroup of students, we have to define our own demographic divisions. We begin by identifying three groups within engineering students, with associated learning profiles.

- *Non-IT engineering student:* Engineering student in a non-IT related field, such as mechanical engineering or civil engineering. These fields may have a wide range of requirements for IT literacy and security knowledge. For most parts we can estimate that the basic minimum requirements are that they can handle themselves in a modern office environment without jeopardizing the security of their employer, and similarly they are capable of understanding and mitigating information security risks in their personal lives. Some disciplines may require more advanced IT literacy and skills than non-engineering students, with associated requirements for security.

- *IT engineering student:* Engineering student in an IT field, such as software engineering, computer engineering, embedded systems engineering or communication systems engineering. Students with this learning profile require in-depth knowledge on one or two IT topics. Additionally a basic understanding on most other areas, from algorithms to databases to semiconductor operating principles, is required. Information security knowledge required can vary between disciplines, but for the most part the students should be able to understand the topics from the previous profile. Deeper understanding of security can be required for some specializations, however.
- *IT engineering student with security focus:* An IT security professional whose main task is to analyze, develop, implement and evaluate different entities in information security. This profile emphasizes specific information security topics in addition to the main aspects from the IT engineering student learning profile.

The learning profiles are cumulative. IT engineering students should incorporate non-IT engineering student learning profile contents to their own, with some limitations. Similarly, a security focused engineering student should have for the most part the same learning profile as an IT engineering student, but with additional requirements and focus areas in security knowledge, and perhaps some relaxations in other areas. Next we will present a thematic division of the field of information security, and map the previously presented profiles to these thematic areas.

8.1.3 *Thematic areas in information security*

Information security as a field spans many disciplines. Before we proceed to define content for the learning profiles, we must first divide the field into manageable sets of compatible concepts and technologies. Previously presented taxonomies on information security are for the most part more focused on modeling things on a more granular level. Fenz & Ecklehart [197], Silic & Back [198], and Blanco *et al.* [199] all approach the problem of classification in their own manner.

When we consider the 12 main themes and 43 sub-themes defined by Silic & Back [198] — their focus is on quantifying directions of research in information security — we consider this approach to be too granular for defining manageable knowledge areas for security education. Instead, we use a division of thematic areas first presented by Blanco *et al.* [199] as a basis for our classification. Blanco *et al.* define three loose knowledge groups: vulnerabilities, threats, attacks; security protocols, mechanisms, policies and controls, countermeasures. We had already identified a similar set of thematic areas on our own, giving support to this division of concepts as very intuitive. The thematic area of data security and information criticality is further added to fully flesh out the thematic areas, which are presented in detail below.

VULNERABILITIES, THREATS, ATTACKS: This thematic area contains all the various vulnerabilities encountered in , whether they are vulnerable software, hardware and wetware (i. e. people), and the threats and methods that leverage and use these vulnerabilities. The concepts of malicious software, different malware types, and their capabilities are also included here. Knowledge about attackers and their goals and motivations, different attack types (active or passive) and softer methods such as social engineering has also an important part in understanding vulnerabilities and threats, so they are included in this thematic area.

SECURITY PROCESSES, MECHANISMS, POLICIES: In this thematic area we have placed more abstract concepts and systems related to information and security system complexity and security policy definition. The idea of the Confidentiality-Integrity-Availability triad and its extensions are also included. Other topics include risk analysis and risk management, threat recognition and mitigation, cyber security and cyber warfare, understanding of critical infrastructure, and legal frameworks.

CONTROLS, COUNTERMEASURES: This thematic area contains the most traditional information security concepts, as they are the most salient for advanced users, and also some of the most complex. They include firewalls, intrusion prevention and detection, cryptography, password security, RFID and NFC technologies, and in general, technological solutions that can be used in implementing information security.

DATA SECURITY AND INFORMATION CRITICALITY: Data is a significant driver, whether it is business, govenment, or private. Understanding different kinds of data (location data, medical data, financial data, personally identifiable or sensitive data, etc.) and its value and applications is critical for many professions and fields of work. Management of data is brought into personal frame of reference with our personal data. Data storage and especially cloud storage are important concepts that by themselves justified a separate thematic area. Network surveillance and privacy, metadata and metadata analysis are also grouped in this thematic area.

8.1.3.1 *Learning profile mapping with thematic area*

Next we map the previously presented learning profiles into our thematic areas, and define – in broad terms – what are the requirements for each profile in each thematic area. For easier referencing, from now on we refer to the learning profiles as Levels 1, 2 and 3, as the like we already noted earlier, the profiles are inclusive when we move forward. Generally on level 1 in all thematic areas, the requirement is to understand the basics, level 2 requires deeper understanding of the background and the capability to apply learned knowledge to practice, and level 3 requires, in addition to the previous levels, the capability to design, analyze and/or implement and manage entities in the thematic areas. Competence on level 1 can generally be assessed with an online exam designed accordingly, but it gets harder as we move further in the framework, as the focus moves from knowledge and understanding to behavior and applying knowledge to practice.

When we examine the thematic areas more closely, we postulate that the requirements on level 1 will be the "new normal" for information security knowledge for the average person in the future. The development of the networked information society has already brought cyber security to the public discussion. This creates a strong motivation for information security education, and places pressure on everyone to adapt secure practices in their everyday lives. We argue that what has previously been required knowledge of information security for experts will be beneficial, if not necessary, for the average person to understand to some

Table 12: Mapping of knowledge areas to different engineering education profiles [14]

	Level 1	Level 2	Level 3
VULNERABILITIES, THREATS, ATTACKS	Aware of existence and capable of identifying when encountered.	Understands fundamental operating principles, and how they work.	Capable of analysis and/or reproduction, has deeper understanding of operating principles.
SECURITY PROCESSES, MECHANISMS, POLICIES	Understands scope and purpose. Knows how to evaluate situations and comply with policy. Understands risk management thinking.	Capable of applying processes and policies to own work.	Capable of evaluating and designing policies and processes. Can monitor and supervise enforcement of policies. Understands effects of human behavior.
CONTROLS, COUNTERMEASURES	Understands basic concepts, functions, limitations and threats. Knows what assets they protect.	Understands underlying principles, capable of applying existing controls and countermeasures in practice.	Capable of managing, analyzing, or designing new countermeasures and security controls.
DATA SECURITY, INFORMATION CRITICALITY	Understands concept of critical data, importance of securing it, and where it is located. Separation and securing of personal and work data.	Able to protect critical data. Capable of applying principles of data protection to own work, deeper understanding of technical ramifications and information value.	Capable of evaluating soundness of data protection and managing mission critical data. Ability to design new data protection methods.

Similarly to function creep

extent. This requirement creep will make something now in the realm of IT professionals to everyday knowledge in the not-so-distant future.

8.1.3.2 *Critical knowledge areas in information security*

Now we will approach the question of what topics in information security can be considered to be of critical importance. Based on the mapping presented in Table 12, we define critical knowledge for an engineering student to be at the very least what is covered in the non-IT engineer column, or level 1, in all categories. What we have considered to be critical knowledge within the thematic areas can be summarized as follows.

CONTROLS, COUNTERMEASURES: network security concepts such as firewalls, antivirus systems, and methods for securing communications, specifically cryptography and its basic principles.

VULNERABILITIES, THREATS, ATTACKS: fundamental knowledge on malware, how they function and what they target, vulnerabilities in software and systems and what kind of attacks can be targeted against them.

SECURITY PROTOCOLS, MECHANISMS, POLICIES: security policy fundamentals, significance of security awareness to overall security, security as a mindset.

DATA SECURITY, INFORMATION CRITICALITY: importance of data in information society, fundamentals of data protection.

8.1.4 *Assessing information security knowledge in students*

So far we have defined information security thematic areas and learning profiles for students. To test our framework, we have used it in first year education of IT engineering students, attempting to target our target demographics with the right level of information. Next, we examine the results of a course designed to teach information security concepts and knowledge to a class composed of students with varying backgrounds. For now, the course is directed at first year IT engineering and computer science students, and to a small percentage of students from other disciplines doing a minor subject in IT. The main method for assessing student performance and learning in the

thematic areas is an online exam that is administered at the same time to all students using an online learning environment. The exam is not the only graded part of the course, contributing about half of the grade on the course, a large group assignment being the other contributing factor.

8.1.4.1 *Online education environment*

The ViLLE platform [200] is an online learning and collaboration platform[2] initially developed for programming education, but has been extended to a full-fledged online learning platform, capable of supporting automatic checking of programming questions, several different types of assignments, and, most importantly for us, a flexible way to do online exams for a course. We decided to use ViLLE for the course because it provided us a flexible platform that supports automatic checking of exams, even for more complex exam types than multiple choice exams. Also, the integrated statistics and data extraction tools gives us good tools for gathering and mining data on student performance. An example screenshot of the ViLLE platform in use is shown in Figure 22.

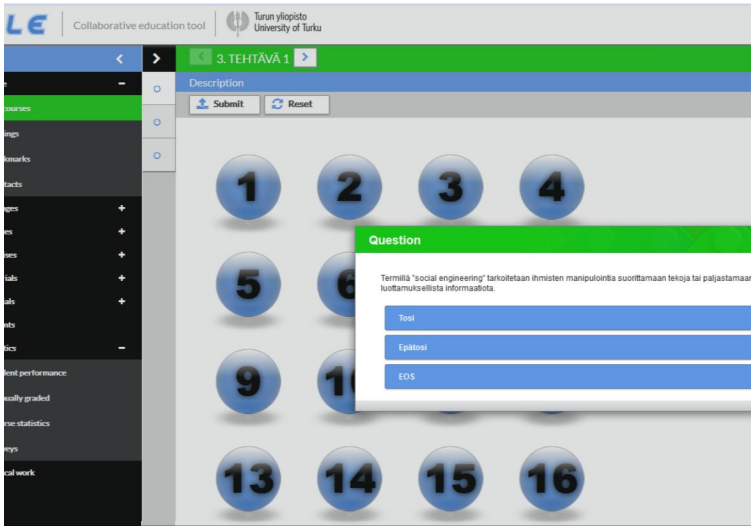

Figure 22: ViLLE online learning and education platform

2 ViLLE system home page. https://ville.cs.utu.fi/ Last accessed 20.6.2017

8.1.4.2 *Measuring learning outcomes*

The exam we use to assess student learning performance and knowledge on subject is implemented on the ViLLE platform as an online test, available to all students on the course at the same time. The exam is location independent, so students can do the exam wherever they want, as long as they have an Internet connection available. The exam is naturally open book, as students will have access to all lecture materials and the Internet, so any advantage students may have from doing the exam in groups is mitigated. The exam is structured according to our framework, so we have made a set of questions for each thematic area, divided into smaller randomized sets that can be assigned individually to students, making each instance of the exam random. The questions themselves are designed to measure knowledge and understanding on level 1 of our framework, because the scope of the course is in first year students. The correct answer gives one point, incorrect answer gives one negative point, and selecting *don't know* gives zero points, with the purpose of conditioning students not to resort to guessing in matters of information security.

Our goal is to expand the use of this framework into education in advanced level courses, where we would have to engineer methods for assessing knowledge, understanding and behavior of students in more demanding levels of competence. This would give us better control on the students' overall learning process from introductory to advanced level courses.

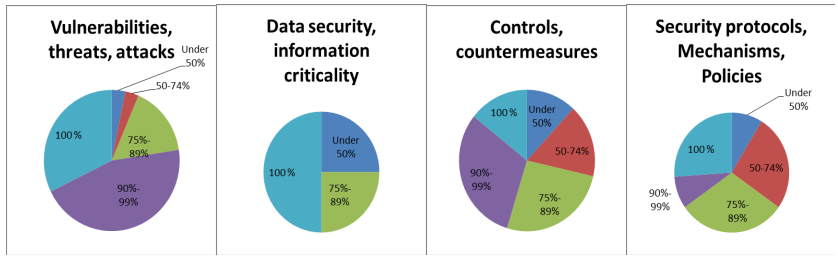

Figure 23: Exam results for students divided by thematic areas.

Data from the exam questions shows us that the most difficult thematic area was controls and countermeasures, and the area that had consistently better results was vulnerabilities, threats and attacks, as seen in Table 13. The other areas fell in between,

Table 13: Correct answers divided by thematic areas

Thematic area	Correct answers	Incorrect answers
Vulnerabilities, threats, attacks	93,4%	6,6%
Data security, information criticality	83,9%	16,1%
Control, countermeasures	79,7%	20,3%
Security protocols, mechanisms, policies	82,2%	17,8%

but with similar percentages in the 80%-90% range. Our goal for the students was to have them score 90% or more overall, and we achieved this only in one category. The goal was set so high because the students had all materials available during the exam, and also had access to the Internet, which corresponds to the normal situation the students have when facing a problem: they Google it. The questions were designed to be not that hard and we focused on basic knowledge in thematic areas, justifying setting the target to be as high as 90%. From Figure 23 we can see, that out of all questions in vulnerabilities, threats and attacks, more than 75% were answered correctly >90% of the time, signifying that nearly all students had a good understanding of the area, contrasted by only about 33% in security protocols, mechanisms and policies. The amount of data in each category is yet too small to yield robust results, but these preliminary results serve as a motivator for refining our framework. Especially the data security and information criticality area has very limited data, and those results should be considered as tentative at best, but this can be attributed to its late emergence after the first analysis on preliminary results from the exam. One key future work aspect is strengthening this thematic area.

The results also reflect the fact that vulnerabilities and data security are categories that students have had more contact with through the group assignment that forced the students to ponder on these issues. Also, based on the data in Figure 23, we assume that students yet lack the technical and mathematical framework to process concepts like cryptography and security policies efficiently. We could perhaps facilitate this by directing the group work in these directions, making the students focus on

these topics even in simplified cases. While we do not yet have sufficient data on this, we hypothesize that data security and information criticality will emerge as a theme which is taken more seriously by older students. This is both because they are further in their studies and understand the role of data better, and because younger people are more comfortable with sharing their data in general.

8.1.4.3 *Discussion and future work*

Our framework for teaching information security concepts has provided us with a clear method for creating sets of exam questions that are able to measure student performance within our defined thematic areas. It can be used for testing student knowledge, and also in the course design phase by making sure that all relevant aspects of information security for each target learning profile are taken into account when designing the course material.

One salient feature we noticed was that after critically examining our chosen exam questions, we found that all questions except one (roughly 99%) were focused on measuring student's knowledge on the topic, instead on their perception on how to behave in a secure context in challenging situations. Some arguably behavioral aspects were present, but in this version of the framework they were approached from a knowledge viewpoint. While behavior is hard to assess in exams of this format, one key future development is developing a test which would also incorporate behavioral issues (i. e. estimating not just knowledge but how students behave in a real-life situation), while being implemented on an online learning platform.

The long-term goal of this framework is to expand it to cover behavioral aspects, and expand the learning profiles outside engineering education, aiming for a generalizable model of security education that can be adapted to different levels of education, with readily made questions in a suitable reference frame and a flexible platform for practical implementation. The first level of security knowledge does not include several important aspects of information security knowledge, and a major future endeavor is to expand the question sets to cover the thematic areas in more detail. Possible expansions include more in-depth

technical topics, legal requirements, standards and security frameworks, which all are in the domain of the IT security engineer, in contrast to the general knowledge aimed for non-IT engineers. More comprehensive analysis on the applicability of this framework is due after more data from further course iterations has been processed. Another potential research direction is to examine how the framework is positioned in relation to Bloom's taxonomy (see [201, 202], for example) and to research potential adaptation of concepts from Bloom's taxonomy.

8.2 INDUSTRY COLLABORATION IN NETWORK SECURITY EDUCATION

In this section we examine the benefits of having an established industry partner in arranging practical exercises and hands-on training for students in network security.

Gaining experience in hands-on laboratory work is essential for engineering students to facilitate their development as future professionals in their specialization. In network security, one cannot become an expert professional in administering firewalls and Intrusion Prevention System (IPS) just by reading textbooks: proper and adequate laboratory experiments are needed. For a brief introduction to firewall and IPS concepts, see e. g. [184]. Unfortunately, building a research and teaching laboratory environment with powerful computing equipment and specialized hardware and software for the target lab works is often costly and a public institution like a university may be reluctant to invest money in an expensive new laboratory. A beneficial solution to the problem is to find an industrial partner from the research area and start negotiations for university-industry collaboration in building a laboratory. All parties of the collaboration benefit from the co-operation: students have the possibility to perform hands-on laboratory work, the university is better able to include work-life relevant education in its curriculum, and the industrial partner gets visibility among students and is able to contribute to university education planning from the educational needs of professional careers point of view.

Our industry collaboration has resulted in building a network security lab for research and education, where modern powerful computing equipment is used together with specialized fire-

wall and IPS hardware and software from a recognized manufacturer to provide students with hands-on laboratory experience and skills on using and administering state-of-the art network security solutions. The hands-on work is organized into a laboratory course where theory learned in lectures is put to test in lab work. At the end of the course, most successful course participants have an opportunity to attempt the vendor's certification as system administrator, firewall architect and IPS architect. The collaboration had been going on – at the time of writing of the original paper – for three years, and the experiences are positive both from the point of view of the university and the industrial partner. Student feedback is also positive, leading us to the conclusion that tight co-operation with an industry partner in organizing hands-on network security laboratories to engineering students is fruitful for all parties and reaches the planned student learning outcomes very well.

8.2.1 *Necessity of network security education for IT engineering students*

Communication networks have spread to all facets of our society. The use of Internet on many things has become ubiquitous and engineering students need to be aware of the challenges posed by this development. These challenges range from basic network security issues to large-scale software security and safety problems, which have demonstrably affected the world in the form of, for example, the high-profile attacks against corporations[3] and government contractors. On a smaller scale, the responsibility of the end user to protect their own data and privacy by understanding the hazards and realities of security has increased. For example, the widespread password leaks in Finland in late 2011, also discussed in Chapter 9, are a prime example of this, as the attackers used the low expertise of message board and website admins to their advantage. Therefore it is imperative that students are familiar with the basic concepts of network security and software security, even though they will not work in the security industry after graduation. The multi-domain skill sets

3 Kim Zetter "Researchers Uncover RSA Phishing Attack, Hiding in Plain Sight". *Wired magazine* 2011. Online, available at http://www.wired.com/threatlevel/2011/08/how-rsa-got-hacked/ Last accessed 22.9.2016.

acquired from this kind of education are important for students working in the IT field [203].

At the University of Turku, the societal need for engineering education in information security was recognized already a long time ago. The Technology Industries of Finland Centennial Foundation decided to start funding an information security minor subject in our department for IT engineering students starting in 2008. In 2010, the university recognized the value and importance of information security engineering education and research, and placed this discipline among its four strategically important research areas in a strong development stage. This development led to the inauguration of our information security engineering Master's programme that had its first student intake in 2011.

Our department's previous curriculum on information security contained only precious little hands-on exercises and focused on traditional lecture teaching. However, previous research has determined that when students are required to find significant amounts of information themselves in their mandatory curricular work, rather than have the instructor provide the information directly, learning results can be expected to improve [204]. Application of self-regulated learning theory to motivate students towards more independent work has also produced good results in the form of more successful learners [205]. In our department, we organized our information security study modules in new ways that require considerably more independent work throughout the course duration than our average lecture courses. The techniques we applied to achieve this were strict requirements for weekly group work, weekly independent reporting and term papers with strict deadlines always enforced by an online learning platform [206, 207]. Even with this increase in the requirements for independent working, we still found our curriculum to not include enough practical hands-on training in key issues. Also the CDIO standards[4] that nowadays guide the curriculum planning in our department emphasize the amount of independent and hands-on work in engineering education. The idea for building an information security laboratory to support our curriculum arose from these needs. The goal of this laboratory environment was to provide the opportunity for a wide

4 CDIO Initiative. Online, available at http://cdio.org/ Last accessed 29.9.2017.

range of practical tasks and exercises for information security students. Information security is a field that can be taught only to a certain extent in the classroom. Practical experience in using all relevant security solutions, software and hardware is an important aspect of building the expertise required for security practitioners. While theoretical studies of the principles of information security such as cryptography, network technologies and software engineering and computer science are naturally important, hands-on experience cannot be ignored. By having up to date laboratory facilities for information security teaching and research, it is possible to provide this opportunity to the students.

Industrial collaboration also provides more opportunities for the students to expand their skill sets. Information security software and especially hardware is expensive, and enterprise solutions are usually out of the consumer market – and thus out of reach for the vast majority of students – due to prohibitive costs and because they are designed for networks and traffic of a whole different scale. Without such an opportunity for the students to familiarize themselves with at least one enterprise solution before entering the job market, they will not have the crucial hands-on experience that potential employers seek in new employees.

8.2.2 *Lab specification*

To make it possible to provide the latest technologies to the teaching environment, we first started with the specifications of the lab. The natural first requirement for such a lab is that it should be expandable and easily scalable to different situations. The information security education environment for the lab is based on the StoneGate firewall and intrusion prevention system[5], so the lab must be able to run the environment without any problems. The first instance of the training environment had physical stand-alone devices as firewalls and IPS devices, and the updated environment is wholly virtualized. The stand-alone devices required less computational capacity from the lab infrastructure,

5 Presently sold under the name Forcepoint. https://www.forcepoint.com/product/network-security/forcepoint-stonesoft-next-generation-firewall

but were more rigid as an infrastructure than the virtualized environment. When industrial collaboration is an important part of an information security lab, the specifications and development plans for the lab are closely tied to the developments at the industrial partner and their products, as our example shows.

The network infrastructure for the lab is another important design aspect. Because of the nature of the work done in the lab, for example handling malicious code and malware, or performing network attacks or disruptive traffic, the lab network must be totally isolated from the rest of the network. On the other hand, a connection to the intranet and the Internet is also required for several basic tasks, so the network must be configured so that it is possible to reconnect it to the Internet easily. Practically this is achieved by either physically removing the connection from the switches, or if all traffic should be routed through one host, using several network interface cards, with one card dedicated to the external connection.

Back in the day, at least. The time from building the original lab to this dissertation is 8 years.

The lab has 12 high-end desktop PCs for students, and one PC for the lab instructor. The backbone of the lab network is a pair of 24-port Gigabit Ethernet switches, with additional smaller switches for isolating smaller network segments for the lab environment. Wireless LAN USB dongles are available for all PCs, with 802.11n-compliant base stations, providing the environment for wireless security projects and research. After three years, it became clear that the hardware requirements of the lab can and will change over time, and this should be accounted for in any plans for hardware acquisitions. This was the case with our switch to a virtualized environment.

The instructor computer was modified to run the VMWare vSphere ESXi 4 hypervisor, which provides the possibility to flexibly add and remove different computers to and from the lab. An another benefit of virtualization was that creating multiple instances of a teaching environment is possible, and it is trivial to roll back changes or change the environment on the fly for different teaching groups, for example. By creating snapshots of virtual machine states, it is possible to have a fine-grained picture of the lab environment, which can be changed in the order of minutes instead of hours. This is a benefit for the lab administrator, as doing infrastructure and operating system changes manually on dedicated hardware is time-consuming and requires more ef-

fort. The streamlined experience achieved with virtualization is a definitive advantage.

The administration of the lab of this size and scope is not a trivial task. Because of the wide use of the lab in different courses and projects, the schedules of courses and projects must be taken into account when planning the teaching curriculum. A good approach to administration and scheduling is to have a well-documented plan on who has access to the lab, what equipment should be available at any given time, and what courses or projects have precedence over others. A single administrator can find the task of managing all aspects of the lab time-consuming, but this is partially remedied by the use of virtualization in the teaching environment.

8.2.3 *Lessons learned from the network security lab*

Our lab has been operational in teaching and research for three years, and some concrete benefits can be readily derived from this experience. All parties – the students, the university and the industrial partner – can be seen to benefit from this co-operation.

8.2.3.1 *Benefits for the students*

The main benefit for the university is the use of enterprise hardware and software in teaching situations. This gives more opportunities for forming a well-rounded teaching curriculum, and given the addition of training materials from the partner, saves time and effort. It must be noted that these materials must be analyzed and supplemented with additional information and theory to broaden their scope and make them viable for use in the classroom. The training materials are specific to the product in question and usually assume that the students have the required background knowledge in the field. We have found that this kind of a hybrid model of industrial and academic training materials works well in teaching situations. The advantage of using industrial state-of-the-art products in classes also gives the lecturers a feel of what is the latest in the field of commercial products. The gap between academic and industrial state-of-the-art can vary significantly. This also guarantees that teaching is actually relevant to what is happening in the industry at any given time.

The students are able to get hands-on experience with state-of-the-art products, something which is not usually possible without industrial co-operation. This is an important factor in their professional profiles, even more so if they have passed the vendor certification tests. The students may be also able to do their master's thesis for the industrial partner, thus gaining important experience from industrial R& D. Usually one supervisor is provided by the company, which means more efficient allocation of university staff supervisory resources.

8.2.3.2 *Benefits for the industry partner*

When software and hardware which are usually available for enterprise users is made available to students and the university, it does not only benefit the students. By enabling students to get familiar with vendor hardware and software during studies, the vendor also increases the visibility of their own products among students and new graduates. The skills learned during studies will obviously affect the way those students work and operate after graduation, and it is not inconceivable that students who have certified themselves will pick that particular vendor's product, should they be in charge of procurement.

An important aspect of the collaboration is the possibility of the industry partner to provide research topics for students. These topics can range from smaller individual projects to master's theses. This frees research resources from the industrial partner to other research tasks, and simultaneously benefits the student as well. Also, the partner receives valuable feedback on the functionality of their products, and is able to benefit from this feedback in their own quality assurance processes.

8.2.4 *Conclusions*

The benefits from a dedicated lab to teaching and research in the field of information security are concrete, and these benefits increase substantially with university-industry co-operation. From the first three years of operation of our lab, it can be assessed that the experience has been positive for students, university staff and industrial partner. From our experience, careful planning in hardware selection, lab schedules and use and is required for

successfully implementing such a lab, with a lot of focus on the requirements from the industrial partner to guarantee that the lab stays operational and up to the specifications. The practical experience gained by students when working with state-of-the-art technologies and industry certificates provide distinct added value to their studies, and researchers have access to an environment where they are able to implement their research in a secure, controllable environment.

The use of cross-discipline education in information security gives students important skill sets that are required when assessing and analyzing security and privacy issues in modern systems. The use of biometrics in travel documents and the problems such an infrastructure has, as discussed in Chapter 5, is a good example of a multidisciplinary problem which requires information security expertise, even though it is not the focal problem. By having the necessary infrastructure in place and having good ties to the industry by means of co-operation, universities can prepare students to face these new challenges and provide them the tools they need to succeed.

8.3 VIRTUALIZATION OF NETWORK SECURITY EDUCATION

In this section we examine how network security can be effectively taught to higher education students. We focus on lessons learned from transferring physical lab infrastructure to a virtual environment that provides better accessibility, modularity and education quality. Hands-on experience in network security is a key aspect of security-oriented communication network engineering studies [208]. Students must be able to study network environments and to achieve this learning goal, a laboratory with sufficient equipment for different network engineering tasks and exercises is necessary. Such laboratories take up a lot of physical space and require a significant amount of equipment, such as computers, network switches, wireless routers and physical network cables, to be able to facilitate well-rounded exercises. The developments in virtualization have made it possible to manage complex systems of virtualized computers, and the application of virtualization to computer network engineering education has already been explored [209], with encouraging results.

Lab sessions requiring no oversight can be time and location independent, only requiring administrator intervention in error situations, facilitating more open and flexible learning experiences for students. Different network environments are easy to simulate, instead of having to physically rewire Ethernet network cables, switches and routers. Reconfiguration and installation of classroom computers is also straightforward with virtual machines, compared to physical computers. Experiences with using virtual environments for education have been positive at our department, for both students and teachers. We can realize significant savings in expenses combined with more dynamic teaching environments, making this an appealing approach to arranging laboratory exercises for students, staff and faculty.

In this section our experiences with virtual environments in university level network security lab education are presented. We discuss the requirements for a network security lab, and outline our design for an environment which enables instructors to teach laboratory sessions without a dedicated computer classroom by leveraging the students' own devices, such as laptops, tablets and even smartphones. We outline the requirements for such an environment, and discuss its positive and negative aspects with regard to student learning experience, education environment engineering and technology, and also consider financial implications.

8.3.1 *Shifting from traditional to virtual environments*

Our experience with virtualization of laboratory education began with the transformation of our industry partner-provided network security learning environment from host-based to virtual environment, where a significant part of resources run on virtual machines [12]. This work is presented in Section 8.2. In this environment, a classroom of interconnected host computers and physical network infrastructure is still required, binding resources in computers and premises. If we consider an average computer classroom, most of the space is taken by physical equipment, leaving very little space for other purposes. This has been one of the main motivators behind our redesign of the lab, but other aspects of lab teaching will also significantly benefit from redesign. Requirements for a network security laboratory

can vary significantly depending upon topic areas and education goals. Next we will examine the requirements for a modern lab environment, describe our current setup and how it matches to the previously presented specification, and finally outline the future lab environment and the solutions behind it.

8.3.1.1 *Lab requirements*

NETWORK INFRASTRUCTURE: A lab should provide means for students to configure, design and implement different network infrastructures, and to actually test their work and receive feedback on their performance. Computer networking can be a difficult topic for students, who may not be used to conceptualizing complex network structures. Actual design and administration of a nontrivial computer network is not an easy task, and while this is not the goal of network security education, those responsible for planning and implementing network security features and also investigating possible failures should have a good grasp on networking, above the usual level of connecting few devices to a home router and calling it a day. Physical network hardware makes this process more concrete, but virtual networks can be so much more versatile that using virtual networking gives significant benefits for education.

SERVER INFRASTRUCTURE: Server hardware is required for running the virtual environment. Before OS virtualization was as common as it is now, laboratory environments required a lot of hardware, as all hosts computers, servers, network switches and access points had to be physically present and properly configured. Computers require power and space, and thus a dedicated space is required for a teaching laboratory. With virtualization, we can run complex environments, with several servers and host computers in different networks, on a single computer, eliminating a lot of space requirements.

AVAILABILITY: A learning environment should be available to the students as often as possible. While contact teaching is important, repeatedly doing exercises and experimenting on the students' own time is also a very important part in internalizing network security concepts.

RESOURCES AND PLATFORM SUPPORT: A learning environment for network security should support as many different OSs and

configurations as possible. In real network environments there are hosts with different OSs and configurations, and a lab environment should be able to reflect this by providing easy support for changing configurations on hosts and adding and removing hosts from the environment. The environment should have sufficient resources, so that it can provide at least adequate performance in a teaching situation. Some extra resources, whether they are hardware or software, should be reserved for situations where unforeseen events strain the environment.

8.3.1.2 *Previous laboratory environment*

The current layout of our security lab can be seen in Figure 24. We have a dedicated classroom for our security lab, which contains 12 classroom computers for students to use during exercises, and two servers responsible for running virtual machines that make up the actual learning environment. They are interconnected by a gigabit Ethernet network set up in the classroom, with an optional connection to the university network and Internet. We also have the option of setting up a wireless network in the classroom for research and education purposes, such as WLAN security audit and testing. The main server runs an instance of the network security learning environment, which is used on a course on firewalls and IPS technologies directed at master's students in network security. Both main and secondary servers also host additional virtual machines used as computational resources for research purposes.

8.3.2 *Requirements analysis for virtualization*

The first thing to take into account when beginning the process of migrating from physical to virtual environments is to define the requirements for the new environment and what kind of tasks can be assigned to the lab. In our case, we need to provide sufficient resources for a firewall and IPS laboratory environment with over 50 virtual machines in one instance and preferably the ability to run more than one instance of the environment simultaneously. Additionally, a penetration testing environment with several virtual machines should be running constantly, with access from outside the university network made possible to stu-

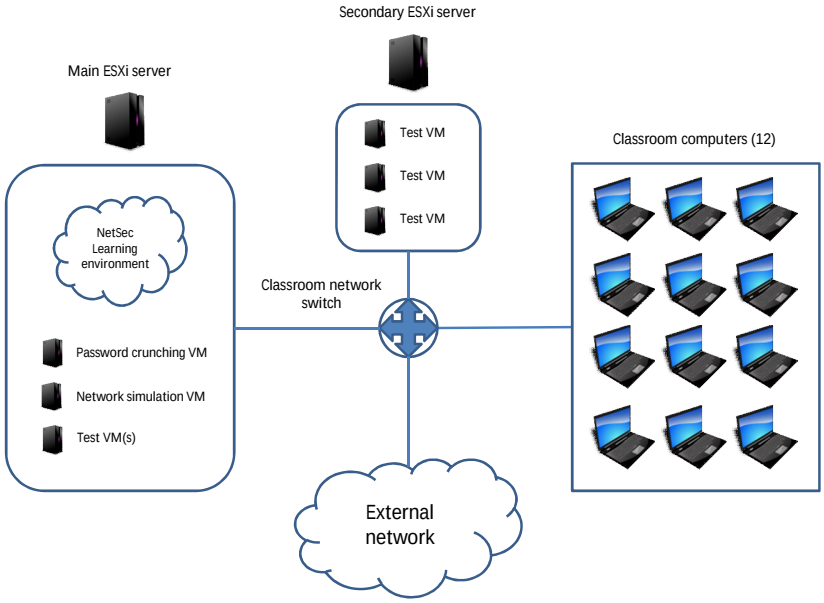

Figure 24: The layout of the old security lab.

dents. Our way of approaching this problem is to define learning environments, of which we run virtual instances in our servers. This makes it possible for many students to be working in the same environment at the same time, thus increasing the capacity of a limited environment.

8.3.2.1 *BYOD and virtualization*

Bring Your Own Device (BYOD) is an IT policy that has been advocated as a way to save costs in IT and provide users a better user experience. Instead of forcing everyone to use a single pre-chosen IT solution, BYOD aims to reduce organizational IT hardware costs by allowing users to use their own preferred devices, whatever they may be, for work. This is also sometimes referred to as Bring Your Own Technology (BYOT), which reflects the inclusion of both hardware and software [210], exemplified by modern mobile devices with different software ecosystems and hardware. BYOD (or BYOT) gives users the interface and software they are familiar with, the capability to use the device they choose with all the benefits, and translates these into produc-

tivity in a work environment. For more discussion on BYOD, its motivations, benefits and challenges, see e. g. [210, 211].

Network security lab education is strictly characterized by software choices made in the lab. Combining this with the goals of BYOD may first seem to be a hard goal to achieve, and some of the known downsides of BYOD such as incompatibility with some infrastructure choices cannot be avoided. But if we wish to reach our goal of doing away with the dedicated lab room for network security, we must provide means for the students to bring their own devices and use them in the lab exercises. This way, the requirements for fixed infrastructure are reduced, as the university does not have to maintain a full classroom of computers for the students to interact with the teaching environment. On the contrary, in a best case scenario, we can eliminate the computer classroom completely, as shown in Figure 25, where various methods for connecting to the virtual environment are shown.

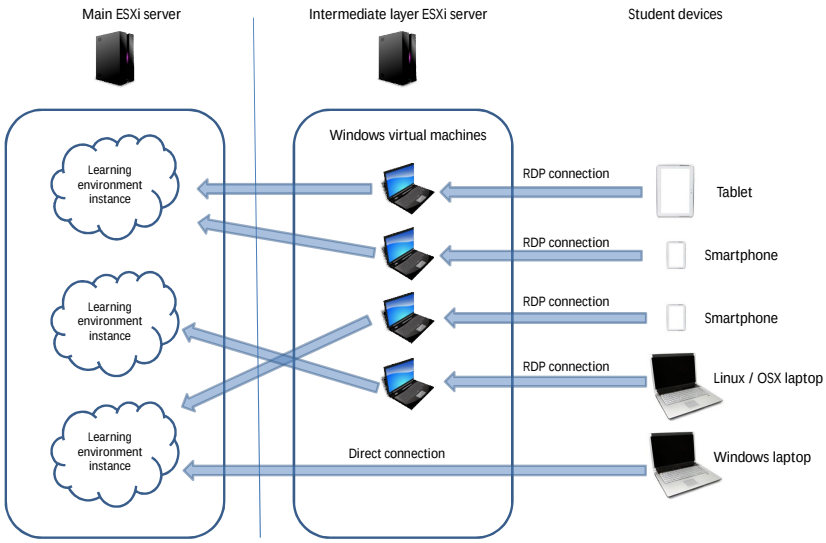

Figure 25: Connection options to the new security lab for students.

Now we can summarize our approach to BYOD security education. We use Windows virtual machines as a translation layer between incompatible devices and lab infrastructure, giving the opportunity to participate in the labs with your own device with most hardware combinations. Windows Remote Desktop Protocol (RDP) is leveraged for this purpose. RDP makes it possible

to remotely control a computer from a device capable of running an RDP client. As RDP has wide support across all operating systems and platforms, and, so far, only Windows is capable of running all necessary software for the laboratory sessions, this makes the combination a natural choice. In Figure 26 we demonstrate this concept with an Android Nexus 5 phone being used to control the lab environment.

Figure 26: A Nexus 5 phone is used to remotely control the lab environment via RDP.

8.3.2.2 *Hypervisor selection*

An important part of the design process is choosing the hypervisor upon which everything will be built. A hypervisor is, simply put, an operating system that is capable of running instances of other operating systems all at the same time. It is the core element of virtualization, and several choices for a hypervisor exist. Virtualization can be achieved with several different solutions, both open source and commercial software. A hypervisor can run on top of a running OS, or on top of "bare metal", i. e. there is no other OS between the physical hardware and the hypervisor. Here we focus on a bare metal hypervisor approach, where we use VMware vSphere[6] as the platform of choice. The same result can be achieved with open source solutions such as XenServer.[7] Configuration, setup and administration would nat-

6 VMware vSphere. Online, available at http://www.vmware.com/products/vsphere.html Accessed 11.6.2017.

7 XenServer. Online, available at https://xenserver.org/ Accessed 11.6.2017.

urally be different, but the basic principle remains the same and both could easily be used to construct such an environment. In our case the choice was driven by administrator familiarity with the system.

8.3.2.3 *Hardware considerations*

We will build our new lab environment[8] using hardware from our previous lab environment, so that we can reuse as much of the previous hardware as possible. We must take into account the fact that the cycle for computer hardware is somewhere around five years, and while obsolete hardware can be reused for some other purpose, for the heavy lifting part we will need to source new hardware.

The current lab consists of 12 Core 2 Duo E8400 computers with 4 GB of memory, a setup consistent with high-end hardware 5 years ago. One E8400 computer has been outfitted with 16 GB of RAM, and it was used as a server for a limited version of the virtual environment. As the main server the lab has an Intel i7-3770 with 32 GB of memory and a SSD hard drive, and this server has been responsible for running the virtual environment thus far. For network connections we have two 24-port gigabit Ethernet switches, several smaller switches, and wireless access points and adapters for adding a wireless network to the environment.

To future-proof our environment for a reasonable time frame, we are considering for the main environment a server with two Intel Xeon E5-2630 processors and 128 GB of memory. We can reuse the old main server and some of the older student computers as servers for the intermediate layer and other auxiliary tasks, giving us more resources for dealing with unforeseen situations.

8.3.2.4 *New network security education environment*

The new education environment is run completely on virtual machines. The only physical components are the main servers

8 Some time has passed between the publication of the original article, and the publication of this dissertation. The development of computer technology knows no mercy, and thus the specifications given here are understandably products of their respective eras of computing.

and a network switch. They can be placed anywhere within university premises, but it is recommended that they are accessible in case of a failure or for maintenance purposes. The student devices use a VPN connection to the university network to gain access to the main server. VPN is a standard technology used for securing connections over insecure networks and providing confidentiality and connectivity to users. One of its most common implementations is granting access to a restricted network through the organization network perimeter firewall to remote users. In our case we use the standard VPN solution for remote access provided by the university IT administration to all students and staff. The lab infrastructure is described in Figure 27.

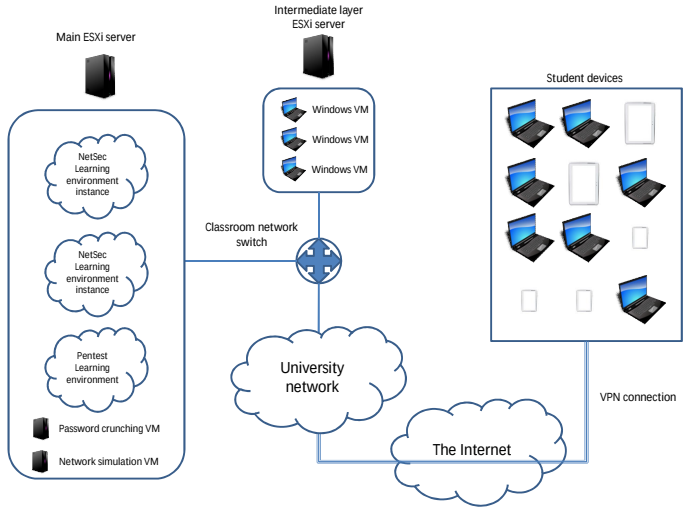

Figure 27: Diagram of the new information security lab.

After this first step, the students can now access the environment. If they require a Windows host for interacting with some components, they can control a Windows VM remotely by using RDP, for which a client is easily available for any operating system and platform. Over the RDP connection the students are essentially controlling a classroom PC, and can use it for interacting with the learning environment.

The learning environments themselves are modular pieces in this architecture. We can modify, add and remove environment

instances as we please, as long as we make sure that the environments are reachable from outside the main server. This is managed through configuring the virtual network adapters in the environment to connect to the correct physical network adapter.

8.3.3 *Effects of virtualization*

The virtualization approach applied to laboratory exercises traditionally done with dedicated hardware and a classroom does also carry some inherent risks and challenges, which we identify and discuss their impact and relevance here.

8.3.3.1 *Student learning experience*

Actually understanding the topology and layout of a network, the location and relations of different host on a virtual network seems to be, in our experience, very difficult for a significant number of students. When we are working with abstract concepts such as virtual machines, it is often hard for a student to understand what computer they are trying to manage at a certain time, or how traffic is routed and networks are connected in virtual environments, as there are no physical references for them to use. When dealing with a physical environment, the instructor can always point out a single computer and say "This computer here runs that server you are now connecting to, and that device over there is running the firewall you are connecting through." Also network connections are easy to visualize when students can physically follow the network cables from host to switch to another host, if necessary.

In our experience, this is significantly harder to visualize for students working with a virtual environment. Students are more often configuring wrong virtual machines, or fail to understand how a network is connected, when asked to debug a non-working configuration. As we can see, losing touch with physical hardware can be a problem for some students. While the truth is that most people do not have to deal with networks more complex than a simple home network, students in the information security and network security field should be at least familiar if not comfortable with administering non-trivial computer net-

works with more than one or two hosts, whether they are made of physical or virtual components.

The User Interface (UI) on a touch-based device differs significantly from an UI designed to be used with a keyboard and a mouse. This can be a problem for some students who use a touch UI to interact with the virtual machines, causing problems with typing on an on screen keyboard, a smaller display, and so on. These kinds of challenges limit the usability of the environment for those who do not have a traditional laptop computer.

Virtualization also makes it possible to integrate lab exercises with regular classroom teaching. Combining lectures and lab sessions normally requires arranging the lecture in a dedicated computer classroom where we can provide a computer for every student with the necessary software, but with the ubiquity of sufficiently powerful devices, we can take advantage of this by embracing BYOD. Now that every student can use their own device to interact with the training environment during a regular lecture, we can integrate exercises into lectures more efficiently. This makes it possible, for example, to use a 30 minute segment of a lecture to talk about firewall rule table creation, and use a 15 minute segment for students to do a related hands-on lab exercise where the students create their own rule tables in the learning environment. With virtualization, this can be done with the student's own devices, in the classroom, without a dedicated computer lab.

Student-instructor interaction is a vital part of successful lab sessions, and it first may seem that transitioning to virtual environments decreases the amount of contact teaching, but it does not have to be so. This will only be an issue with lab exercises that students can do any time they wish, and even then there are several methods for interaction between students and faculty. Options range from on-call hours at the office, Skype or other video conferencing meetings with students, message boards, e-mail messages, to instant messaging software and Internet Relay Chat (IRC) chat rooms, all depending on the circumstances. These means do not supersede or make obsolete contact teaching, but rather make it possible for at least some kind of interaction in lab exercises with independent schedules.

8.3.3.2 *Education environment engineering*

Virtual machines are much more flexible and usable from a lab engineer point of view. Simply the ability to have snapshots of a virtual machine state, having more than one snapshot available, and being able to switch between snapshots makes it possible to maintain several instances of an education environment concurrently. If we are working with normal computers, when a student performs a lab sessions, manually reversing all changes that have been made can be prohibitively time-consuming. Non-persistent virtual machines offer computers that can be used for a lab session and can be reverted to their original state for the next student simply by restarting them, as no changes made to the system are actually written on disk. After several years of experience on administering both virtual machines and physical hosts, we can with reasonable confidence say that administration and environment engineering is significantly more streamlined and less taxing in a virtual environment. The instructor workload in supervising students can be difficult to compare between these scenarios, but in our estimation it is approximately equal. Virtual machines give the opportunity for the instructor to remotely assume control of a host for troubleshooting purposes without moving around in the classroom, but otherwise the workloads can be estimated to be similar.

8.3.3.3 *Technical considerations*

We must also consider the requirements this setup poses for the network used by the students to connect to the environment. If we have 20-30 students all connecting via the same Wireless Local Area Network (WLAN) access point, this will definitely cause traffic congestion in the network. This would be a typical situation for a standard classroom setting, so steps must be taken to ensure that we have sufficient bandwidth available for all students. Some mitigating factors can be students who use their own ISP connections on their own devices, and dedicated WLAN access points configured to provide access for students. Network connection quality and availability is essential for this virtualization scenario to work, so we should assign sufficient resources to ensuring that the network will not become the bottleneck.

Of course some students will not be able to participate with their own devices due to technical difficulties, malfunction, not having a suitable device in the first place, or forgetting to bring one to a session. This problem is easy to manage with a small number of laptop computers that can be temporarily assigned to students who would otherwise not be able to participate in a session.

Perhaps the most salient threat scenario for this kind of environment is hardware failure. As it is with any hardware, equipment failure can result in critical data loss, possibly erasing all progress made by students in a session. Especially when we are working with a server that runs critical software, catastrophic failure is a possibility. A distinct benefit for standalone hardware for all hosts is that the risk of hardware failure is distributed, leading to a lower probability of data loss in case of failure. This risk can be mitigated with redundancy in server hardware, but only to a certain degree. We must eventually accept some level of risk in such an environment regardless, and with good industry standard practices these risks can be mitigated sufficiently.

8.3.3.4 *Financial considerations*

While our foremost goal is not to consider financial incentives, we must note that depending on circumstances, this can also be considered as a driver for virtualization. For departments that have to spend money on facilities, reducing the amount of real estate on lease can bring cost savings. In our case, these savings are in the order of tens of thousands of euros, but this naturally depends on facility costs that can and do vary between locations. Also, maintenance of a computer classroom, replacing outdated and malfunctioning hardware, also takes resources, both in the form of money and time spent on administering the classroom computers

8.3.4 *Conclusion*

By making laboratory infrastructure platform agnostic by using virtual machines as a translation layer, we can provide a similar experience on all client devices to all students, regardless of the platform that each student has chosen. This not only provides

means to offer lab education independent from a classroom, but also makes it possible to students to do the lab exercises on any platform, any time they choose. The benefits are clear both for students and faculty. The students get a more flexible approach to the exercises and the possibility to use their own preferred devices in education, and the faculty gets better education results, increased throughput for student exercises, and possible financial savings in reduced premises and classroom computer maintenance. The risks of the virtualization approach have been identified to be mostly dependent on hardware failure, and with sufficient backups and good practices in place, we can conclude that the benefits of virtualization in network security laboratory education far surpass the risks.

8.4 STUDENT PERCEPTIONS ON INFORMATION SECURITY

Before we can begin to improve information security education, we must establish a baseline on how university students are capable of perceiving information security issues and threats. For this purpose, an analysis is made on the perceptions of security and privacy from mostly first year university students. The study is performed on data from an introductory level course (Tietotekniikka ja yhteiskunta, tr. *Information technology and society*), with students mostly from computer science and software engineering degree programs. A small minority of students are from other degree programs as well. The language of education on the course is Finnish, and all students are native speakers.

During the course, the attitudes and perceptions of students were measured with a weekly writing task where the students had to provide answers to questions relating to privacy and security. The data is gathered from two course iterations, in 2012 and 2013. After 2013, the course has changed somewhat, including a change in the course name, but the questions for group work and individual assignments are relatively similar in nature. Analysis of the data from these later instances of the course is outside the scope of this thesis.

8.4.1 *Research methodology*

The original data that is used as the basis of this analysis is in the form of free-form text, originally submitted by students as answers to written questions. To make it possible to analyze this kind of data, *content analysis* [212] is used to extract data from these student answers. Each answer to a question is thus referred to as a document, and each group task answer is considered to be an individual document, with the exception of the analysis of the individual World Wide Web (WWW) footprint, where each student in a group has provided their own answer, and each answer is thus treated as their own document. Each document has their own document identifier, making it possible to easily differentiate documents from each other.

The data used in this research is gathered from group assignments done during the course, with 12 answers in the 2012 course instance and 10 in 2013. The total sample size is small *(n=22)*, so these results cannot be reasonably generalized without more data and further research into the topic material.

All in all, eight measurable quantities were identified. The following student capabilities, perceptions, or characteristics were identified as vital for security education on the introductory level, and were subsequently analyzed based on the data. First the capability for understanding security and related issues was measured.

1. CAPABILITY FOR SECURITY THREAT ASSESSMENT. How well students have demonstrated the capability to recognize and analyze security threats in various scenarios.

2. CAPABILITY FOR OBSERVING FROM ANOTHER PERSPECTIVE. How well the students are capable of distancing themselves from their own perceptions, and adopting a different viewpoint in the process of security analysis.

Another aspect of security that was measured from the data was the size of the WWW footprint students, measured on an individual level, and the group level understanding of the significance of that footprint.

3. SIZE OF WWW FOOTPRINT. The size of the personal WWW footprint of each individual student in the group.

4. UNDERSTANDING OF THE WWW FOOTPRINT. How well the students have understood the significance of a WWW footprint, and what security implications various WWW footprints have.

Additional analysis was performed on student perceptions on surveillance. Here, the sample size is even smaller *(n=10)*, so the results are only preliminary in nature. The following perception categories were identified from the data.

5. PERCEPTION OF ETHICAL CHALLENGES WITH SURVEILLANCE. How well students are able to identify and analyze ethical problems with surveillance and mass surveillance.

6. GENERAL PERCEPTION OF SURVEILLANCE. What is the general disposition of the students towards acceptance of surveillance and mass surveillance.

Finally, the students were asked to provide their insight on surveillance. From their answers, positive and negative consequences of surveillance in society are ascertained. Three main consequences, if more than one was mentioned, from both categories are considered. These answers are analyzed based on the ranking (based on order of appearance and/or contextual significance) and thematic area (societal or technological).

7. POSITIVE CONSEQUENCES OF SURVEILLANCE. What are the positive effects of surveillance?

8. NEGATIVE CONSEQUENCES OF SURVEILLANCE. What are the negative effects of surveillance?

8.4.1.1 *Student assignments*

The assignments that were given to students are described here briefly. Instead of providing the student assignments verbatim, only the essential content, translated from Finnish to English, of each assignment is detailed.

- ASSIGNMENT A: Describing interesting applications in generally understandable terms, and analyzing success factors and technological requirements.

Table 14: Aspects measured from student answers.

Measured aspect	Scale used	Assignment(s)
Multiple perspectives	Likert 1-3	A,B
Threat assessment	Likert 1-3	D
WWW footprint size	Likert 1-3	C
WWW footprint understanding	Likert 1-3	C
Perception of ethical challenges	Likert 1-3	E
General acceptance of surveillance	Likert 1-3	E
Surveillance positives	Ordinal	E
Surveillance negatives	Ordinal	E

- Assignment B: Describing a hypothetical application or product, technologies it would require, and potential user demographics.
- Assignment C: Measuring the WWW footprint of group members.
- Assignment D: Security analysis of interesting applications.
- Assignment E: Analyzing mass surveillance from technological, societal and political viewpoints.

8.4.1.2 *Methodological issues*

The analysis is performed by a single reviewer in this study, so there is no cross-referencing of reviewer results and no difference matrices are derived from the results. This limits the applicability of these results due to limitations of the methodology. This is taken into account in the analysis of the results, but it is clear that this is a preliminary analysis — a proof-of-concept study — and that more comprehensive effort is needed.

8.4.2 *Results*

In the study it was found that the students had a relatively good grasp of potential security threats, but lacked skills in considering issues from various perspectives. This was evidenced by

Table 15: Aspects measured from all group answers (n in parenthesis)

Category	High	Medium	Low
Multiple perspectives (22)	7	10	5
Threat assessment (22)	14	6	2
WWW footprint size (100)	14	45	41
WWW footprint understanding (22)	12	8	2
Perception of ethical challenges (10)	6	4	0
General acceptance of surveillance (10)	3	4	3

good capability to assess threats and to understand potential threats by caused by personal WWW footprints. The size of the web footprint — we refer to any publicly available information of a person that can be connected to a person as the web footprint — of each individual student was assessed from the source documents. 41 students had a small to non-existent footprint, 45 had a moderate web footprint, but only 14 had an extensive web footprint. Students in computer science and computer engineering thus have on average a relatively modest digital footprint.

Student perceptions on the ethical challenges of surveillance were very acute. No student group was considered to have a poor understanding of ethical issues associated with surveillance, discussed earlier in this thesis in Chapter 2. The perception on the acceptability of surveillance was evenly distributed, with both critical and supportive views expressed in the source documents. This result would seem to imply that CS and CE students have a good understanding of the associated ethical issues, but the opinion on whether surveillance is seen as a resource or as a threat to society is clearly divided. The results are shown in detail in Table 15.

Perceptions on consequences of surveillance on a societal level were measured in the final question. Student groups were asked to provide up to three likely consequences, both positive and negative, of surveillance ranked in order of importance. The free-form answers were gathered and grouped under categories shown in Table 16.

Each category was given a score weighted simply by the order of importance, and then sorted according to the end result.

Table 16: Student perceptions on positive and negative consequences of surveillance.

Consequence	First	Second	Third	Score
Solving crime	4	2	0	16
Preventing terrorism	3	1	0	11
Increased security	2	0	0	6
Emphasis on societal issues	1	1	0	5
Potential data mining	0	0	1	1
Erosion of societal trust	3	3	2	17
Chilling effect	3	0	0	9
Abuse of information	1	2	2	9
Expensive to maintain	2	0	0	6
Reduced societal cohesion	1	1	0	5

An answer ranked first is worth three points, second gets two points and an answer ranked third gets one point. By summing the points in each category, we are able to ascertain the most important positive and negative societal level consequences for surveillance from the point of view of CS and CE students. Potential for fighting crime and terrorism were seen as highly positive results. Negative consequences were dominated by erosion of societal trust and the potential chilling effect caused by surveillance.

It is evident from these results that the students in this sample have a realistic perception on topics related to information security, both on a personal and also on the societal abstraction level. More generalized conclusions cannot be made from this source data.

8.5 CONCLUSIONS AND FUTURE WORK

In this chapter we have presented discussion and improvements to university-level information security education. A well-defined security knowledge framework is presented and has been tested in practical security education. Hands-on exercises are an important part of an information security curriculum, and experiences

from university-industry cooperation in providing this opportunity to students is discussed. Virtual environments for realizing network security education are designed and the design process is discussed at length.

Future work in the area of information security education includes taking the presented model of security education to primary and secondary level education, and measuring how the security awareness and behavior evolves. The best scenario is to have a populace aware of network and information security threats, making the work of malicious actors difficult in the networked information society.

Student digital footprints could be estimated more robustly in future studies, using for example methods that create results that are comparable between different studies. In this study this was not done because content analysis as a method does not facilitate deriving sufficient data from free-form student answers. Thus a new, improved design for this experiment is needed. The analysis of student perceptions is an initial study, mainly aimed at validating the research methods used in the study and pointing into potential research directions. More data is needed regardless in order to have potentially generalizable results. Further analysis based on the available source material using better statistical analysis methods and more extensive verification of the research model is left for future work for researchers in the field.

9 TOOLS FOR SOFTWARE DEFINED CRYPTOGRAPHY

> SHEPARD. *"Remember the old days when you could just slap omni-gel on everything?"*
> LIARA. *"That security upgrade made a lot of people unhappy."*
>
> – *Mass Effect 2* (2010)

In this chapter we present tools, methods and concepts that can be used to improve the security of network communications in the networked information society. Communication security for IoT and similar applications with low power requirements combined with security is enhanced by creating a framework and design flow for secure embedded devices. A survey of Instruction Set Extension (ISE) based methods for accelerating cryptographic primitives is presented, and the concept of adaptive cryptography is discussed as a potential solution to the attacks against cryptography described in Chapter 3. This chapter is based on the authors' contribution[1] to the following publications: [16, 18, 17].

9.1 SOFTWARE DEFINED PLATFORMS FOR ADAPTIVE CRYPTOGRAPHY

The rise of cyber crime in the first decade of the 21st century has provided evidence that there are adversaries willing to compro-

1 Hakkala, Antti and Isoaho, Jouni and Virtanen, Seppo. Towards Adaptive Cryptography and Security with Software Defined Platforms. *Computing Platforms for Software-Defined Radio*, Springer, 2017. Used with permission of Springer.

mise any and all assumptions of security, and they do it for profit. Cyber crime in itself has grown to become a multi-billion business, and their methods range from social engineering to capitalizing on existing vulnerabilities in software and hardware [213]. In the battle against cyber crime the main focus should be in building security in the systems and maintaining the ability to adapt to discovered vulnerabilities, thus swaying the balance in the continuous security arms race to the defenders' benefit.

When a vulnerability is detected in software, it is patched with a software update. If the flaw is in a hardware implementation of a standardized algorithm, and that particular algorithm in itself is found to be compromised, the flaw cannot be fixed without changing the hardware itself. It is hard to convince consumers to upgrade their hardware because of such a security flaw. It is in most cases also not cost effective in the consumer device segment. In more specialized or mission critical hardware, however, replacement can be worth serious consideration. This can be true even in cases where consumer devices would be left in a vulnerable state through their lifetime, e.g. WLAN access points supporting only outdated encryption options such as WEP.

Traditionally different types of communication have required specific terminal equipment and infrastructure, all specific to a single application. Placing phone calls typically required a phone and either a wired or wireless telephone connection. Web browsing and reading email required a computer and a network connection. Watching television required a television set, and so on. Nowadays the situation is significantly more complex. While we do have handheld battery-powered devices capable of handling all of these communication scenarios, the pressure for security in communications has not yet reached the point where all communications regardless of type and platform are encrypted by default.

Ubiquitous encryption places even more requirements on multipurpose devices such as smartphones, but the current generation devices are already more than capable of handling this. One significant caveat is whether better security and privacy should come at the cost of end user experience. In a consumer product this could be considered unacceptable by both device manufacturers and users alike.

Future communication systems need to support a variety of signal processing algorithms, radio interfaces, and communication and security protocols in the same device, with the capability to adapt to changing situations. The present trend towards provision of high processing speed for a wide variety of applications by handheld battery-powered devices has been observed by Björkqvist *et al.* [214]. This trend can already be realized in current smartphones and software ecosystems. With the advent of Internet of Things (IoT) and Wireless Sensor Network (WSN) systems, smaller devices with less computational capacity than current handheld devices will have to fulfill the same requirements for secure communication as present handheld devices. We see that a new trend driven by data security and privacy concerns will require more security related functionality from all devices, regardless of their position in the communication chain – whether they are measuring, processing or just transmitting data. For example, sensors used in healthcare applications of IoT must use strong methods to secure medically sensitive data, and military WSN applications must adapt to changing environments and corresponding security requirements.

9.1.1 *Previous research on embedded security*

The field of embedded security is under intense research. Some challenges, as identified by Isoaho *et al.* [215], are e.g. securing the data content and data transportation within the device, data exchange with the recipient device at the other end of the wireless link, and data transportation between end-to-end applications. These are vital capabilities and indispensable built-in features for any communication device that aims to serve as a credible solution for current and future communication applications. Software Defined Radio (SDR) and Open Wireless Architecture (OWA) have previously been established as potential design paradigms to address the needs for higher system integration and multiple communication interface support.

We specify a design methodology and process as a framework for enabling resource constrained devices both to implement necessary security features, and to preserve processing power to the actual target application. We refer to this as the Software-Defined Secure Communications (SDSC) framework. We examine

the augmentation of the SDR and OWA concepts with built-in security support, and present our approach to building security in the framework. We identify and analyze the requirements set by modern security algorithms and applications to future communication systems, and further specify a design methodology which leverages SDR principles, adaptive cryptography as well as software and reconfigurable hardware for such systems. As a result of the analysis, we discuss the potential benefits of this approach in the future, where IoT and sensor networks are ubiquitous. Essentially we see our method as a flexible, programmable and security oriented design framework for security enabled SDR applications that can be used to create proprietary hardware and software designs, which can later be adapted to several different functions with software. We also see benefits of our method in shortened time-to-market and potential cost savings in increased volume and customization options.

In this work we use the programmable and parameterizable Tools for Application-specific hardware/software CO-design (TACO) hardware platform [15, 216] for protocol processing applications. TACO has previously been extended with a Digital Signal Processing (DSP) extension [217], making it possible to run protocol processing and DSP tasks in parallel on the same programmable processor with optimized execution units. SDR capabilities have also been introduced to the architecture by Anwar *et al.* [218].

9.1.2 *TACO platform for application domain specific programmable acceleration*

The TACO hardware platform is based on the Transport Triggered Architecture (TTA) paradigm [219, 220] and is aimed at developing application domain specific programmable processors. A processor customized for a specific application does not only outperform a general purpose core, but is also expected to consume less power and circuit area. Application-specific hardware can reduce the computational load associated with complicated operations that are cumbersome when executed on traditional processors. A customized processor can meet the desired throughput requirements at a lower clock speed, thus reducing energy expense. Unlike a general purpose execution unit, a highly tailored processor will not implement any extra functionality, except for

the subset needed for the desired application domain. Because of the TTA based architecture, TACO protocol processors are programmable and configurable, providing dedicated hardware to deal efficiently with application-specific tasks.

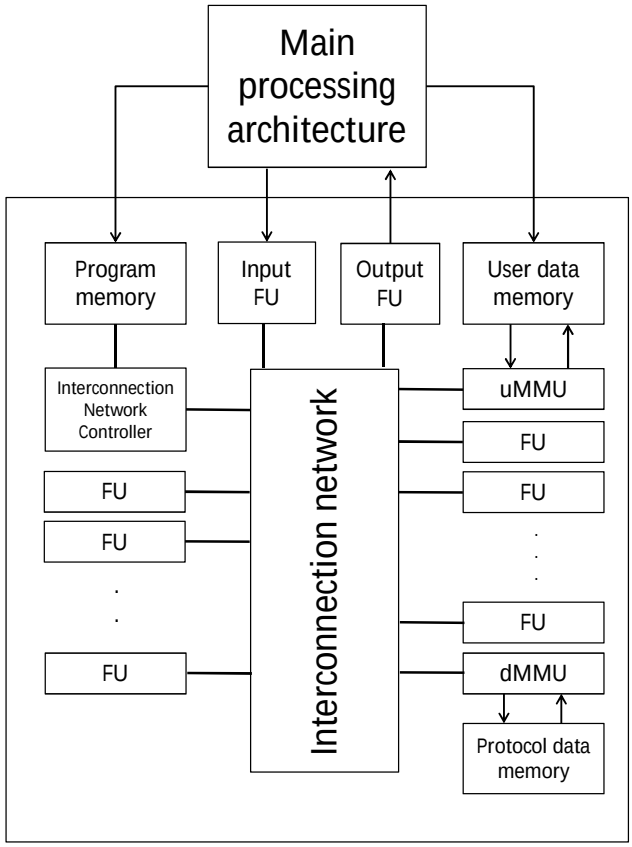

Figure 28: General TACO processor architecture. Used with permission of Springer.

As seen in Figure 28, a TACO processor has a main component: the Functional Unit (FU). A TACO processor consists of a set of FUs connected by an interconnection network of one or more buses. There may be multiple FU instances of the same type in a TACO processor, allowing parallel execution of respective operations. This resource redundancy, depending on circumstances, may serve either to speed up the execution of a particular single application or to provide real multitasking for a set of simultaneously run applications. The interconnection network is com-

posed of one or more buses, supported by an interconnection network controller. The interconnection network is presented in more detail in Figure 29, where the individual connections between FUs and buses are shown. The number of buses directly limits instruction level parallelism, as it determines the maximum number of simultaneous data moves. As TACO is a modular architecture, FUs are designed separately and are provided standard interfaces to the interconnection network. This also makes it easy to automate the design process to a high degree. Each FU implements a set of functions, and the final configuration is defined by the choice of appropriate FUs to the target application. FUs can be modified internally as long as they provide the same interface to the sockets connecting them to the interconnection network. This makes module re-usability straightforward and simple, and it is an integral part of the TACO hardware platform and design methodology.

The performance of the architecture can be scaled up by adding extra FUs, buses, or by increasing the capacity of the data transports and storage. Increasing the number of existing FUs provides more performance for parallel and pipelinable tasks, while adding FUs with new functions in turn provides support for new application types. Performance of TTA processors is at least equal to that of their conventional or DSP counterparts (i.e. processors of different types but with functionally and quantitatively equivalent capacity). According to [221, 222], this performance is also achieved at far lower cost, which encourages further development of optimized functionality on the TACO platform.

TACO is programmed by only one type of instruction, *move*, which specifies independent data transports on each of the defined buses. In contrast to traditional processors, data transports trigger operations, and not vice versa. The operation of the processor occurs as a side-effect of the transports between functional units, transparently to the control unit and its instructions. The architecture allows for one processor cycle to fit several bus transports. The interconnection network controller implements the data transports on the buses. The controller uses a program memory, from which it fetches instructions, splits them into sub-instructions, and dispatches each sub-instruction onto the corresponding bus. Each move sub-instruction has a source and a destination field, and the data is carried accordingly from one FU

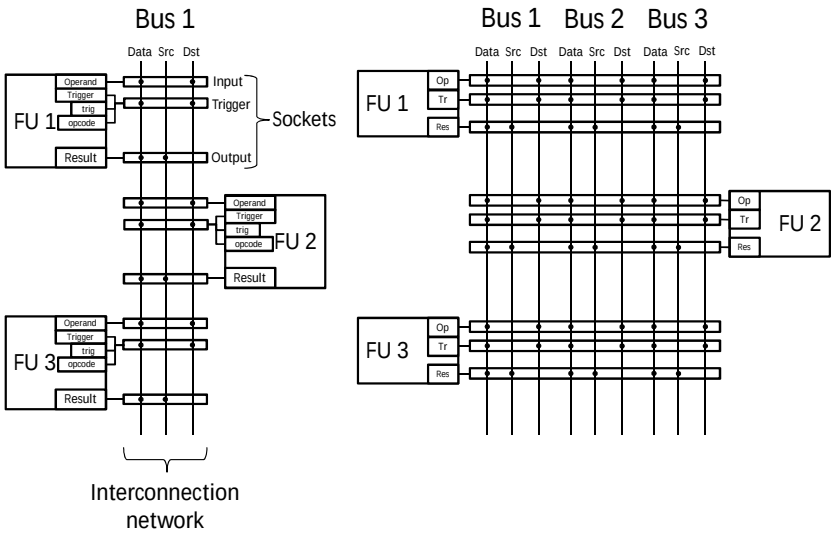

Figure 29: TACO processor interconnection network. Each interconnection bus consists of a data bus, a source address bus and a destination address bus. *Left:* Three FUs and one bus: only one data transport per clock cycle. *Right:* Three FUs and three buses: three data transports per clock cycle are possible, so each functional unit can be triggered to function on each clock cycle. Used with permission of Springer.

register to another on the respective bus. When adding new FUs to a design, the instruction format does not change as long as the existing FUs are addressable by the length of source and destination addresses, and unless more parallelism is required. The excellent pipelining capacity of TTA, together with the emphasis on moving data make it a good platform for data-intensive applications. It is exceptionally well suited to protocol processing and for all communication-related tasks in general.

9.1.2.1 *Protocol Processing Application Domain*

The first application domain addressed in the development of TACO was protocol processing. The following discussion on building the TACO platform for this application domain is based on earlier research in [220, 15, 216].

In order to develop customized devices for the protocol processing application domain commonly used protocols and pro-

tocol processing applications were carefully studied. Special emphasis was placed on finding functionality that varies very little from one protocol to another or remains virtually the same across a set of protocols. Studies by Jantsch *et al.* [223], Virtanen [15], and Truscan *et al.* [220] identified characteristic functionality that, if parameterized, is generalizable to several protocols instead of just one specific protocol. These studies identified, for example, the need for bit pattern matching and masking, substring replacement in bit strings, Boolean comparisons, counters and checksum calculations. They are also functionality that can be parameterized. One finding was also that protocol processing can be implemented using only unsigned arithmetic (i.e., there is no need for managing negative values), making hardware implementation considerably simpler. The emphasis in these studies was on analyzing protocols that can be regarded as layer 1-3 protocols (physical, data link, and network layer) in the OSI reference model. In these layers the protocols are not end-to-end protocols, but require intermediate stations (e.g., repeaters, bridges, switches, routers etc.) between the source and destination terminals. Thus, these protocols present a clear need for application-specific hardware systems in addition to application software, whereas the end-to-end protocols in OSI layers 4 and above are often completely implemented as software running on networked workstations. However, the higher layer end-to-end protocols can also potentially take advantage of application-specific processors, particularly in mobile devices and low-power sensor networks. In terms of the work presented in this chapter, we argue that we cannot concentrate solely on the three lowest OSI layers, as security-related protocols reside on the upper layer of the OSI model. Especially important are protocols providing end-to-end encryption, an essential part of information security.

9.1.2.2 *TACO Extension to Digital Signal Processing Domain*

Paakkulainen [217] explored possibilities for augmenting the TACO hardware platform to support DSP operations, in addition to the pre-existing support for protocol processing. As the first step towards bridging the gap between the protocol processing and DSP domains, the TACO hardware platform was enhanced

to support finite impulse response (FIR) filtering. To enable efficient filtering the multiply-and-accumulate operation typical to many DSP applications was implemented. The approach chosen in this study required remarkably few modifications to the TACO hardware platform. The process of designing and implementing a new FU for the multiply-and-accumulate operation demonstrated the flexibility of our hardware platform in terms of adding support for new functionality even across the boundaries of supported application domains. As our initial starting point in the design of the DSP domain multiply-and-accumulate FU, we used an existing protocol processing FU, namely the Internet check sum calculation unit presented in [15].

Fine-tuning of the TACO modules allowed us to enhance our protocol processor architecture with support for some DSP applications without loss of performance in either domain, and at a very low cost. We determined that the Multiply and Accumulate (MAC) FU consumes about 5-10 % of the total chip area of a typical TACO processor implementation. With a single MAC FU hardware implementation we were able to specify multiple system implementation schemes with different optimization factors for the target application. The conducted analysis suggests that the obtained enhancement of TACO platform allows efficient parallel execution of application software requiring both protocol processing and DSP operations, and this was also demonstrated by experimental results. This work laid foundation for further studies on parallel execution of operations from DSP, protocol processing and other domains on our hardware platform. With these encouraging results we proceed to the software defined radio domain.

9.1.3 *TACO as a platform for Software Defined Radio*

Anwar *et al.* [218] took the cross-domain customization of TACO further and explored both methodological and technical aspects of adding support for Software Defined Radio domain to the architecture. Their goal was to deliver a novel software defined approach for designing and implementing common baseband processing tasks, with focus on exploring the algorithmic and architectural design spaces of 3G and 4G systems. Specifically, we identified computational and geometric structures shared by di-

verse coding schemes, services and hardware platforms related to SDR domain, and integrated physical layer support for them in extensible hardware. While the primary goal was to enhance the TACO platform with the features of a programmable SDR enabled processor, in this process the methodology of extending the TACO platform to support new application domains was also realized. This process is done in two phases. First, the initial design is done using standard design principles and methods for the original target application. After this the design is extended to support a new set of requirements by identifying target components in the design and modifying them to fulfill the requirements from the new application domain.

The approach incorporates control structures, component abstractions and parameterization, and architectural optimization into a system design process. We found the proposed platform and design methodology to be very suitable for SDR domain applications, as they have strict requirements on power consumption, chip size, processing speed and scalability for other existing and future applications and standards. Since the TACO platform has the scalability and modularity required by SDR applications, the design framework allows the designer to focus on the internals of the required hardware components due to the well-defined structural entities of the hardware platform. A key benefit of our approach is the support for making FUs definable in software. Once a new radio application has been mapped and partitioned to existing FUs, the designer typically needs to make only minor changes or additions to the internal implementation of existing units. In many cases, modifying the external interfaces of the functional units is not necessary when adding support for novel functionality. These characteristics result in a shortened turnaround time and less design effort in the process of upgrading the platform to support new (radio) applications. On the architectural abstraction level, the designer typically does not have to modify the platform implementation, except to choose the utilized level of execution parallelism and of data transport capacity. Parallel execution of operations can be realized by adding several FUs of the same type to a design, and data transport capacity is affected by the number of buses and the interconnection network.

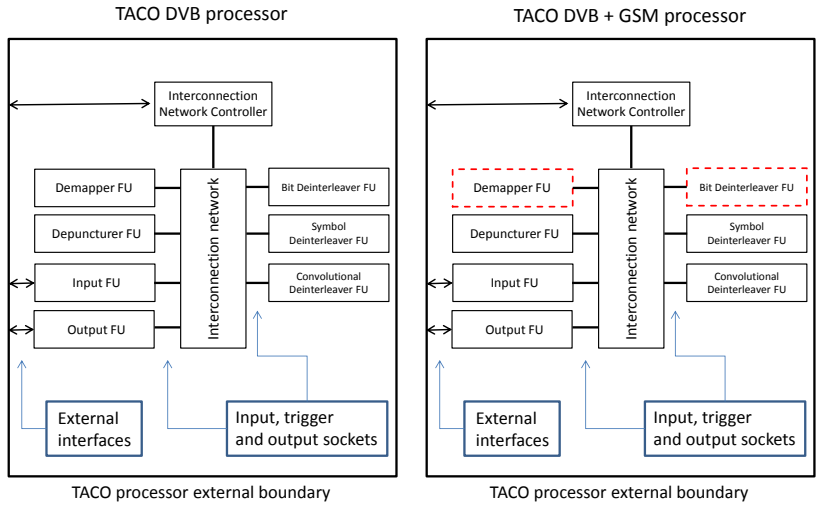

Figure 30: TACO DVB + GSM processor. Used with permission of Springer.

Anwar *et al.* [218] designed a TACO platform for Digital Video Broadcasting (DVB) and then modified the system to additionally support GSM. The platform implementation required 4000 lines of VHDL code. As expected, it was found that extending the platform from supporting only DVB to also support GSM did not considerably increase the amount of code, as most additions were needed into specific processes inside the FU implementations. Only two FUs needed to be enhanced for GSM support: demapper and bit deinterleaver. These FUs are marked with a dashed line in Figure 30. The cost of extending one functional unit to support additional, software definable functionality from another standard was an area increase of only about 10-15 % and a power increase of 23-26 %, showing that these existing units already supported most of the new functionality. At the processor level, the GSM extension produced only 2.4 % of total processor area and 2.7 % of overall power consumption. The extension required four days of hardware design and verification, two days for each FU that needed modifications.

9.1.4 *Towards Software Defined Secure Communication — Multi-domain Integration for secure network applications*

Now the TACO platform and design methodology is extended towards a comprehensive communications solution. The Software-Defined Secure Communications (SDSC) platform will provide support for all aspects of communication, and we start outlining the platform by examining necessary building blocks for secure communication systems, examining the characteristic functionality found in security algorithms, and discussing on how to adapt to a changing security environment.

9.1.4.1 *Cryptography Fundamentals*

Cryptography is a core element of security. All secure protocols and communication standards include methods for guaranteeing data confidentiality, integrity and authenticity. They are the essential cornerstones of information security, so we will have to extend the TACO platform to handle these vital areas. Confidentiality is provided by encryption, while integrity and authenticity are provided by error correction, cryptographic hash functions, Message Authentication Codes and digital signatures. The choice of a cryptographic solution for a target application is not trivial, as there are many choices, constraints and metrics for the designer to consider. The communication cipher suite for an IoT healthcare solution will differ significantly from a smartphone due to differences in *i.a.* the operational environment, security requirements, and physical limitations of the devices.

Cryptographic ciphers can be broadly classified into three distinct groups:[2] symmetric algorithms, asymmetric algorithms, and cryptographic hash algorithms [73]. Symmetric algorithms use the same key for encryption and decryption, while asymmetric algorithms have two keys; a public key for encryption and a private key for decryption. Hash algorithms compress an arbitrary length input to a fixed length output, and attempt to provide unique outputs for each input. All three types of algorithms are

2 The terminology used can be inconsistent. Some refer to symmetric algorithms as private key algorithms, and asymmetric algorithms as public key algorithms. In turn, asymmetric algorithms have a key pair that consists of a private key and a public key.

important for secure communication protocols. Asymmetric algorithms are used for initiating a secure communication channel over an insecure medium. Keys for a symmetric algorithm are then exchanged over this secure channel, and that symmetric algorithm is used for bulk data encryption. The integrity of communications can be verified using secure hash functions, which can provide proof that a message has not been altered in transition. Cryptographic hash functions are also often used to derive session keys from master encryption keys, using key derivation schemes such as PBKDF2[3], and bcrypt [224].

9.1.4.2 *Accelerating cryptography*

There are several different approaches to accelerating cryptography algorithms [225]. Because of large operands and complicated operations that are ill-suited for standard computer architectures, software cryptography implementations tend to be slow when compared to hardware. Thus the natural approach is to build cipher-specific hardware. Dedicated hardware solutions are orders of magnitude faster than software implementations, but they have some significant drawbacks. Hardware implementation of an algorithm cannot be changed later, should the need arise, thus locking us to that algorithm implementation. Reconfigurable hardware such as the Field Programmable Gate Array (FPGA) have been used to solve this issue.

Hakkala and Virtanen [16] have surveyed the methods for accelerating different cryptographic algorithms. Exponentiation operations are very common in cryptography, and combined with modular arithmetic, they provide the backbone of ciphers such as RSA [226]. Since these operations can be regarded as bottlenecks, accelerating them will boost the performance. Known efficient multiplication algorithms and different number representations such as Montgomery representation [227] for modular arithmetic can be leveraged for significant gains. The Chinese Remainder Theorem can be applied specifically to RSA, providing up to four times faster operation [228]. Finite field arithmetic is crucial to Elliptic Curve Cryptography (ECC) [229, 230]. ECC operates in finite prime or extension fields. In prime fields, operands are reduced modulo a prime, and in extension fields the reduc-

3 Defined in IETF RFC 2898, https://tools.ietf.org/html/rfc2898

tion is done modulo an irreducible polynomial. A representation of the arithmetic hierarchy for ECC is presented in [231], where the relationship between the arithmetic operations and underlying primitives is shown. Elliptic curve operations are all based on modular arithmetic, and this makes modular arithmetic an essential underlying operation for ECC as well.

Similarly, finite field arithmetic is also essential for AES [75], in which calculations are done in the binary extension field $GF(2^8)$. In some cases leveraging different coordinate systems for representing field elements can eliminate the need for computationally costly operations. As an example of reuse of cryptographic primitives, AES calculations can be accelerated with an Elliptic Curve instruction set designed to accelerate calculations in finite fields of characteristic 2^n [232]. Another example is the use of the Intel AES instruction set [233] to accelerate AES-based SHA-3 hash function candidates, an approach which leverages existing hardware on processors to accelerate the execution of structurally similar algorithms. A summary of essential operations in various cryptography algorithms is presented in Table 17.

9.1.4.3 *Previous Work on TTA and Cryptography*

As far as we are aware, there is very little literature on the use of transport triggered architecture based systems in the cryptography domain. TTA has mostly been used for protocol processing, and so far the available literature clearly reflects this. Hardware implementations of IWEP, RC4, and 3DES all based on TTA are analyzed in [221]. The paper aims to explore Instruction Level Parallelism (ILP) for the previously mentioned algorithms, finding that parallelism can be exploited in some limited scope. The suitability of TTA for efficient AES encryption is studied in [243], where AES on different platforms, including TTA based processor, for wireless sensor networks was examined in a survey study. A coprocessor for elliptic curve cryptography acceleration based on TTA is presented in [244]. Here, the presented system is targeted at the Chinese standard ECC algorithm SM2. Their implementation was capable of scalar multiplication in 3 ms at 80 MHz for a 192 bit elliptic curve. A TTA hardware implementation of RSA based on the Chinese Remainder Theorem is presented in [245], where the presented implementation is capable

Table 17: Summary of essential features and associated methods in cryptography

BASE OPERATION	TECHNIQUES AND METHODS
Fast exponentiation	Fast multiplication algorithms, windowing methods (See e.g. [234])
Fast modular arithmetic	Montgomery [227] and Barrett [235] modular reduction, and derivatives[236, 237, 238], Itoh-Tsujii inversion [239], repeated squarings [240], Extended Euclidean Algorithm, Chinese Remainder Theorem
Fast point scalar multiplication	Various coordinate systems, fast multiplication algorithms (See e.g. [241])
Fast substitution, rotation and permutation	Operation-specific Instruction Set Extensions [242], hardware implementation

of decrypting RSA at the rate of 106 kbps at 100 MHz. A TTA based hardware solution for RSA key expansion is presented in [246], where the presented hardware is capable of generating three 1024 bit RSA keys per second at 100 MHz, when implemented in VHDL and synthesized in a 0.18 μm process.

Our approach with the TACO platform had been primarily directed towards protocol processing, but both the advantages of the TTA based platform and the requirements of future networks support extending the platform to the cryptography domain. The development of communication systems towards IoT and sensor networks also put pressure on providing significant cryptography resources with flexible cipher suites and broad communication protocol support on hardware.

9.1.4.4 *The Pyramid Model of Security*

Security design can be generalized as a pyramid with different layers representing abstraction levels of security systems [247]. At the highest abstraction level is the security protocol architecture, which contains the protocols used for security purposes. The next abstraction level contains the actual algorithms such as RSA, DSA and SHA. The third abstraction level contains the "building blocks" of algorithms, derived from number theory. The fourth abstraction level contains cycle accurate platform independent representations of these algorithms, and the fifth level contains the physical implementation of the algorithms. If approached from a completely different point of view, the security pyramid model can also be seen as a service architecture. It incorporates a distinct customer-service relationship between the layers, where the upper layers issue directives and requirements for the layers below, whereas the latter provide service to the layers above, according to the received requirements [16]. The model is illustrated in Figure 31.

When we consider different protocols and algorithms for the SDSC platform and design methodology, we are mostly concerned with the first, second and third layer of the security pyramid. These layers constitute the manual part of the design flow, while the fourth layer is handled by the TACO design toolkit. The mapping of the pyramid model to the TACO design flow is examined later.

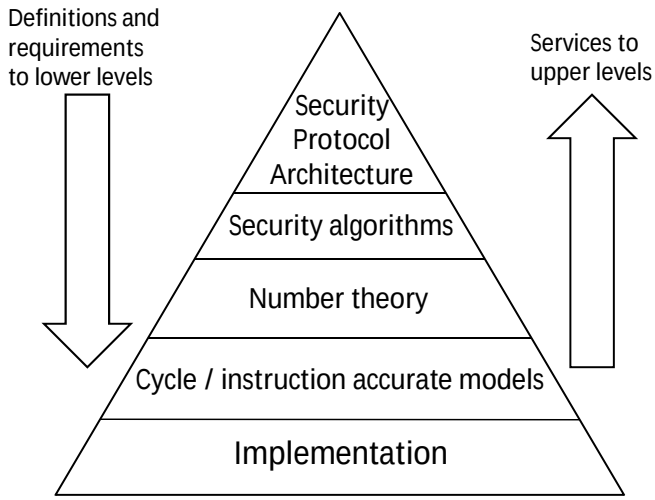

Figure 31: Pyramid model of security. Adapted from [247]. Used with permission of Springer.

9.1.4.5 *Adaptive Cryptography for Dynamic Security Dimensions*

Security requirements for future communication systems will continue to evolve, and when applications and use cases develop, security features and capabilities must evolve and adapt with them. One source of motivation for strengthening security features is the widespread surveillance of communication networks, affecting everyone on the Internet. This affects a significant part of the world's population and businesses.

The security landscape for IoT is challenging. When all devices can be – and many are – used to transfer sensitive information and all of it is subject to potential eavesdropping by malicious parties, securing these devices is a difficult but necessary task. IoT also brings an additional layer of complexity by making the number of communicating devices grow significantly in the future. Another strong driver are sensor networks, which require robust and adaptive security features as well.

Our chosen approach to this problem is to first recognize the underlying primitive operations for cryptography algorithms and then selecting a representative set of primitives to implement in hardware. In this way we can provide hardware acceleration to the computationally most challenging operations, and also

provide flexibility with regard to chosen algorithms and cipher suites, as they are implemented on a higher abstraction level.

Adaptive cryptography has been defined as a cryptographic library from where different implementations of algorithms can be selected to adapt to changes in requirements [248, 249]. The target application in the referred papers was IPSec, for which an FPGA implementation of a large cipher suite was made. We consider this approach to be inadequate in the face of current developments regarding network surveillance and the efforts to undermine cryptography standards by intelligence agencies. With this in mind, to respond to these challenges we will redefine adaptive cryptography as

> adapting to changing cryptography requirements on the fly, without replacing physical hardware, even when the changes warrant using a *completely new* algorithm.

We suggest an approach where the important underlying operations for cryptography are implemented in hardware, and the higher abstraction level is implemented in software. In this manner we will take a performance hit when compared to pure hardware implementations, but we see that this performance hit is acceptable with the gain of better flexibility and adaptability to unforeseen situations.

With adaptive cryptography techniques, we can create systems that are more versatile and less subject to catastrophic failure in case of a serious threat to the underlying security infrastructure and algorithms. We will discuss one particular promising implementation approach in Section 9.1.6.1, where we discuss the concept of security dimensions.

In Figure 32 we show the landscape of various encryption application areas and targets. Adaptive cryptography targets the application area where resources are constrained but flexibility is required. This application area may be of more interest in the future, based on the direction of societal development and how communication systems are used. The advent of big data has made nearly everything that is capable of collecting data a viable data source, and this data is communicated through the Internet. At the same time, important societal functions are also being transformed by the Internet, leading to more and more

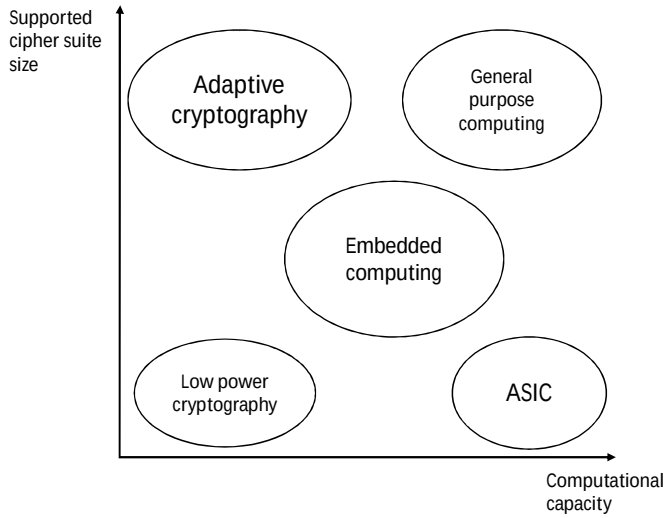

Figure 32: Landscape of different encryption application areas. Used with permission of Springer.

aspects of our lives being online. This in turn leads to developments that are not as positive, such as online crime, cyberterrorism and mass Internet surveillance — all are phenomena that have surfaced strongly in the 2010s. The TACO SDSC platform aims to provide a starting point for building systems capable of answering to these difficult challenges.

9.1.5 *Design methodology and flow*

In this section we define the design methodology and flow for the SDSC platform. The methodology is built on and extended from the SDR platform methodology proposed in [218]. We see the SDSC concept as a viable platform for next-generation communication devices. Parameterizability, programmability and run-time reconfigurability are all properties that make designing new hardware and software for emerging applications easier. This also provides better product lifetime, as product platforms are more extensible. Modularity allows reusing blocks both at design phase and at run-time, helps optimizing resource utilization and it is crucial for the design flow itself. This is especially important in energy- and space-constrained environments with high performance demands.

9.1.5.1 *Target Application and Domain-specific Requirement Analysis*

Several design stage issues must be taken into account. To identify potential points of application-domain specific optimization, we need to identify computationally heaviest portions of the code and find potential sets of functionality that are common throughout the application domain. These are priority targets for acceleration and thus are subject to being broken down into primitives at platform level. The analysis of the target application starts at the pseudo code level, where we start to identify recurring and parallel operations.

With looped operations we must verify that the operations are truly independent from each other. In this case it is possible to accelerate functions by adding additional FUs to the design and calculating each iteration of the operation on an individual FU in parallel. This is demonstrated in Figure 33, where a loop of 3 iterations on a single FU is replaced by three FUs in parallel. For

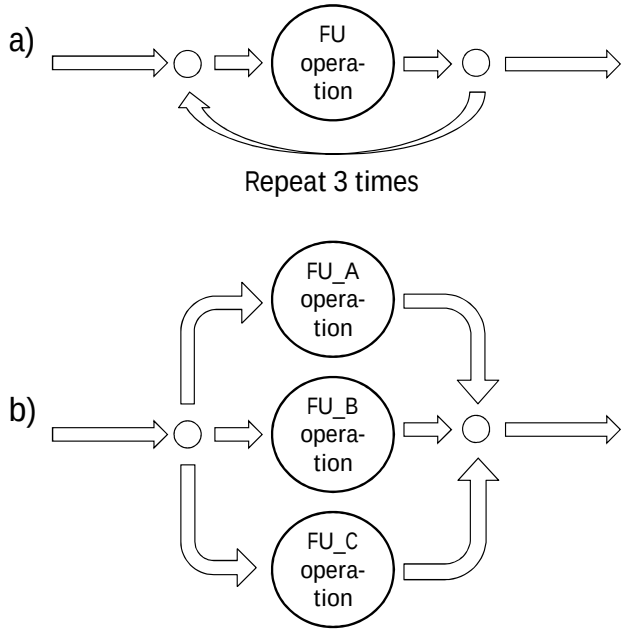

Figure 33: Illustration of parallelism with multiple FUs. a) Repeated operation with single FU. b) Parallel implementation using several FUs. Used with permission of Springer.

some of our intended applications in the SDSC domain, such as

cryptography, repeated operations are often dependent of each other. As an example, it would be intuitive to have a design with several FUs which implement one round of a symmetric cipher such as AES, thus increasing encryption and decryption speed. But because the input to a round is dependent on the output of the previous round in most cipher modes, this approach cannot be taken. One exception to this is the Electronic Code book (ECB) mode, where the rounds are independent, but it is seldom used in any real-world applications because it lacks semantic security [250]. This in combination with the findings in [221] suggest that while parallelism does exist, it is not common enough in cryptography applications. Acceleration of primitives and different arithmetic techniques thus are often better options for optimizing cryptography performance.

Reuse of previously designed hardware instances, whether whole or partial, is a key design goal in this design methodology. This approach provides significant advantage in design time. Using existing design instances that have similar functionality as templates is a practical approach that provides the designer with a baseline design with low initial cost, sparing a lot of time and effort. After template initialization, all the necessary functional units are then added to the design. This is based on a careful analysis of the target application and what are the essential operations that need to be implemented on the hardware level.

9.1.5.2 *Design flow*

The proposed SDSC design methodology is described in Figure 34. The first step is analysis of the target application, and the result of this analysis is a list of functionality requirements for that application. This makes the methodology a top-town approach. When these key functions has been identified, the necessary algorithms required for implementing them must be chosen. The application domain of the target application will heavily influence this process. It also provides a chance to group similar functions into hardware units and augmenting already existing template components, thus promoting reuse of existing components. If no such functionality is available in the platform's library of ex-

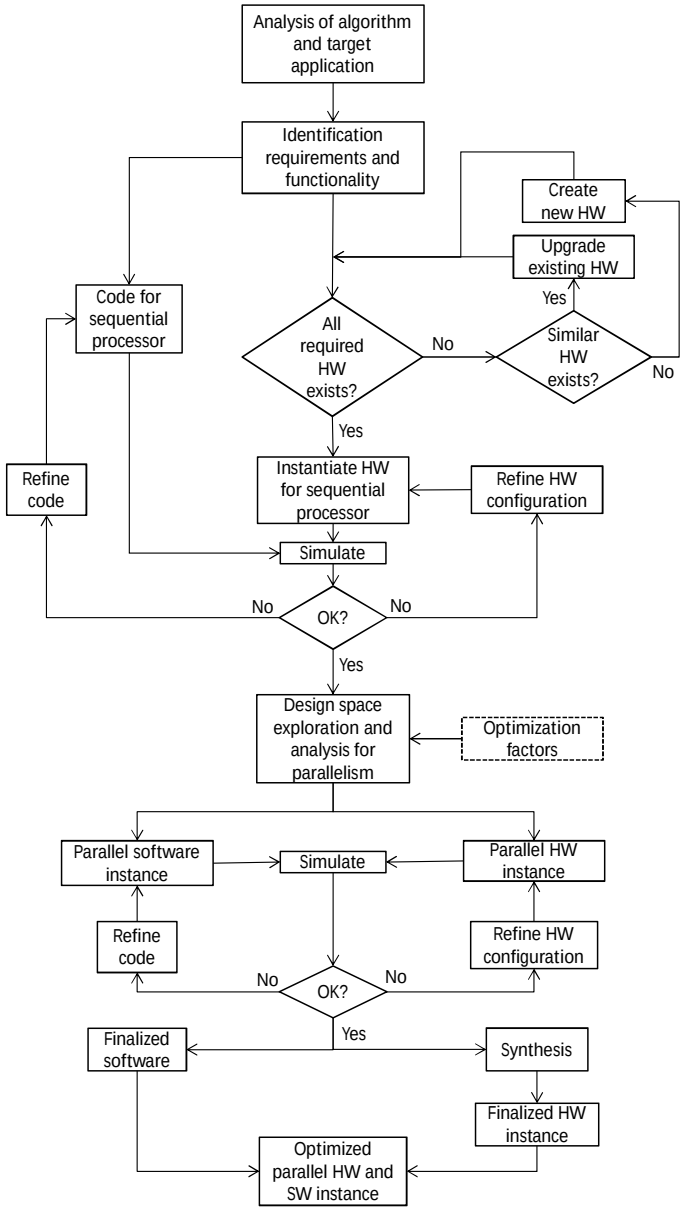

Figure 34: Design flow for hardware-software codesign in TACO SDSC. Used with permission of Springer.

isting hardware templates, a new hardware template must be designed for this purpose.

Next, a hardware platform instance for sequential processing is specified. In this instance, only one interconnection bus and

one FU of each required FU type are included. This gives us an initial setup that satisfies functional processing requirements of the target application, but processing speed requirements are not typically met at this stage. This process is shown in Figure 35, where the system level model is mapped to VHDL library components. This sequential instance, along with associated sequential program code is then simulated, testing for correct operation. Once the functionality is verified, the next stage in the method-

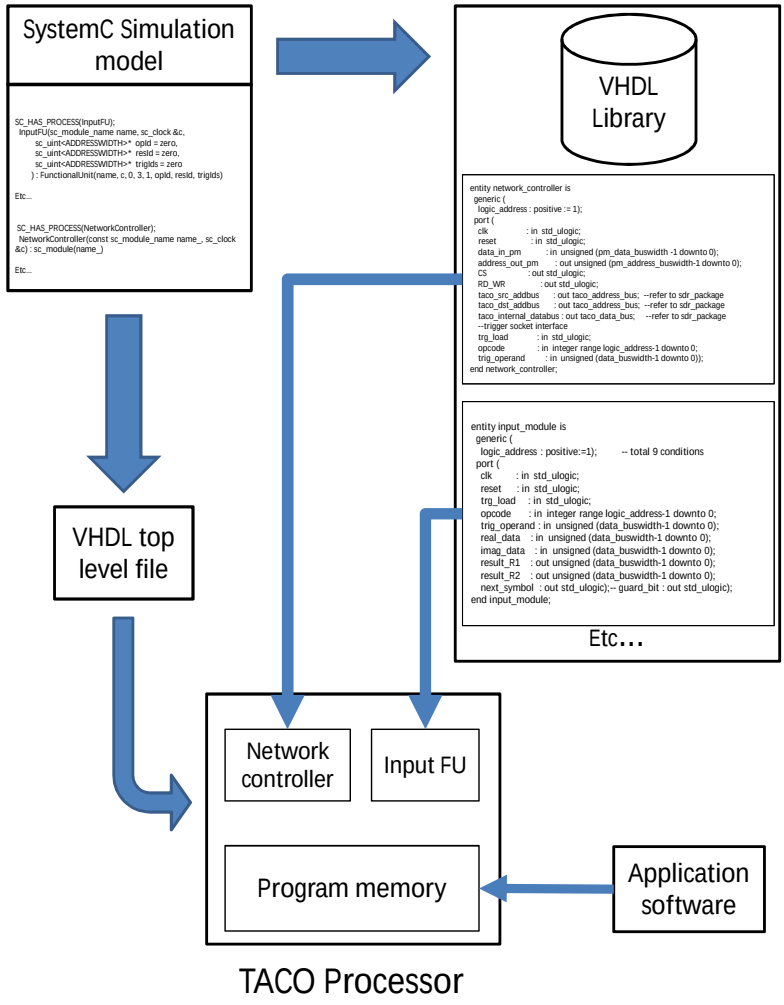

Figure 35: Process flow for mapping of system-level designs to VHDL hardware library components and their relation to the final platform instance, complemented by the co-designed application software. Used with permission of Springer.

ology is designer-driven design space exploration. Here, parallel processing capacity is added to the model to meet the processing speed requirements of the target application. The designer typically specifies 2-3 additional hardware instances, varying the number of data transport buses (data transportation parallelism) and the number of FUs of the same type in the hardware instance (FU level parallelism). Both methods increase processing performance. The target application and its computational capacity requirement greatly affects the set of feasible hardware configuration instances. This process is not automatic. Optimization factors derived from power consumption and chip size requirements, as well as current processing performance requirements and even future extensibility – are all manually defined by the designer. When we are satisfied with the obtained model we test the optimized version for functional correctness and processing speed, and estimate its hardware characteristics (power consumption and chip size) based on the data gathered from the VHDL component library. When all tests are approved, the hardware is synthesized and the software is finalized for that hardware instance.

9.1.5.3 *Cryptography reconfiguration approaches*

In the case of a fixed hardware implementation, adversaries immediately know what algorithms are used and what are available for use in different applications. With essential operations implemented in hardware, and encryption on the algorithm level implemented in software with hardware acceleration, the software designer has a lot of control on the actual security features of the system. Reacting to vulnerabilities is similarly significantly more flexible, as we can first approach the problem as a software problem, even in cases when hardware level changes are normally warranted.

An effective way for an attacker to approach a hardware cryptography implementation is to try to map its behavior and to find potential side channel attacks [251], which allow the attacker to extract information about the encrypted data from the operation of the physical device itself. If a hardware implementation of an algorithm is found to be vulnerable to a side channel attack, mitigating this threat can be next to impossible. With

adaptive cryptography, responding to side channel attacks and timing attacks changes from impossible to plausible due to the possibility of reprogramming the software and still using the same hardware acceleration blocks as before. Naturally, if the hardware blocks have been implemented in a fundamentally insecure manner and the side-channel attack is not dependent on leveraging the way the algorithm has been implemented, this approach cannot be taken.

If we are using a reconfigurable hardware platform such as FPGA, it expands the options available for cryptography reconfiguration. In addition to redesigning the software, we can also change hardware configuration without actually replacing the hardware. In such use cases where the target application changes with a non-negligible probability, this can be a valid approach. Also the cost of reconfigurable hardware can be prohibitive in many applications, where the unit cost of a computation platform must be pushed as low as possible.

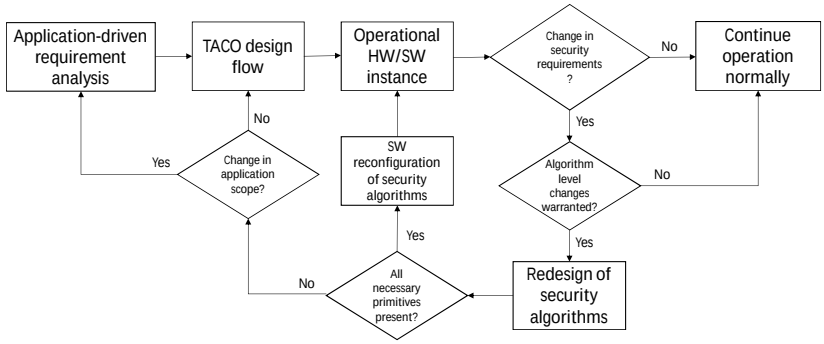

Figure 36: Design flow for reconfiguring security functions. Used with permission of Springer.

The design flow for reconfiguring cryptography functions and the relation of this process to the TACO design flow is presented in Figure 36. In the case of a change in the security environment that warrants algorithm level changes to the device, we first enter the software segment of the flow, where the security algorithms are re-evaluated and redesigned as necessary. After this we examine the suitability of the current hardware platform. On one hand, if all necessary primitives are present, no hardware reconfiguration is necessary and we can resume operation with new, redesigned software. On the other hand, if hardware imple-

mentations of necessary primitives are not present, we enter the hardware segment of the flow. Here we will first assess whether the initial application requirements have changed and need to be re-evaluated, or if we can just add the necessary hardware instances and go through design space exploration on the existing design. Regardless, we will re-iterate the TACO design flow to get a new design suitable for the changed situation. If we already have hardware reconfiguration capability, reacting to changes in application scope and required cryptography primitives is significantly more streamlined, as we can adapt the existing hardware to a new configuration that we get as the result of a new iteration of the TACO design flow without replacing hardware.

9.1.5.4 *Product customization on the market*

A potential benefit from using the TACO SDSC design flow is the opportunity for swift design of hardware implementations that can be mass produced at large volume and later customized to the specific configuration required by the target application on the software layer. In this approach a single hardware platform capable of supporting several functions and cryptography cipher sets is designed, and the actual product differentiation is done with software. Another significant benefit is that when primitive operations are implemented on hardware and their relationships are defined on the software layer, we can fix possible bugs in implementation by rewriting the software. This makes it possible to fix a bug in production with a software update that would require product recall if implemented fully in hardware.

Reconfigurable hardware naturally makes this process less vital, but the target applications for this approach would probably be those where it is not cost effective or feasible to use reconfigurable hardware. Regardless of the type of the hardware platform, this level of flexibility in customization after manufacture allows for larger manufacturing volumes and cost savings, while at the same time shortening the time-to-market for a product.

9.1.6 *Application and implementation scenarios*

In this section we will briefly discuss some interesting and potentially noteworthy applications and implementation scenarios.

9.1.6.1 *Dynamic security dimensions and adaptive cryptography*

In [18] we have explored the concept of dynamic security dimensions and their effects on sensor network security. The concept is based on the various target applications that sensor networks have, and that security requirements can and do change on the fly. We aim to identify all relevant security dimensions, and model them in a manner which facilitates the construction of an adaptive holistic security system for low-power and resource-constrained devices. Sensor networks and IoT devices can benefit from this approach.

To summarize, we first model the security landscape of sensor networks as individual "dimensions" that represent distinct, measurable characteristics. The key dimensions that we have identified are energy (d_e), processing capacity (d_{proc}), memory capacity (d_{me}), data coherence (d_{co}), data lifetime (d_{li}), location security level (d_{loc}), and mobility (d_{mo}). These dimensions, or dimension vectors, are modified by the corresponding scalar weight factors (w). This enables modeling potential he changes in the environment. Each dimension vector can be either a scalar based on qualitative assessment of the security landscape, or can have a distinct unit, if desired. For example, the location security can be assigned as follows. If the system processes top secret data, $d_{loc} = 100$; if the data is confidential, $d_{loc} = 50$; and if the data is public, $d_{loc} = 10$. One approach is to normalize all dimensions to a common scale, for example from 0 to 100.

A model with seven distinct dimensions can be complex to visualize and assess, so reducing the complexity of the model, if possible, is justified. By analyzing the nature of the dimensions and grouping those that are related together, the result is reduced complexity in modeling the security environment. Energy is the only dimension left uncompounded, given its importance in resource-constrained systems. Other dimensions are clearly interconnected; location privacy and node mobility can be grouped together, as they are both spatial functions. Memory and processing capacity are both hardware related dimensions, and data coherence and lifetime are heavily dependent on the

target application. We define the final set of compound dimensions as the 4-tuple

$$(\alpha, \beta, \gamma, \delta) = (d_e w_e, \frac{d_{me} w_{me} + d_{proc} w_{proc}}{w_{me} + w_{proc}}, \frac{d_{mo} w_{mo} + d_{loc} w_{loc}}{w_{mo} + w_{loc}}, \frac{d_{co} w_{co} + d_{li} w_{li}}{w_{co} + w_{li}}). \quad (9)$$

These dimensions can be used to model the security environment by assigning a mapping between different real world scenarios and the 4-dimensional security landscape. For example, a certain region of the security space where energy consumption is not an issue, and the coherence and lifetime of data are high, using the strongest available encryption algorithm supported by the current configuration is justifiable. Another region could have high limits for energy consumption and low data lifetime and coherence, so the security environment would allow data to be transmitted even without any confidentiality measures in place. For some situations even integrity preserving features can be turned off to conserve energy.

Changes in dimensions can force changes in operating parameters, and this may require changing operational security features in the system. These changes can happen quickly or gradually. Even the slow kind of change can be dramatic, if for example a certain set of cryptographic ciphers is compromised. While unlikely, this not impossible, as cryptanalysis on existing systems is always ongoing and intelligence organizations and other equivalent governmental organizations have capabilities that can force changes in this field [252]. The ability to adapt even older hardware to a new situation can be a preferable approach, compared to having to replace all hardware with new designs, having to wait for them to be designed and verified and manufactured, and so on.

The main motivation of this approach is to design a dynamically adaptive sensor network security framework, capable of assessing available resources and changing application requirements in order to provide optimal protection with minimal overhead and cost. The framework observes the network continuously using an intrusion detection/prevention system and also uses location related information and feedback from a trust management module responsible for mapping different trust scenar-

ios and conditions on the sensor network and its surroundings. With this information, the end system based on this framework takes action to achieve optimal security-energy point. This translates into maximum security under a given resource and context. The security framework will be developed based on a platform vision to enable cost-effective and application-independent realization. The scalability feature of the framework is also an important aspect since the number of nodes in low-power wireless networks varies from a few to many nodes based on the application requirement.

We see that incorporating adaptive cryptography as a concept is vital to the evolution of such a framework, and the TACO SDSC platform will give us a robust and efficient HW-SW platform for development and eventual implementation.

9.1.6.2 *Building security in the design*

The current version of the TACO SDSC methodology does not yet take into account some aspects of secure processor design as outlined in [253] and [254], such as secure handling of cryptographic keys. These security issues have been addressed in designs such as SAFES [255], and similar care should be taken when designing cryptography functionality for TACO SDSC processors. This is a clear improvement target for the design process in order to incorporate security and safety into the core of the methodology.

In a broader scope, the concept of building security and trust into systems that we use every day is a vital aspect of future communication systems. The TACO SDSC platform and methodology make designing communication systems that can adapt to changing environments and requirements significantly easier. For a single hardware instance, we can support several different software implementations of the same functionality. By doing the changes only on the software level, we provide an additional layer of resistance to attacks by malicious parties. If a hardware instance is deemed to be vulnerable to an attack, by reconfiguring the system on the software level we can extend our mitigation capabilities substantially.

9.2 CONCLUSION

When we consider the requirements of future communication systems, the need for flexible encryption in all communication layers is evident. With the advent of IoT, ubiquitous computing, and a data intensive information society, devices and sensors with strict requirements for both efficient communications and robust security will be in demand. The SDR concept can be leveraged and enhanced to answer these requirements.

A design methodology and process for a Software-Defined Secure Communications framework has been presented in this chapter. As a platform, SDSC can provide secure, flexible and lightweight communication to sensor networks and IoT devices without significantly compromising performance, and the presented framework provides the methods and processes for exploring this further. In simulated tests, a TACO processor with separate FUs for modular arithmetic implements reduced key length RSA with a 25% increase in expended clock cycles and a 31% increase in execution time when compared to a dedicated hardware solution implemented on the same platform [256]. The performance of this system can be potentially increased with further improvement and optimization, a promising avenue for future research in the field.

Additional benefits of this design approach include potential cost savings in larger production volumes and shorter time-to-market for products when we can design devices, that are highly customizable with software and implement only key primitive operations in hardware. We see that our design model coupled with the concept of adaptive cryptography could be leveraged towards developing flexible and yet secure application-specific devices for future communication networks.

10 CONCLUSION

Had to be me. Someone else might have gotten it wrong.

– Mordin Solus
Mass Effect 3 (2012)

Information security is inherently made of compromise. Truly, security in itself is a complicated trade-off between security and various other factors [257, p. 11]. A goal of absolute security is unattainable, so the designers of systems must take into account certain trade-offs in security versus usability. An educated risk evaluation based on the target system and environment is critical in creating trustworthy systems and services. This is true for eGovernment systems and other critical infrastructure, and this work is a part of the process of creating a holistic picture of the security and privacy issues in the networked information society.

10.1 MAPPING RESULTS TO RESEARCH OBJECTIVES

In this thesis we have examined information security in the networked information society. We identified a set of research problems and have devised solutions or shown potential directions from where a solution can be found. To reiterate, the Research Objectives (RO) defined in the beginning of this thesis were:

RO1: Building security in to the networked information society; and

RO2: Providing parameters for further development of a safe and secure networked information society.

To summarize the contribution of this thesis, the following ten Key Results (KR) of this thesis, and the RO they map to, are as follows.

KR1: Mass Internet surveillance is harmful to the networked information society. The technological methods used to conduct systematic information gathering and surveillance endanger the security and safety of all Internet users. (RO2)

KR2: Forced trust, detrimental to the networked information society, in the infrastructure of the networked information society and associated Critical Governmental Information System (CGIS) exists. (RO2)

KR3: Biometric passport systems are problematic in the context of the networked information society due to the threat to privacy and security of citizens through biometric information. (RO2)

KR4: Current frameworks for modeling Internet voting systems lack adversary models that take sufficiently and explicitly into account the capabilities that can be observed in implementation of modern mass surveillance. (RO2)

KR5: A framework for necessary security knowledge for engineering students is defined and tested in university level security education. (RO1)

KR6: Network security education is improved through virtual education environments and university-industry cooperation. (RO1)

KR7: Datenherrschaft can be used as framework for managing person-derived data used in legitimate surveillance activities. (RO1;RO2)

KR8: Based on analysis of Finnish web user passwords, natural language does not provide extra security in general when used in passwords, as users tend to only use base forms of words. (RO1)

KR9: Devices in the era of Internet of Things (IoT) have reduced computational capacity. Changes in cryptographic algorithms can happen due to change in target application, compromise or obsolescence of the algorithm, or introduction of new methods. Computationally restricted devices with

potential hardware implementations of cryptography algorithms cannot adapt to such changes. Adaptive cryptography implementation with well-chosen cryptography primitives is a potential tool in achieving end-to-end security for everyone in the networked information society. (RO1;RO2)

KR10: The TACO toolkit and design methodology, together with the improved design flow presented in this thesis, provides a robust method for fast design of embedded devices that are capable of implementing adaptive cryptography requirements in communication. This is realized through efficient reuse of existing components, and the capability to add new functionality to a device quickly, resulting in a short design cycle. (RO1)

This thesis work has provided an opportunity to help the ongoing effort to form a coherent view of a vast research area – designing, implementing and educating information security in the networked information society. The contribution of this thesis is an additional voice in the discussion on how we should continue on building a secure infrastructure for the future world. This thesis not only presents engineering solutions to the problems of the future, but also gives potential directions for later research and poses new questions that are complex and difficult — perhaps even impossible — to answer fully. This is not a bad thing: it is what good questions should do — pose even more challenging questions so that others can endeavor to answer them in the future. As to whether the questions that have been posed here in this thesis are the *right* questions, that only time will tell.

BIBLIOGRAPHY

[1] C. West Churchman. Free for all. *Management Science*, 14 (4):B–141–B–146, 1967. doi: 10.1287/mnsc.14.4.B141. URL http://dx.doi.org/10.1287/mnsc.14.4.B141.

[2] John Scott and Gordon Marshall. *A dictionary of sociology*. Oxford University Press, USA, 2009.

[3] Nick Moore. The information society. *World information report*, 98:271–284, 1997.

[4] Frank Webster. *Theories of the information society, Third edition*. Routledge, New York, United States of America, 2006.

[5] Monica Anderson. *Technology Device Ownership: 2015*. Pew Research Center, October 2015. Available online at http://www.pewinternet.org/2015/10/29/the-demographics-of-device-ownership/ Accessed 10.02.2016, 2015.

[6] Olli I. Heimo, Jani S. S. Koskinen, and Kai K. Kimppa. Responsibility in acquiring critical governmental information systems: Whose fault is failure? In *ETHICOMP 2013 Proceedings*, pages 213–217, 2013.

[7] Bruce Schneier. *Data and Goliath*. W. W. Norton & Co., New York, United States of America, 2015.

[8] Antti Hakkala, Olli I. Heimo, Kai K. Kimppa, and Sami Hyrynsalmi. Security, Privacy'); DROP TABLE users; – and Forced Trust in the Information Age? In *In Ethicomp 2017 conference, Turin, Italy, 5.-8.6.2017*, 2017.

[9] Olli I. Heimo, Antti Hakkala, and Kai K. Kimppa. The problems with security and privacy in egovernment — case: Biometric passports in finland. In *Ethicomp 2011 Conference. Sheffield-Hallam University, UK, September 14.–16.*, 2011.

[10] Olli I. Heimo, Antti Hakkala, and Kai K. Kimppa. How to abuse biometric passport systems. *Journal of Information, Communication and Ethics in Society*, 10(2):68–81, 2012.

[11] Jani S. S. Koskinen. *Datenherrschaft – An Ethically Justified Solution to the Problem of Ownership of Patient Information*. PhD thesis, Turku School of Economics; University of Turku, March 2016.

[12] Antti Hakkala and Seppo Virtanen. University-industry collaboration in network security education for engineering students. In *Proceedings of the International Conference on Engineering Education ICEE 2012*, 2012.

[13] Antti Hakkala and Seppo Virtanen. Virtualization of laboratory education in network security engineering. In *Proceedings of the International Conference on Engineering Education ICEE 2015*, 2015.

[14] Antti Hakkala and Jouni Isoaho. Defining and measuring key expertise areas in information security for higher education students. In *Proceedings of the International Conference on Engineering Education ICEE 2015*, 2015.

[15] Seppo Virtanen. *A framework for rapid design and evaluation of protocol processors*. PhD thesis, University of Turku, 2004.

[16] Antti Hakkala and Seppo Virtanen. Accelerating cryptographic protocols: A review of theory and technologies. In *Proceedings of CTRQ 2011 : The Fourth International Conference on Communication Theory, Reliability, and Quality of Service*, 2011.

[17] Antti Hakkala, Jouni Isoaho, and Seppo Virtanen. Towards adaptive cryptography and security with software defined platforms. In Hussain Waqar, Jari Nurmi, Jouni Isoaho, and Fabio Garzia, editors, *Computing Platforms for Software-Defined Radio*. Springer, 2017. ISBN 978-3-319-49678-8.

[18] Ethiopia Nigussie, Antti Hakkala, Seppo Virtanen, and Jouni Isoaho. Energy-aware adaptive security management for wireless sensor networks. In *IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks*. IEEE, July 2014.

[19] Roger Clarke. Information technology and dataveillance. *Commun. ACM*, 31(5):498–512, May 1988. ISSN 0001-0782. doi: 10.1145/42411.42413. URL http://doi.acm.org/10.1145/42411.42413.

[20] Columbia University School of Law. Surveillance, dataveillance and personal freedoms - use and abuse of information technology - a symposium, 1973.

[21] Statistics Finland. Official Statistics of Finland (OSF): Use of information and communications technology by individuals [e-publication]. Available online at http://www.stat.fi/til/sutivi/index_en.html, 2015. Accessed 19.1.2016.

[22] Oscar H. Gandy. The surveillance society: information technology and bureaucratic social control. *Journal of Communication*, 39(3):61–76, 1989.

[23] Colin J. Bennett. The public surveillance of personal data: a cross-national analysis. *Computers, surveillance, and privacy*, pages 237–259, 1996.

[24] Heng Xu, Tamara Dinev, H. Jeff Smith, and Paul Hart. Examining the formation of individual's privacy concerns: toward an integrative view. *ICIS 2008 Proceedings*, page 6, 2008.

[25] John O. Koehler. *STASI: The untold story of the East German secret police*. Basic Books, 2008.

[26] Gary Bruce. Access to secret police files, justice, and vetting in east germany since 1989. *German Politics & Society*, 26(1):82–111, 2008.

[27] Steven Pfaff. The Limits of Coercive Surveillance – Social and Penal Control in the German Democratic Republic. *Punishment & Society*, 3(3):381–407, 2001.

[28] Brandon C. Welsh, David P. Farrington, and Sema A. Taheri. Effectiveness and social costs of public area surveillance for crime prevention. *Annual Review of Law and Social Science*, 11:111–130, 2015.

[29] Leon Hempel and Eric Töpfer. On the threshold to Urban Panopticon: Analysing the Employment of CCTV in European Cities and Assessing its Social and Political Impacts. CCTV in Europe - Final report. Technical report, Centre for Technology and Society, Technical University Berlin, 2004.

[30] Clive Norris, Mike McCahill, and David Wood. Editorial. the growth of cctv: a global perspective on the international diffusion of video surveillance in publicly accessible space. *Surveillance & Society*, 2(2/3):111, 2004.

[31] Greg Walton. *China's golden shield: corporations and the development of surveillance technology in the People's Republic of China*. Rights & Democracy, 2001.

[32] David Lyon. *Surveillance studies: An overview*. Polity, 2007.

[33] David Lyon, Kirstie Ball, and Kevin D. Haggerty. *Routledge handbook of surveillance studies*. Routledge, 2012.

[34] Jeremy Bentham and Miran Božovič (ed.). *The Panopticon and Other Prison Writings (Wo Es War)*. Verso Books, 1995.

[35] Michel Foucault. *Disclipine and Punish: The birth of the prison*. Random House LLC, 1977.

[36] Lu Xia, Chia-Chih Chen, and Jake K. Aggarwal. Human detection using depth information by kinect. In *Computer Vision and Pattern Recognition Workshops (CVPRW), 2011 IEEE Computer Society Conference on*, pages 15–22. IEEE, 2011.

[37] Xu Xu, Raymond W. McGorry, Li-Shan Chou, Jia-hua Lin, and Chien-chi Chang. Accuracy of the Microsoft Kinect for measuring gait parameters during treadmill walking. *Gait & posture*, 42(2):145–151, 2015.

[38] Brook Galna, Gillian Barry, Dan Jackson, Dadirayi Mhiripiri, Patrick Olivier, and Lynn Rochester. Accuracy of the Microsoft Kinect sensor for measuring movement in people with Parkinson's disease. *Gait & Posture*, 39(4):1062 – 1068, 2014. doi: http://dx.doi.org/10.1016/j.gaitpost.2014.01.008. URL http://www.sciencedirect.com/science/article/pii/S0966636214000241.

[39] Zygmunt Bauman. *Globalization: The human consequences*. Columbia University Press, 1998.

[40] Roy Boyne. Post-panopticism. *Economy and Society*, 29(2): 285–307, 2000.

[41] Kevin D. Haggerty. Tear down the walls: On demolishing the panopticon. In David Lyon, editor, *Theorizing Surveillance. The panopticon and beyond*, chapter 2, pages 23–45. Routledge, 2007.

[42] Daniel Kahneman and Amos Tversky. Prospect theory: An analysis of decision under risk. *Econometrica: Journal of the Econometric Society*, pages 263–291, 1979.

[43] Frederick Schauer. Fear, risk and the first amendment: Unraveling the chilling effect. *BUL Rev.*, 58:685, 1978.

[44] Elizabeth Stoycheff. Under Surveillance: Examining Facebook's Spiral of Silence Effects in the Wake of NSA Internet Monitoring. *Journalism & Mass Communication Quarterly*, pages 1–16, March 2016. doi: 10.1177/1077699016630255.

[45] Elisabeth Noelle-Neumann. The spiral of silence a theory of public opinion. *Journal of communication*, 24(2):43–51, 1974.

[46] Ray Bull and Elizabeth Gibson-Robinson. The influences of eye-gaze, style of dress, and locality on the amounts of money donated to a charity. *Human Relations*, 34(10):895–905, 1981.

[47] Robert Kurzban. The social psychophysics of cooperation: Nonverbal communication in a public goods game. *Journal of Nonverbal Behavior*, 25(4):241–259, 2001.

[48] Kevin J. Haley and Daniel M.T. Fessler. Nobody's watching?: Subtle cues affect generosity in an anonymous economic game. *Evolution and Human behavior*, 26(3):245–256, 2005.

[49] Terence C. Burnham and Brian Hare. Engineering human cooperation. *Human Nature*, 18(2):88–108, 2007.

[50] Max Ernest-Jones, Daniel Nettle, and Melissa Bateson. Effects of eye images on everyday cooperative behavior: a field experiment. *Evolution and Human Behavior*, 32(3):172–178, 2011.

[51] Damien Francey and Ralph Bergmüller. Images of eyes enhance investments in a real-life public good. *PLoS One*, 7(5):e37397, 2012.

[52] Daniel Nettle, Zoe Harper, Adam Kidson, Rosie Stone, Ian S. Penton-Voak, and Melissa Bateson. The watching eyes effect in the dictator game: it's not how much you give, it's being seen to give something. *Evolution and Human Behavior*, 34(1):35–40, 2013.

[53] Adam Sparks and Pat Barclay. Eye images increase generosity, but not for long: The limited effect of a false cue. *Evolution and Human Behavior*, 34(5):317–322, 2013.

[54] Mathias Ekström. Do watching eyes affect charitable giving? evidence from a field experiment. *Experimental Economics*, 15(3):530–546, 2012.

[55] Open Science Collaboration et al. Estimating the reproducibility of psychological science. *Science*, 349(6251): aac4716, 2015.

[56] Daniel J. Solove. "i've got nothing to hide" and other misunderstandings of privacy. *San Diego law review*, 44:745–772, 2007.

[57] David Lyon. An electronic panopticon? a sociological critique of surveillance theory. *The Sociological Review*, 41(4): 653–678, 1993.

[58] Shoshana Zuboff. *In the age of the smart machine: The future of work and power*. Basic Books, 1988.

[59] Peter Bain and Phil Taylor. Entrapped by the 'electronic panopticon'? worker resistance in the call centre. *New technology, work and employment*, 15(1):2–18, 2000.

[60] Jonathan Mayer and Patrick Mutchler. MetaPhone: the sensitivity of telephone metadata. Web Policy, available

online at http://webpolicy.org/2014/03/12/metaphone-the-sensitivity-of-telephone-metadata/ Accessed 8.9.2016, 2014.

[61] Giuseppe Zappalà. Killing by metadata: Europe and the surveillance–targeted killing nexus. *Global Affairs*, 2015. doi: 10.1080/23340460.2015.1080035.

[62] Kevin D. Haggerty and Richard V. Ericson. The surveillant assemblage. *The British journal of sociology*, 51(4):605–622, 2000.

[63] Gilles Deleuze and Felix Guattari. A thousand plateaus (b. massumi, trans.). *Minneapolis: University of Minnesota Press.(Original work published 1980)*, 1987.

[64] Darren Ellis, Ian Tucker, and David Harper. The affective atmospheres of surveillance. *Theory & Psychology*, page 0959354313496604, 2013.

[65] Kenneth C. Laudon. Ethical concepts and information technology. *Communications of the ACM*, 38(12):33–39, 1995.

[66] James Cox. *Canada and the Five Eyes Intelligence Community*. Canadian Defence & Foreign Affairs Institute, 2012.

[67] Andrew Clement. NSA Surveillance: Exploring the Geographies of Internet Interception. In *iConference 2014 Proceedings*, pages 412–425, 2014.

[68] David Lyon. Surveillance, snowden, and big data: capacities, consequences, critique. *Big Data & Society*, 1(2): 2053951714541861, 2014.

[69] David Lyon. *Surveillance after Snowden*. John Wiley & Sons, 2015.

[70] Carna Botnet. Internet census 2012: Port scanning /0 using insecure embedded devices. *Published online, available at http://internetcensus2012.github.io/InternetCensus2012/paper.html*, 2013.

[71] Nicolas Falliere, Liam O. Murchu, and Eric Chien. W32.stuxnet dossier. *White paper, Symantec Corp., Security Response*, 5(6), 2011.

[72] Data Encryption Standard. Federal information processing standards publication 46. *National Bureau of Standards, US Department of Commerce*, 1977.

[73] Bruce Schneier. *Applied Cryptography: Protocols, Algorithms and Source code in C, 2nd edition*. John Wiley and Sons, 1996.

[74] Electronic Frontier Foundation. *Cracking DES: Secrets of Encryption Research, Wiretap Politics and Chip Design*. O'Reilly & Associates, Inc., Sebastopol, CA, USA, 1998. ISBN 1565925203.

[75] National Institute of Standards and Technology (NIST). Advanced Encryption Standard. *Federal Information Processing Standards Publication 197*, November 26 2001.

[76] A. Michael Froomkin. The metaphor is the key: cryptography, the clipper chip, and the constitution. *University of Pennsylvania Law Review*, 143(3):709–897, 1995.

[77] Daniel J. Bernstein, Tanja Lange, and Ruben Niederhagen. Dual EC: A standardized back door. In *The New Codebreakers*, pages 256–281. Springer, 2016.

[78] Andrew Roos. A Class of Weak Keys in the RC4 Stream Cipher. Available online at `http://impic.org/papers/WeakKeys-report.pdf` Accessed 7.2.2017., 1995.

[79] Jovan Dj. Golić. Linear statistical weakness of alleged RC4 keystream generator. In *International Conference on the Theory and Applications of Cryptographic Techniques*, pages 226–238. Springer, 1997.

[80] Scott Fluhrer, Itsik Mantin, and Adi Shamir. Weaknesses in the key scheduling algorithm of RC4. In *International Workshop on Selected Areas in Cryptography*, pages 1–24. Springer, 2001.

[81] Itsik Mantin and Adi Shamir. A practical attack on broadcast RC4. In *International Workshop on Fast Software Encryption*, pages 152–164. Springer, 2001.

[82] Andreas Klein. Attacks on the RC4 stream cipher. *Designs, Codes and Cryptography*, 48(3):269–286, 2008.

[83] Pouyan Sepehrdad, Serge Vaudenay, and Martin Vuagnoux. Discovery and exploitation of new biases in RC4. In *International Workshop on Selected Areas in Cryptography*, pages 74–91. Springer, 2010.

[84] Takanori Isobe, Toshihiro Ohigashi, Yuhei Watanabe, and Masakatu Morii. Full plaintext recovery attack on broadcast RC4. In *International Workshop on Fast Software Encryption*, pages 179–202. Springer, 2013.

[85] Ilya Mironov. (Not so) random shuffles of RC4. In *Annual International Cryptology Conference*, pages 304–319. Springer, 2002.

[86] Mathy Vanhoef and Frank Piessens. All Your Biases Belong to Us: Breaking RC4 in WPA-TKIP and TLS. In *USENIX Security*, volume 2015, 2015.

[87] Andrea Bittau, Mark Handley, and Joshua Lackey. The final nail in WEP's coffin. In *Security and Privacy, 2006 IEEE Symposium on*, pages 15–pp. IEEE, 2006.

[88] Nadhem AlFardan, Daniel J Bernstein, Kenneth G Paterson, Bertram Poettering, and Jacob CN Schuldt. On the security of RC4 in TLS. In *Presented as part of the 22nd USENIX Security Symposium (USENIX Security 13)*, pages 305–320, 2013.

[89] Bodo Möller, Thai Duong, and Krzysztof Kotowicz. This poodle bites: exploiting the ssl 3.0 fallback, 2014.

[90] David Adrian, Karthikeyan Bhargavan, Zakir Durumeric, Pierrick Gaudry, Matthew Green, J. Alex Halderman, Nadia Heninger, Drew Springall, Emmanuel Thomé, Luke Valenta, et al. Imperfect forward secrecy: How Diffie-Hellman fails in practice. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*, pages 5–17. ACM, 2015.

[91] Whitfield Diffie and Martin Hellman. New directions in cryptography. *IEEE transactions on Information Theory*, 22 (6):644–654, 1976.

[92] J. Ronald Prins. Diginotar certificate authority breach óperation black tulip. Available online at https://www.ncsc.nl/binaries/content/documents/ncsc-en/current-topics/news/interim-report-on-diginotar-digital-breach-published/1/Interim%2Breport%2Bon%2BDigiNotar%2Bdigital%2BCertificate%2Bbreach.pdf Accessed 7.2.2017., September 5th 2011.

[93] James B. Rule. *Privacy in peril: How we are sacrificing a fundamental right in exchange for security and convenience*. Oxford University Press, 2007.

[94] Naomi Gilens. The nsa has not been here: Warrant canaries as tools for transparency in the wake of the snowden disclosures. *Harv. JL & Tech.*, 28:525, 2014.

[95] Bruce Schneier. *Liars and outliers: enabling the trust that society needs to thrive*. Wiley, 2012.

[96] Stephen Marsh and Mark R. Dibben. The role of trust in information science and technology. *Annual Review of Information Science and Technology*, 37(1):465–498, 2003.

[97] William Dutton, Gerardo A. Guerra, Daniel J. Zizzo, and Malcolm Peltu. The cyber trust tension in e-government: Balancing identity, privacy, security. *Information Polity*, 10 (1):13–23, 2005.

[98] Daniel Metlay. Institutional trust and confidence: A journey into a conceptual quagmire. *Social trust and the management of risk*, 100:116, 1999.

[99] Diomidis Spinellis. Reflections on trusting trust revisited. *Communications of the ACM*, 46(6):112, 2003.

[100] Lars Hopland Nestas and Kjell J. Hole. Building and maintaining trust in internet voting. *Computer*, 45(5):74–80, 2012.

[101] Peter E. Hart and Ziming Liu. Trust in the preservation of digital information. *Communications of the ACM*, 46(6): 93–97, 2003.

[102] Paul B. de Laat. Trusting virtual trust. *Journal of Ethics and Information Technology*, 7(3):167–180, 2005.

[103] David Gefen, Elena Karahanna, and Detmar W Straub. Trust and tam in online shopping: an integrated model. *MIS quarterly*, 27(1):51–90, 2003.

[104] Erkki Patokorpi and Kai K. Kimppa. Dynamics of the key elements of consumer trust building online. *Journal of Information, Communication and Ethics in Society*, 4(1):17–26, 2006.

[105] Daniel W. Manchala. Trust metrics, models and protocols for electronic commerce transactions. In *Distributed Computing Systems, 1998. Proceedings. 18th International Conference on*, pages 312–321. IEEE, 1998.

[106] Roger C. Mayer, James H. Davis, and F. David Schoorman. An integrative model of organizational trust. *Academy of management review*, 20(3):709–734, 1995.

[107] Juha Vuorinen. *Parasitic Order Machine. A Sociology and Ontology of Information Securing*. Annales universitatis turkuensis b392, University of Turku, Finland, 2014.

[108] Stephen Marsh and Mark R. Dibben. Trust, untrust, distrust and mistrust–an exploration of the dark (er) side. In *Trust Management*, pages 17–33. Springer, 2005.

[109] Matteo Dell'Amico, Pietro Michiardi, and Yves Roudier. Password strength: An empirical analysis. In *INFOCOM*, volume 10, pages 983–991, 2010.

[110] Anatol Rapoport and Albert M. Chammah. *Prisoner's dilemma: A study in conflict and cooperation*, volume 165. University of Michigan press, 1965.

[111] Garrett Hardin. The tragedy of the commons. *Science*, 162 (3859):1243–1248, 1968.

[112] Donald A. MacKenzie. *Inventing accuracy: a historical sociology of nuclear missile guidance*. MIT Press, 1990.

[113] Antti Hakkala. Analysis of an iris recongition system based on partial iris patterns. Master's thesis, University of Turku, 2009.

[114] Brendan F. Klare, Ben Klein, Emma Taborsky, Austin Blanton, Jordan Cheney, Kristen Allen, Patrick Grother, Alan Mah, Mark Burge, and Anil K. Jain. Pushing the frontiers of unconstrained face detection and recognition: IARPA Janus Benchmark A. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1931–1939. IEEE, 2015.

[115] Haoqiang Fan and Erjin Zhou. Approaching human level facial landmark localization by deep learning. *Image and Vision Computing*, 47:27–35, 2016.

[116] James L. Wayman. Fundamentals of biometric authentication technologies. *International Journal of Image and Graphics*, 1(1):93–113, 2001.

[117] John Daugman. Probing the uniqueness and randomness of iriscodes: Results from 200 billion iris pair comparisons. *Proceedings of the IEEE*, 94(11):1927–1935, Nov. 2006. ISSN 0018-9219. doi: 10.1109/JPROC.2006.884092.

[118] Benjamin Miller. Vital signs of identity. *IEEE Spectrum*, 31 (2):22–30, Feb 1994.

[119] Alphonse Bertillon. *La couleur de l'iris*. Masson, 1886.

[120] Anil K. Jain, Ruud Bolle, and Sharath Pankati, editors. *Biometrics: Personal Identification in Networked Society*. Kluwer, New York, 1999.

[121] Anil K. Jain, Arun Ross, and Salil Prabhakar. An introduction to biometric recognition. *IEEE Transactions on Circuits and Systems for Video Technology*, 14(1):4–20, Jan 2004.

[122] John D. Woodward. Biometrics: privacy's foe or privacy's friend? *Proceedings of the IEEE*, 85(9):1480–1492, Sep 1997. ISSN 0018-9219. doi: 10.1109/5.628723.

[123] James L. Wayman. Technical testing and evaluation of biometric identification devices. In Anil K. Jain; Ruud

Bolle; Sharath Pankati, editor, *Biometrics: Personal Identification in Networked Society*. Kluwer, New York, 1999.

[124] John Daugman. Biometric decision landscapes. Technical Report UCAM-CL-TR-482, Computer Laboratory, University of Cambridge, 2000.

[125] Nalini K. Ratha, Jonathan H. Connell, and Ruud M. Bolle. Enhancing security and privacy in biometrics-based authentication systems. *IBM Systems Journal*, 40(3):614–634, 2001.

[126] Ari Juels and Madhu Sudan. A fuzzy vault scheme. *Designs, Codes and Cryptography*, 38(2):237–257, 2006.

[127] Umut Uludag, Sharath Pankanti, and Anil K. Jain. Fuzzy vault for fingerprints. In *Audio-and Video-Based Biometric Person Authentication*, pages 310–319. Springer, 2005.

[128] Raffaele Cappelli, Alessandra Lumini, Dario Maio, and Davide Maltoni. Fingerprint image reconstruction from standard templates. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 29(9):1489–1503, 2007.

[129] Michael Cherry and Edward Imwinkelried. A cautionary note about fingerprint analysis and reliance on digital technology. *Judicature*, 89:334, 2006.

[130] Itiel E. Dror, David Charlton, and Ailsa E. Péron. Contextual information renders experts vulnerable to making erroneous identifications. *Forensic Science International*, 156 (1):74–78, 2006.

[131] International Civil Aviation Organization. Machine readable travel documents. Available online at http://www.icao.int/Security/mrtd/Pages/Document9303.aspx, accessed 22.10.2012, 2006.

[132] Wojciech Mostowski and Erik Poll. Electronic passports in a nutshell. Technical report, Technical Report ICIS–R10004, Radboud University Nijmegen, 2010.

[133] International Civil Aviation Organization. Doc 9303. Machine Readable Travel Documents, Seventh Edition. Part

10: Logical Data Structure (LDS) for Storeage of Biometrics and Other Data in the Contactless Integrated Circuit (IC). Available online at https://www.icao.int/publications/pages/publication.aspx?docnum=9303, 2015.

[134] International Civil Aviation Organization. Doc 9303. Machine Readable Travel Documents, Seventh Edition. Part 11: Security Mechanisms for MRTDs. Available online at https://www.icao.int/publications/pages/publication.aspx?docnum=9303, 2015.

[135] Jean Monnerat, Serge Vaudenay, and Martin Vuagnoux. About machine-readable travel documents. In *In Proceedings of the International Conference on RFID Security 2007*. Citeseer, 2007.

[136] Federal Office for Information Security (BSI). Advanced Security Mechanisms for Machine Readable Travel Documents and eIDAS Token – Part 2: Protocols for electronic IDentification, Authentication and trust Services (eIDAS) . Available online at https://www.bsi.bund.de/EN/Publications/TechnicalGuidelines/TR03110/BSITR03110.html, 2016.

[137] Rafik Chaabouni and Serge Vaudenay. The extended access control for machine readable travel documents. *BIOSIG 2009, Biometrics and Electronic Signatures, LNI, Gesellschaft für Informatik (GI), Bonn, Germany*, pages 93–103, 2009.

[138] Joan Daemen and Vincent Rijmen. The block cipher Rijndael. In *CARDIS*, volume 1820, pages 277–284. Springer, 1998.

[139] Jaap-Henk Hoepman, Engelbert Hubbers, Bart Jacobs, Martijn Oostdijk, and Ronny Wichers Schreur. Crossing borders: Security and privacy issues of the european e-passport. In *Advances in Information and Computer Security*, pages 152–167. Springer, 2006.

[140] Ari Juels, David Molnar, and David Wagner. Security and privacy issues in e-passports. In *Security and Privacy*

for Emerging Areas in Communications Networks, 2005. SecureComm 2005. First International Conference on, pages 74–88. IEEE, 2005.

[141] Henning Richter, Wojciech Mostowski, and Erik Poll. Fingerprinting passports. In *NLUUG spring conference on security*, pages 21–30, 2008.

[142] Boris Danev, Thomas S Heydt-Benjamin, and Srdjan Capkun. Physical-layer identification of rfid devices. In *Proceedings of the USENIX Security Symposium*, pages 199–214, 2009.

[143] Chaoming Song, Zehui Qu, Nicholas Blumm, and Albert-László Barabási. Limits of predictability in human mobility. *Science*, 327(5968):1018–1021, 2010.

[144] Hui Zang and Jean Bolot. Anonymization of location data does not work: A large-scale measurement study. In *Proceedings of the 17th Annual International Conference on Mobile Computing and Networking*, MobiCom '11, pages 145–156, New York, NY, USA, 2011. ACM. ISBN 978-1-4503-0492-4. doi: 10.1145/2030613.2030630. URL http://doi.acm.org/10.1145/2030613.2030630.

[145] Gerrit Hornung. The european regulation on biometric passports: Legislative procedures, political interactions, legal framework and technical safeguards. *SCRIPT-ed*, 4(3): 246–262, 2007.

[146] European Court of Human Rights (ECHR). The retention of the fingerprints of a person who had not been convicted breached his right to respect for his private life. Available online at http://hudoc.echr.coe.int/webservices/content/pdf/003-4332390-5192548 Last accessed 3.5.2013, 2013.

[147] Gildas Avoine, Antonin Beaujeant, Julio Hernandez-Castro, Louis Demay, and Philippe Teuwen. A survey of security and privacy issues in epassport protocols. *ACM Computing Surveys (CSUR)*, 48(3):47, 2016.

[148] Leonard Flom and Aran Safir. Iris recognition system. United States Patent No. 4641349, U.S. Government Printing Office, Washington, DC, 1987.

[149] John Daugman. High confidence visual recognition of persons by a test of statistical independence. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15(11):1148–1161, Nov 1993. ISSN 0162-8828. doi: 10.1109/34.244676.

[150] R.P. Wildes, J.C. Asmuth, G.L. Green, S.C. Hsu, R.J. Kolczynski, J.R. Matey, and S.E. McBride. A system for automated iris recognition. In *Applications of Computer Vision, 1994., Proceedings of the Second IEEE Workshop on*, pages 121–128, Dec 1994. doi: 10.1109/ACV.1994.341298.

[151] Rebecca T. Mercuri. *Electronic vote tabulation checks and balances*. PhD thesis, University of Pennsylvania, 2001.

[152] Richard F. Celeste, Dick Thornburgh, and Herbert Lin. *Asking the right questions about electronic voting*. National Academy Press, 2005.

[153] Roy G. Saltman. Accuracy, integrity and security in computerized vote-tallying. *Communications of the ACM*, 31(10): 1184–1191, 1988.

[154] Tadayoshi Kohno, Adam Stubblefield, Aviel D. Rubin, and Dan S. Wallach. Analysis of an electronic voting system. In *Security and Privacy, 2004. Proceedings. 2004 IEEE Symposium on*, pages 27–40. IEEE, 2004.

[155] Juhani Karhumäki, Arto Lepistö, Sami Mäkelä, Hannu Nurmi, Tommi Penttinen, Ari Renvall, Petri Salmela, and Seppo Virtanen. Auditointiraportti kunnallisvaalien sähköisen äänestyksen pilotista. National Publication 15, Turku Centre for Computer Science TUCS, October 2008.

[156] Krishna Sampigethaya and Radha Poovendran. A framework and taxonomy for comparison of electronic voting schemes. *Computers & Security*, 25(2):137–153, 2006.

[157] Dimitris A. Gritzalis. Principles and requirements for a secure e-voting system. *Computers & Security*, 21(6):539 – 556, 2002. ISSN 0167-4048. doi: 10.1016/S0167-4048(02)

01014-3. URL http://www.sciencedirect.com/science/article/pii/S0167404802010143.

[158] Costas Lambrinoudakis, Dimitris A. Gritzalis, Vassilis Tsoumas, Maria Karyda, and Spyros Ikonomopoulos. Secure electronic voting: The current landscape. *Secure electronic voting*, pages 101–122, 2003.

[159] Orhan Çetinkaya. *Verifiability and receipt-freeness in cryptographic voting systems*. PhD thesis, Middle East Technical University, 2007.

[160] Lucie Langer, Axel Schmidt, Johannes Buchmann, Melanie Volkamer, and Alexander Stolfik. Towards a framework on the security requirements for electronic voting protocols. In *Requirements Engineering for e-Voting Systems (REVOTE), 2009 First International Workshop on*, pages 61–68. IEEE, 2010.

[161] Lucie Langer, Axel Schmidt, Johannes Buchmann, and Melanie Volkamer. A taxonomy refining the security requirements for electronic voting: analyzing helios as a proof of concept. In *Availability, Reliability, and Security, 2010. ARES'10 International Conference on*, pages 475–480. IEEE, 2010.

[162] Lucie Langer. *Privacy and verifiability in electronic voting*. PhD thesis, TU Darmstadt, 2010.

[163] Feng Yumeng, Tian Liye, Liu Fanbao, and Gan Chong. Electronic voting: A review and taxonomy. In *Industrial Control and Electronics Engineering (ICICEE), 2012 International Conference on*, pages 912–917. IEEE, 2012.

[164] Stephan Neumann and Melanie Volkamer. A holistic framework for the evaluation of internet voting systems. *Design, Development, and Use of Secure Electronic Voting Systems*, pages 76–91, 2014.

[165] Drew Springall, Travis Finkenauer, Zakir Durumeric, Jason Kitcat, Harri Hursti, Margaret MacAlpine, and J. Alex Halderman. Security analysis of the estonian internet voting system. In *Proceedings of the 2014 ACM SIGSAC Confer-*

ence on Computer and Communications Security, pages 703–715. ACM, 2014.

[166] Stephan Neumann, Melanie Volkamer, Jurlind Budurushi, and Marco Prandini. Secivo: a quantitative security evaluation framework for internet voting schemes. *Annals of Telecommunications*, 71(7-8):337–352, 2016.

[167] Margaret Cocker and Paul Sonne. Ukraine: Cyberwar's hottest front. *The Wall Street Journal*, November 9th 2015. URL http://www.wsj.com/articles/ukraine-cyberwars-hottest-front-1447121671. Accessed 09.03.2016.

[168] Danny Dolev and Andrew Yao. On the security of public key protocols. *IEEE Transactions on information theory*, 29 (2):198–208, 1983.

[169] Thomas H. Davenport, Robert G. Eccles, and Laurence Prusak. Information politics. *The strategic management of intellectual capital*, pages 101–120, 1998.

[170] Paul A. Strassmann. *The politics of information management*. The Information Economics Press, 1994.

[171] Thomas C. Redman. The impact of poor data quality on the typical enterprise. *Communications of the ACM*, 41(2): 79–82, 1998.

[172] Ville Kainu and Jani Koskinen. Between public and personal information-not prohibited, therefore permitted? *Privacy and Surveillance-current aspects and future perspectives*, pages 45–59, 2012.

[173] Jani S. S. Koskinen, Ville Kainu, and Kai K. Kimppa. The concept of datenherrschaft of patient information from a lockean perspective. *Journal of Information, Communication and Ethics in Society*, 14(1):70–86, 2016.

[174] Charles Matthew Weir. *Using probabilistic techniques to aid in password cracking attacks*. PhD thesis, Florida State University, 2010.

[175] Joseph Bonneau. The science of guessing: analyzing an anonymized corpus of 70 million passwords. In *Security*

and Privacy (SP), 2012 IEEE Symposium on, pages 538–552. IEEE, 2012.

[176] Joseph Bonneau and Ekaterina Shutova. Linguistic properties of multi-word passphrases. In *Financial Cryptography and Data Security*, pages 1–12. Springer, 2012.

[177] Thorsten Brantz and Alex Franz. The google web it 5-gram corpus. Technical Report LDC2006T13, Lingusitic Data Consortium, 2006.

[178] Dino Schweitzer, Jeff Boleng, Colin Hughes, and Louis Murphy. Visualizing keyboard pattern passwords. In *Visualization for Cyber Security, 2009. VizSec 2009. 6th International Workshop on*, pages 69–73. IEEE, 2009.

[179] Matt Weir, Sudhir Aggarwal, Breno De Medeiros, and Bill Glodek. Password cracking using probabilistic context-free grammars. In *Security and Privacy, 2009 30th IEEE Symposium on*, pages 391–405. IEEE, 2009.

[180] Tommi Laine. Suomalaisten salasanojen vahvuus ja salasanamallit. B.sc. thesis, University of Turku, Department of Information Technology, 2012.

[181] James M. Anderson. Why we need a new definition of information security. *Computers & Security*, 22(4):308–313, 2003.

[182] Aleph One. Smashing the stack for fun and profit. *Phrack magazine*, 7(49):14–16, 1996.

[183] François Koeune and François-Xavier Standaert. A tutorial on physical security and side-channel attacks. In *Foundations of Security Analysis and Design III*, pages 78–108. Springer, 2005.

[184] William Stallings. *Network security essentials: applications and standards, Fourth edition*. Pearson Prentice Hall, 2011.

[185] Gary McGraw. *Software security: building security in*. Pearson Education, 2006.

[186] Basie Von Solms and Rossouw Von Solms. The 10 deadly sins of information security management. *Computers & Security*, 23(5):371–376, 2004.

[187] Harold F. Tipton and Micki Krause. *Information security management handbook*. CRC Press, 2003.

[188] Kevin D Mitnick and William L Simon. *The art of deception: Controlling the human element of security*. John Wiley & Sons, 2011.

[189] Christopher Hadnagy. *Social engineering: The art of human hacking*. John Wiley & Sons, 2010.

[190] MikkoT Siponen. Five dimensions of information security awareness. *Computers and society*, 31(2):24–29, 2001.

[191] Tejaswini Herath and H Raghav Rao. Encouraging information security behaviors in organizations: Role of penalties, pressures and perceived effectiveness. *Decision Support Systems*, 47(2):154–165, 2009.

[192] Steve Sheng, Mandy Holbrook, Ponnurangam Kumaraguru, Lorrie Faith Cranor, and Julie Downs. Who falls for phish?: a demographic analysis of phishing susceptibility and effectiveness of interventions. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 373–382. ACM, 2010.

[193] Yacine Rezgui and Adam Marks. Information security awareness in higher education: An exploratory study. *Computers & Security*, 27(7):241–253, 2008.

[194] Max M. North, Roy George, and Sarah M. North. Computer security and ethics awareness in university environments: A challenge for management of information systems. In *Proceedings of the 44th annual Southeast regional conference*, pages 434–439. ACM, 2006.

[195] Maslin Masrom, Zuraini Ismail, and Ramlah Hussein. Ethical awareness of computer use among undergraduate students. *ACM SIGCAS Computers and Society*, 39(1):27–40, 2009.

[196] Mansur Aliyu, Nahel A.O. Abdallah, Nojeem A. Lasisi, Dahir Diyar, and Ahmed M. Zeki. Computer security and ethics awareness among iium students: An empirical study. In *Information and Communication Technology for the Muslim World (ICT4M), 2010 International Conference on*, pages A52–A56. IEEE, 2010.

[197] Stefan Fenz and Andreas Ekelhart. Formalizing information security knowledge. In *Proceedings of the 4th international Symposium on information, Computer, and Communications Security*, pages 183–194. ACM, 2009.

[198] Mario Silic and Andrea Back. Information security: Critical review and future directions for research. *Information Management & Computer Security*, 22(3):279–308, 2014.

[199] Carlos Blanco, Joaquín Lasheras, Eduardo Fernández-Medina, Rafael Valencia-García, and Ambrosio Toval. Basis for an integrated security ontology according to a systematic review of existing proposals. *Computer Standards & Interfaces*, 33(4):372–388, 2011.

[200] Mikko-Jussi Laakso, Erkki Kaila, and Teemu Rajala. ViLLE - Designing and utilizing a collaborative learning environment. *Submitted to Computers & Education*, 2015.

[201] Lorin W. Anderson, David R. Krathwohl, Peter W. Airasian, Kathleen A. Cruikshank, Richard E. Mayer, PPaul R. Pintrich, James Raths, and Merlin C. Wittrock. A taxonomy for learning, teaching and assessing: A revision of Bloom's taxonomy. *New York. Longman Publishing. Artz, AF, & Armour-Thomas, E.(1992). Development of a cognitive-metacognitive framework for protocol analysis of mathematical problem solving in small groups. Cognition and Instruction*, 9 (2):137–175, 2001.

[202] David R. Krathwohl. A revision of Bloom's taxonomy: An overview. *Theory into practice*, 41(4):212–218, 2002.

[203] Jaroslav Král and Michal Žemlička. Bottleneck of knowledge society. In *World Summit on Knowledge Society*, pages 83–91. Springer, 2008.

[204] John F. Sanford and Les M. Sztandera. Thoughts on the future of education in information technology. *International Journal of Teaching and Case Studies*, 1(1-2):23–32, 2007.

[205] Hisham Al-Mubaid. Designing and managing intervention methods to promote self-regulated learning. *International Journal of Teaching and Case Studies*, 1(3):224–233, 2008.

[206] Seppo Virtanen. Improving the learning process of engineering students by deployment of activating icts. In *World Summit on Knowledge Society*, pages 328–333. Springer, 2008.

[207] Seppo Virtanen. Increasing the self-study effort of higher education engineering students with an online learning platform. *International Journal of Knowledge and Learning*, 4(6):527–538, 2008.

[208] Giovanni Vigna. Teaching network security through live exercises. In *Security education and critical infrastructures*, pages 3–18. Springer, 2003.

[209] Dalibor Dobrilovic, Vesna Jevtic, Zeljko Stojanov, and Borislav Odadzic. Usability of virtual network laboratory in engineering education and computer network course. In *Interactive Collaborative Learning (ICL), 2012 15th International Conference on*, pages 1–6. IEEE, 2012.

[210] Keith W. Miller, Jeffrey Voas, and George F. Hurlburt. Byod: Security and privacy considerations. *It Professional*, 14(5):53–55, 2012.

[211] Gordon Thomson. BYOD: enabling the chaos. *Network Security*, 2012(2):5–8, 2012.

[212] Klaus Krippendorff. *Content analysis: An introduction to its methodology*. Sage, 2012.

[213] Tyler Moore, Richard Clayton, and Ross Anderson. The economics of online crime. *The Journal of Economic Perspectives*, 23(3):3–20, 2009.

[214] Jerker Björkqvist and Seppo Virtanen. Convergence of hardware and software in platforms for radio technologies. *Communications Magazine, IEEE*, 44(11):52–57, November 2006.

[215] Jouni Isoaho, Seppo Virtanen, and Juha Plosila. Current challenges in embedded communication systems. *International journal of embedded and real-time communication systems*, 1(1):1–21, 2010.

[216] Seppo Virtanen, Tero Nurmi, Jani Paakkulainen, and Johan Lilius. A system-level framework for designing and evaluating protocol processor architectures. *International journal of embedded systems*, 1(12):78–90, 2006.

[217] Jani Paakkulainen, Seppo Virtanen, and Jouni Isoaho. Tuning a protocol processor architecture towards dsp operations. In Timo D. Hämäläinen, Andy D. Pimentel, Jarmo Takala, and Stamatis Vassiliadis, editors, *Embedded Computer Systems: Architectures, Modeling, and Simulation: 5th International Workshop, SAMOS 2005, Samos, Greece, July 18-20, 2005. Proceedings*, pages 132–141. Springer Berlin Heidelberg, Berlin, Heidelberg, 2005. ISBN 978-3-540-31664-0. doi: 10.1007/11512622_15. URL http://dx.doi.org/10.1007/11512622_15.

[218] Muhammad I. Anwar, Seppo Virtanen, and Jouni Isoaho. A software defined approach for common baseband processing. *Journal of Systems Architecture*, 54(8):769–786, 2008.

[219] Henk Corporaal. *Microprocessor Architectures from VLIW to TTA*. John Wiley and Sons Ltd., Chichester, England, 1998.

[220] Dragos Truscan, Seppo Virtanen, and Johan Lilius. Protocol processor design issues. In Jari Nurmi, editor, *Processor design: System-on-Chip computing for ASICs and FPGAs*, pages 257–285. Springer, Dordrecht, the Netherlands, 2007.

[221] Panu Hämäläinen, Marko Hännikäinen, Timo D. Hämäläinen, Henk Corporaal, and Jukka Saarinen. Implementation of encryption algorithms on transport triggered architectures. In *Proceedings of IEEE International Symposium on Circuits and Systems*, volume 4, pages 726–729, 2001.

[222] Panu Hämäläinen, Jari Heikkinen, Marko Hännikäinen, and Timo D. Hämäläinen. Design of transport triggered architecture processors for wireless encryption. In *Proceedings of 8th Euromicro Conference on Digital System Design*, 2005.

[223] Axel Jantsch, Johnny Öberg, and Ahmed Hemani. Is there a niche for a general protocol processor core? In *Proceedings of the 16th IEEE Norchip Conference*, pages 93–100, 1998.

[224] Niels Provos and David Mazieres. A future-adaptable password scheme. In *USENIX Annual Technical Conference, FREENIX Track*, pages 81–91, 1999.

[225] Paul Kocher, Ruby Lee, Gary McGraw, and Anand Raghunathan. Security as a new dimension in embedded system design. In *Proceedings of the 41st Annual Design Automation Conference*, 2004.

[226] Ronald L. Rivest, Adi Shamir, and Leonard Adleman. A method for obtaining digital signatures and public-key cryptosystems. *Communications of the ACM*, 21(2):126, 1978.

[227] Peter L. Montgomery. Modular Multiplication Without Trial Division. *Mathematics of Computation*, 44(170):519–521, 1985.

[228] Klaus Hansen, Troels Larsen, and Kim Olsen. On the effciency of fast RSA variants in modern mobile phones. *International Journal of Computer Science and Information Security*, 6(3):136–140, 2009.

[229] Victor Miller. Use of elliptic curves in cryptography. *Proceedings of the Advances in Cryptology—CRYPTO'85*, 1986.

[230] Neil Koblitz. Elliptic curve cryptosystems. *Mathematics of Computation*, 48:203–209, 1987.

[231] Mohamed Khalil-Hani and Yuan Wen Hau. SystemC HW/SW co-design methodology applied to the design of an elliptic curve crypto system on chip. In *Microelectronics, 2008. ICM 2008. International Conference on*, pages 147–150. IEEE, 2008.

[232] Stefan Tillich and Johann Großschädl. Accelerating AES using instruction set extensions for elliptic curve cryptography. In *Proceedings of Computational Science and Its Applications ICCSA 2005*, pages 665–675, 2005.

[233] Ryad Benadjila, Olivier Billet, Shay Gueron, and Matt J.B. Robshaw. The Intel AES Instructions Set and the SHA-3 Candidates. In *ASIACRYPT 2009: Proceedings of the 15th International Conference on the Theory and Application of Cryptology and Information Security*, pages 162–178. Springer-Verlag, 2009.

[234] Daniel M. Gordon. A survey of fast exponentiation methods. *Journal of algorithms*, 27(1):129–146, 1998.

[235] Paul Barrett. Implementing the Rivest Shamir and Adleman public key encryption algorithm on a standard digital signal processor. In *Advances in Cryptology—CRYPTO'86*, pages 311–323. Springer, 1987.

[236] Seong Ming Hong, Sang-Yeop Oh, and Hyunsoo Yoon. New modular multiplication algorithms for fast modular exponentiation. In *Advances in Cryptology — EUROCRYPT'96*, pages 166–177. Springer, 1996.

[237] William Hasenplaugh, Gunnar Gaubatz, and Vinodh Gopal. Fast modular reduction. *18th IEEE Symposium on Computer Arithmetic (ARITH'07)*, 2007. ISSN 1063-6889.

[238] Chia-Long Wu. An efficient common-multiplicand-multiplication method to the Montgomery algorithm for speeding up exponentiation. *Information Sciences*, 179(4): 410 – 421, 2009.

[239] Toshiya Itoh and Shigeo Tsujii. A fast algorithm for computing multiplicative inverses in GF(2^m) using normal bases. *Information and Computation*, 78(3):171–177, 1988.

[240] Kimmo U. Järvinen. On repeated squarings in binary fields. In *Proceedings of the 16th International Workshop on Selected Areas in Cryptography, SAC 2009*, volume 5867 of *Lecture Notes in Computer Science*, pages 331–349. Springer-Verlag, 2009.

[241] Darrel Hankerson, Alfred J. Menezes, and Scott Vanstone. *Guide to elliptic curve cryptography*. Springer Science & Business Media, 2006.

[242] Sandro Bartolini, Roberto Giorgi, and Enrico Martinelli. Instruction Set Extensions for Cryptographic Applications. In Çetin Kaya Koç, editor, *Cryptographic Engineering*, pages 191–233. Springer, 2009.

[243] Panu Hämäläinen, Marko Hännikäinen, and Timo D. Hämäläinen. Review of hardware architectures for advanced encryption standard implementations considering wireless sensor networks. In *Embedded Computer Systems: Architectures, Modeling, and Simulation*, pages 443–453. Springer, 2007.

[244] Yanhua Liu, Wei Guo, Ya Tan, Jizeng Wei, and Dazhi Sun. An efficient scheme for implementation of sm2 digital signature over gf (p). In *Contemporary Research on E-business Technology and Strategy*, pages 250–258. Springer, 2012.

[245] Wei Guo, Yaling Liu, Songhui Bai, Jizeng Wei, and Dazhi Sun. Hardware architecture for RSA cryptography based on residue number system. *Transactions of Tianjin University*, 18:237–242, 2012.

[246] Jingwei Hu, Wei Guo, Jizeng Wei, Yisong Chang, and Dazhi Sun. A novel architecture for fast rsa key generation based on rns. In *Parallel Architectures, Algorithms and Programming (PAAP), 2011 Fourth International Symposium on*, pages 345–349. IEEE, 2011.

[247] Patrick Schaumont and Ingrid Verbauwhede. A reconfiguration hierarchy for elliptic curve cryptography. In *ASILOMAR Conference on signals, systems and computers*, volume 1, pages 449–453. IEEE; 1998, 2001.

[248] Andreas Dandalis and Viktor K. Prasanna. An adaptive cryptographic engine for internet protocol security architectures. *ACM Transactions on Design Automation of Electronic Systems*, 9(3):333–353, July 2004. ISSN 1084-4309. doi: 10.1145/1013948.1013952. URL http://doi.acm.org/10.1145/1013948.1013952.

[249] Andreas Dandalis, Viktor K. Prasanna, and Jose D.P. Rolim. An adaptive cryptographic engine for IPSec architectures. In *Field-Programmable Custom Computing Machines, 2000 IEEE Symposium on*, pages 132–141. IEEE, 2000.

[250] Shafi Goldwasser and Silvio Micali. Probabilistic encryption. *Journal of computer and system sciences*, 28(2):270–299, 1984.

[251] Paul C. Kocher. Timing attacks on implementations of diffie-hellman, rsa, dss, and other systems. In *Advances in Cryptology—CRYPTO'96*, pages 104–113. Springer, 1996.

[252] Stephen Checkoway, Matthew Fredrikson, Ruben Niederhagen, Adam Everspaugh, Matthew Green, Tanja Lange, Thomas Ristenpart, Daniel J. Bernstein, Jake Maskiewicz, and Hovav Shacham. On the practical exploitability of Dual EC in TLS implementations. In *USENIX Security*, volume 1, 2014.

[253] Lilian Bossuet, Michael Grand, Lubos Gaspar, Viktor Fischer, and Guy Gogniat. Architectures of flexible symmetric key crypto engines – a survey: From hardware coprocessor to multi-crypto-processor system on chip. *ACM Computing Surveys (CSUR)*, 45(4):41, 2013.

[254] Lubos Gaspar, Viktor Fischer, Florent Bernard, Lilian Bossuet, and Pascal Cotret. Hcrypt: a novel concept of crypto-processor with secured key management. In *Reconfigurable Computing and FPGAs (ReConFig), 2010 International Conference on*, pages 280–285. IEEE, 2010.

[255] Guy Gogniat, Tilman Wolf, Wayne Burleson, Jean-Philippe Diguet, Lilian Bossuet, and Romain Vaslin. Reconfigurable hardware for high-security/high-performance embedded systems: the safes perspective. *Very Large Scale Integration (VLSI) Systems, IEEE Transactions on*, 16(2):144–155, 2008.

[256] Pramesh Khanal. Modeling and Simulating ASIPs for Cryptography in SystemC. Master's thesis, University of Turku, 2013.

[257] Bruce Schneier. *Beyond fear: thinking sensibly about security in an uncertain world*. Copernicus books, 2003.

Turku Centre for Computer Science
TUCS Dissertations

1. **Marjo Lipponen**, On Primitive Solutions of the Post Correspondence Problem
2. **Timo Käkölä**, Dual Information Systems in Hyperknowledge Organizations
3. **Ville Leppänen**, Studies on the Realization of PRAM
4. **Cunsheng Ding**, Cryptographic Counter Generators
5. **Sami Viitanen**, Some New Global Optimization Algorithms
6. **Tapio Salakoski**, Representative Classification of Protein Structures
7. **Thomas Långbacka**, An Interactive Environment Supporting the Development of Formally Correct Programs
8. **Thomas Finne**, A Decision Support System for Improving Information Security
9. **Valeria Mihalache**, Cooperation, Communication, Control. Investigations on Grammar Systems.
10. **Marina Waldén**, Formal Reasoning About Distributed Algorithms
11. **Tero Laihonen**, Estimates on the Covering Radius When the Dual Distance is Known
12. **Lucian Ilie**, Decision Problems on Orders of Words
13. **Jukkapekka Hekanaho**, An Evolutionary Approach to Concept Learning
14. **Jouni Järvinen**, Knowledge Representation and Rough Sets
15. **Tomi Pasanen**, In-Place Algorithms for Sorting Problems
16. **Mika Johnsson**, Operational and Tactical Level Optimization in Printed Circuit Board Assembly
17. **Mats Aspnäs**, Multiprocessor Architecture and Programming: The Hathi-2 System
18. **Anna Mikhajlova**, Ensuring Correctness of Object and Component Systems
19. **Vesa Torvinen**, Construction and Evaluation of the Labour Game Method
20. **Jorma Boberg**, Cluster Analysis. A Mathematical Approach with Applications to Protein Structures
21. **Leonid Mikhajlov**, Software Reuse Mechanisms and Techniques: Safety Versus Flexibility
22. **Timo Kaukoranta**, Iterative and Hierarchical Methods for Codebook Generation in Vector Quantization
23. **Gábor Magyar**, On Solution Approaches for Some Industrially Motivated Combinatorial Optimization Problems
24. **Linas Laibinis**, Mechanised Formal Reasoning About Modular Programs
25. **Shuhua Liu**, Improving Executive Support in Strategic Scanning with Software Agent Systems
26. **Jaakko Järvi**, New Techniques in Generic Programming – C++ is more Intentional than Intended
27. **Jan-Christian Lehtinen**, Reproducing Kernel Splines in the Analysis of Medical Data
28. **Martin Büchi**, Safe Language Mechanisms for Modularization and Concurrency
29. **Elena Troubitsyna**, Stepwise Development of Dependable Systems
30. **Janne Näppi**, Computer-Assisted Diagnosis of Breast Calcifications
31. **Jianming Liang**, Dynamic Chest Images Analysis
32. **Tiberiu Seceleanu**, Systematic Design of Synchronous Digital Circuits
33. **Tero Aittokallio**, Characterization and Modelling of the Cardiorespiratory System in Sleep-Disordered Breathing
34. **Ivan Porres**, Modeling and Analyzing Software Behavior in UML
35. **Mauno Rönkkö**, Stepwise Development of Hybrid Systems
36. **Jouni Smed**, Production Planning in Printed Circuit Board Assembly
37. **Vesa Halava**, The Post Correspondence Problem for Market Morphisms
38. **Ion Petre**, Commutation Problems on Sets of Words and Formal Power Series
39. **Vladimir Kvassov**, Information Technology and the Productivity of Managerial Work
40. **Frank Tétard**, Managers, Fragmentation of Working Time, and Information Systems

41. **Jan Manuch**, Defect Theorems and Infinite Words
42. **Kalle Ranto**, Z_4-Goethals Codes, Decoding and Designs
43. **Arto Lepistö**, On Relations Between Local and Global Periodicity
44. **Mika Hirvensalo**, Studies on Boolean Functions Related to Quantum Computing
45. **Pentti Virtanen**, Measuring and Improving Component-Based Software Development
46. **Adekunle Okunoye**, Knowledge Management and Global Diversity – A Framework to Support Organisations in Developing Countries
47. **Antonina Kloptchenko**, Text Mining Based on the Prototype Matching Method
48. **Juha Kivijärvi**, Optimization Methods for Clustering
49. **Rimvydas Rukšėnas**, Formal Development of Concurrent Components
50. **Dirk Nowotka**, Periodicity and Unbordered Factors of Words
51. **Attila Gyenesei**, Discovering Frequent Fuzzy Patterns in Relations of Quantitative Attributes
52. **Petteri Kaitovaara**, Packaging of IT Services – Conceptual and Empirical Studies
53. **Petri Rosendahl**, Niho Type Cross-Correlation Functions and Related Equations
54. **Péter Majlender**, A Normative Approach to Possibility Theory and Soft Decision Support
55. **Seppo Virtanen**, A Framework for Rapid Design and Evaluation of Protocol Processors
56. **Tomas Eklund**, The Self-Organizing Map in Financial Benchmarking
57. **Mikael Collan**, Giga-Investments: Modelling the Valuation of Very Large Industrial Real Investments
58. **Dag Björklund**, A Kernel Language for Unified Code Synthesis
59. **Shengnan Han**, Understanding User Adoption of Mobile Technology: Focusing on Physicians in Finland
60. **Irina Georgescu**, Rational Choice and Revealed Preference: A Fuzzy Approach
61. **Ping Yan**, Limit Cycles for Generalized Liénard-Type and Lotka-Volterra Systems
62. **Joonas Lehtinen**, Coding of Wavelet-Transformed Images
63. **Tommi Meskanen,** On the NTRU Cryptosystem
64. **Saeed Salehi**, Varieties of Tree Languages
65. **Jukka Arvo**, Efficient Algorithms for Hardware-Accelerated Shadow Computation
66. **Mika Hirvikorpi**, On the Tactical Level Production Planning in Flexible Manufacturing Systems
67. **Adrian Costea**, Computational Intelligence Methods for Quantitative Data Mining
68. **Cristina Seceleanu**, A Methodology for Constructing Correct Reactive Systems
69. **Luigia Petre**, Modeling with Action Systems
70. **Lu Yan**, Systematic Design of Ubiquitous Systems
71. **Mehran Gomari**, On the Generalization Ability of Bayesian Neural Networks
72. **Ville Harkke**, Knowledge Freedom for Medical Professionals – An Evaluation Study of a Mobile Information System for Physicians in Finland
73. **Marius Cosmin Codrea**, Pattern Analysis of Chlorophyll Fluorescence Signals
74. **Aiying Rong**, Cogeneration Planning Under the Deregulated Power Market and Emissions Trading Scheme
75. **Chihab BenMoussa**, Supporting the Sales Force through Mobile Information and Communication Technologies: Focusing on the Pharmaceutical Sales Force
76. **Jussi Salmi**, Improving Data Analysis in Proteomics
77. **Orieta Celiku**, Mechanized Reasoning for Dually-Nondeterministic and Probabilistic Programs
78. **Kaj-Mikael Björk**, Supply Chain Efficiency with Some Forest Industry Improvements
79. **Viorel Preoteasa**, Program Variables – The Core of Mechanical Reasoning about Imperative Programs
80. **Jonne Poikonen**, Absolute Value Extraction and Order Statistic Filtering for a Mixed-Mode Array Image Processor
81. **Luka Milovanov**, Agile Software Development in an Academic Environment
82. **Francisco Augusto Alcaraz Garcia**, Real Options, Default Risk and Soft Applications
83. **Kai K. Kimppa**, Problems with the Justification of Intellectual Property Rights in Relation to Software and Other Digitally Distributable Media
84. **Dragoş Truşcan**, Model Driven Development of Programmable Architectures
85. **Eugen Czeizler**, The Inverse Neighborhood Problem and Applications of Welch Sets in Automata Theory

86. **Sanna Ranto**, Identifying and Locating-Dominating Codes in Binary Hamming Spaces
87. **Tuomas Hakkarainen**, On the Computation of the Class Numbers of Real Abelian Fields
88. **Elena Czeizler**, Intricacies of Word Equations
89. **Marcus Alanen**, A Metamodeling Framework for Software Engineering
90. **Filip Ginter**, Towards Information Extraction in the Biomedical Domain: Methods and Resources
91. **Jarkko Paavola**, Signature Ensembles and Receiver Structures for Oversaturated Synchronous DS-CDMA Systems
92. **Arho Virkki**, The Human Respiratory System: Modelling, Analysis and Control
93. **Olli Luoma**, Efficient Methods for Storing and Querying XML Data with Relational Databases
94. **Dubravka Ilić**, Formal Reasoning about Dependability in Model-Driven Development
95. **Kim Solin**, Abstract Algebra of Program Refinement
96. **Tomi Westerlund**, Time Aware Modelling and Analysis of Systems-on-Chip
97. **Kalle Saari**, On the Frequency and Periodicity of Infinite Words
98. **Tomi Kärki**, Similarity Relations on Words: Relational Codes and Periods
99. **Markus M. Mäkelä**, Essays on Software Product Development: A Strategic Management Viewpoint
100. **Roope Vehkalahti**, Class Field Theoretic Methods in the Design of Lattice Signal Constellations
101. **Anne-Maria Ernvall-Hytönen**, On Short Exponential Sums Involving Fourier Coefficients of Holomorphic Cusp Forms
102. **Chang Li**, Parallelism and Complexity in Gene Assembly
103. **Tapio Pahikkala**, New Kernel Functions and Learning Methods for Text and Data Mining
104. **Denis Shestakov**, Search Interfaces on the Web: Querying and Characterizing
105. **Sampo Pyysalo**, A Dependency Parsing Approach to Biomedical Text Mining
106. **Anna Sell**, Mobile Digital Calendars in Knowledge Work
107. **Dorina Marghescu**, Evaluating Multidimensional Visualization Techniques in Data Mining Tasks
108. **Tero Säntti**, A Co-Processor Approach for Efficient Java Execution in Embedded Systems
109. **Kari Salonen**, Setup Optimization in High-Mix Surface Mount PCB Assembly
110. **Pontus Boström**, Formal Design and Verification of Systems Using Domain-Specific Languages
111. **Camilla J. Hollanti**, Order-Theoretic Mehtods for Space-Time Coding: Symmetric and Asymmetric Designs
112. **Heidi Himmanen**, On Transmission System Design for Wireless Broadcasting
113. **Sébastien Lafond**, Simulation of Embedded Systems for Energy Consumption Estimation
114. **Evgeni Tsivtsivadze**, Learning Preferences with Kernel-Based Methods
115. **Petri Salmela**, On Commutation and Conjugacy of Rational Languages and the Fixed Point Method
116. **Siamak Taati**, Conservation Laws in Cellular Automata
117. **Vladimir Rogojin**, Gene Assembly in Stichotrichous Ciliates: Elementary Operations, Parallelism and Computation
118. **Alexey Dudkov**, Chip and Signature Interleaving in DS CDMA Systems
119. **Janne Savela**, Role of Selected Spectral Attributes in the Perception of Synthetic Vowels
120. **Kristian Nybom**, Low-Density Parity-Check Codes for Wireless Datacast Networks
121. **Johanna Tuominen**, Formal Power Analysis of Systems-on-Chip
122. **Teijo Lehtonen**, On Fault Tolerance Methods for Networks-on-Chip
123. **Eeva Suvitie**, On Inner Products Involving Holomorphic Cusp Forms and Maass Forms
124. **Linda Mannila**, Teaching Mathematics and Programming – New Approaches with Empirical Evaluation
125. **Hanna Suominen**, Machine Learning and Clinical Text: Supporting Health Information Flow
126. **Tuomo Saarni**, Segmental Durations of Speech
127. **Johannes Eriksson**, Tool-Supported Invariant-Based Programming

128. **Tero Jokela**, Design and Analysis of Forward Error Control Coding and Signaling for Guaranteeing QoS in Wireless Broadcast Systems
129. **Ville Lukkarila**, On Undecidable Dynamical Properties of Reversible One-Dimensional Cellular Automata
130. **Qaisar Ahmad Malik**, Combining Model-Based Testing and Stepwise Formal Development
131. **Mikko-Jussi Laakso**, Promoting Programming Learning: Engagement, Automatic Assessment with Immediate Feedback in Visualizations
132. **Riikka Vuokko**, A Practice Perspective on Organizational Implementation of Information Technology
133. **Jeanette Heidenberg**, Towards Increased Productivity and Quality in Software Development Using Agile, Lean and Collaborative Approaches
134. **Yong Liu**, Solving the Puzzle of Mobile Learning Adoption
135. **Stina Ojala**, Towards an Integrative Information Society: Studies on Individuality in Speech and Sign
136. **Matteo Brunelli**, Some Advances in Mathematical Models for Preference Relations
137. **Ville Junnila**, On Identifying and Locating-Dominating Codes
138. **Andrzej Mizera**, Methods for Construction and Analysis of Computational Models in Systems Biology. Applications to the Modelling of the Heat Shock Response and the Self-Assembly of Intermediate Filaments.
139. **Csaba Ráduly-Baka**, Algorithmic Solutions for Combinatorial Problems in Resource Management of Manufacturing Environments
140. **Jari Kyngäs**, Solving Challenging Real-World Scheduling Problems
141. **Arho Suominen**, Notes on Emerging Technologies
142. **József Mezei**, A Quantitative View on Fuzzy Numbers
143. **Marta Olszewska**, On the Impact of Rigorous Approaches on the Quality of Development
144. **Antti Airola**, Kernel-Based Ranking: Methods for Learning and Performace Estimation
145. **Aleksi Saarela**, Word Equations and Related Topics: Independence, Decidability and Characterizations
146. **Lasse Bergroth**, Kahden merkkijonon pisimmän yhteisen alijonon ongelma ja sen ratkaiseminen
147. **Thomas Canhao Xu**, Hardware/Software Co-Design for Multicore Architectures
148. **Tuomas Mäkilä**, Software Development Process Modeling – Developers Perspective to Contemporary Modeling Techniques
149. **Shahrokh Nikou**, Opening the Black-Box of IT Artifacts: Looking into Mobile Service Characteristics and Individual Perception
150. **Alessandro Buoni**, Fraud Detection in the Banking Sector: A Multi-Agent Approach
151. **Mats Neovius**, Trustworthy Context Dependency in Ubiquitous Systems
152. **Fredrik Degerlund**, Scheduling of Guarded Command Based Models
153. **Amir-Mohammad Rahmani-Sane**, Exploration and Design of Power-Efficient Networked Many-Core Systems
154. **Ville Rantala**, On Dynamic Monitoring Methods for Networks-on-Chip
155. **Mikko Pelto**, On Identifying and Locating-Dominating Codes in the Infinite King Grid
156. **Anton Tarasyuk**, Formal Development and Quantitative Verification of Dependable Systems
157. **Muhammad Mohsin Saleemi**, Towards Combining Interactive Mobile TV and Smart Spaces: Architectures, Tools and Application Development
158. **Tommi J. M. Lehtinen**, Numbers and Languages
159. **Peter Sarlin**, Mapping Financial Stability
160. **Alexander Wei Yin**, On Energy Efficient Computing Platforms
161. **Mikołaj Olszewski**, Scaling Up Stepwise Feature Introduction to Construction of Large Software Systems
162. **Maryam Kamali**, Reusable Formal Architectures for Networked Systems
163. **Zhiyuan Yao**, Visual Customer Segmentation and Behavior Analysis – A SOM-Based Approach
164. **Timo Jolivet**, Combinatorics of Pisot Substitutions
165. **Rajeev Kumar Kanth**, Analysis and Life Cycle Assessment of Printed Antennas for Sustainable Wireless Systems
166. **Khalid Latif**, Design Space Exploration for MPSoC Architectures

167. **Bo Yang**, Towards Optimal Application Mapping for Energy-Efficient Many-Core Platforms
168. **Ali Hanzala Khan**, Consistency of UML Based Designs Using Ontology Reasoners
169. **Sonja Leskinen**, m-Equine: IS Support for the Horse Industry
170. **Fareed Ahmed Jokhio**, Video Transcoding in a Distributed Cloud Computing Environment
171. **Moazzam Fareed Niazi**, A Model-Based Development and Verification Framework for Distributed System-on-Chip Architecture
172. **Mari Huova**, Combinatorics on Words: New Aspects on Avoidability, Defect Effect, Equations and Palindromes
173. **Ville Timonen**, Scalable Algorithms for Height Field Illumination
174. **Henri Korvela**, Virtual Communities – A Virtual Treasure Trove for End-User Developers
175. **Kameswar Rao Vaddina**, Thermal-Aware Networked Many-Core Systems
176. **Janne Lahtiranta**, New and Emerging Challenges of the ICT-Mediated Health and Well-Being Services
177. **Irum Rauf**, Design and Validation of Stateful Composite RESTful Web Services
178. **Jari Björne**, Biomedical Event Extraction with Machine Learning
179. **Katri Haverinen**, Natural Language Processing Resources for Finnish: Corpus Development in the General and Clinical Domains
180. **Ville Salo**, Subshifts with Simple Cellular Automata
181. **Johan Ersfolk**, Scheduling Dynamic Dataflow Graphs
182. **Hongyan Liu**, On Advancing Business Intelligence in the Electricity Retail Market
183. **Adnan Ashraf**, Cost-Efficient Virtual Machine Management: Provisioning, Admission Control, and Consolidation
184. **Muhammad Nazrul Islam**, Design and Evaluation of Web Interface Signs to Improve Web Usability: A Semiotic Framework
185. **Johannes Tuikkala**, Algorithmic Techniques in Gene Expression Processing: From Imputation to Visualization
186. **Natalia Díaz Rodríguez**, Semantic and Fuzzy Modelling for Human Behaviour Recognition in Smart Spaces. A Case Study on Ambient Assisted Living
187. **Mikko Pänkäälä**, Potential and Challenges of Analog Reconfigurable Computation in Modern and Future CMOS
188. **Sami Hyrynsalmi**, Letters from the War of Ecosystems – An Analysis of Independent Software Vendors in Mobile Application Marketplaces
189. **Seppo Pulkkinen**, Efficient Optimization Algorithms for Nonlinear Data Analysis
190. **Sami Pyöttiälä**, Optimization and Measuring Techniques for Collect-and-Place Machines in Printed Circuit Board Industry
191. **Syed Mohammad Asad Hassan Jafri**, Virtual Runtime Application Partitions for Resource Management in Massively Parallel Architectures
192. **Toni Ernvall**, On Distributed Storage Codes
193. **Yuliya Prokhorova**, Rigorous Development of Safety-Critical Systems
194. **Olli Lahdenoja**, Local Binary Patterns in Focal-Plane Processing – Analysis and Applications
195. **Annika H. Holmbom**, Visual Analytics for Behavioral and Niche Market Segmentation
196. **Sergey Ostroumov**, Agent-Based Management System for Many-Core Platforms: Rigorous Design and Efficient Implementation
197. **Espen Suenson**, How Computer Programmers Work – Understanding Software Development in Practise
198. **Tuomas Poikela**, Readout Architectures for Hybrid Pixel Detector Readout Chips
199. **Bogdan Iancu**, Quantitative Refinement of Reaction-Based Biomodels
200. **Ilkka Törmä**, Structural and Computational Existence Results for Multidimensional Subshifts
201. **Sebastian Okser**, Scalable Feature Selection Applications for Genome-Wide Association Studies of Complex Diseases
202. **Fredrik Abbors**, Model-Based Testing of Software Systems: Functionality and Performance
203. **Inna Pereverzeva**, Formal Development of Resilient Distributed Systems
204. **Mikhail Barash**, Defining Contexts in Context-Free Grammars
205. **Sepinoud Azimi**, Computational Models for and from Biology: Simple Gene Assembly and Reaction Systems
206. **Petter Sandvik**, Formal Modelling for Digital Media Distribution

207. **Jongyun Moon**, Hydrogen Sensor Application of Anodic Titanium Oxide Nanostructures
208. **Simon Holmbacka**, Energy Aware Software for Many-Core Systems
209. **Charalampos Zinoviadis**, Hierarchy and Expansiveness in Two-Dimensional Subshifts of Finite Type
210. **Mika Murtojärvi**, Efficient Algorithms for Coastal Geographic Problems
211. **Sami Mäkelä**, Cohesion Metrics for Improving Software Quality
212. **Eyal Eshet**, Examining Human-Centered Design Practice in the Mobile Apps Era
213. **Jetro Vesti**, Rich Words and Balanced Words
214. **Jarkko Peltomäki**, Privileged Words and Sturmian Words
215. **Fahimeh Farahnakian**, Energy and Performance Management of Virtual Machines: Provisioning, Placement and Consolidation
216. **Diana-Elena Gratie**, Refinement of Biomodels Using Petri Nets
217. **Harri Merisaari**, Algorithmic Analysis Techniques for Molecular Imaging
218. **Stefan Grönroos**, Efficient and Low-Cost Software Defined Radio on Commodity Hardware
219. **Noora Nieminen**, Garbling Schemes and Applications
220. **Ville Taajamaa**, O-CDIO: Engineering Education Framework with Embedded Design Thinking Methods
221. **Johannes Holvitie**, Technical Debt in Software Development – Examining Premises and Overcoming Implementation for Efficient Management
222. **Tewodros Deneke**, Proactive Management of Video Transcoding Services
223. **Kashif Javed**, Model-Driven Development and Verification of Fault Tolerant Systems
224. **Pekka Naula**, Sparse Predictive Modeling – A Cost-Effective Perspective
225. **Antti Hakkala**, On Security and Privacy for Networked Information Society – Observations and Solutions for Security Engineering and Trust Building in Advanced Societal Processes

Turku Centre *for* Computer Science

http://www. tucs.fi

tucs@abo.fi

University of Turku
Faculty of Mathematics and Natural Sciences
- Department of Information Technology
- Department of Mathematics and Statistics

Turku School of Economics
- Institute of Information Systems Science

Åbo Akademi University
Faculty of Science and Engineering
- Computer Engineering
- Computer Science

Faculty of Social Sciences, Business and Economics
- Information Systems

ISBN 978-952-12-3607-5
ISSN 1239-1883

Antti Hakkala

On Security and Privacy for Networked Information Society