University of Zurich^UZH

Mapping Boundaries: An Analytical Dive into AirTags and respective Privacy Concerns

Dominic Vogel
Zurich, Switzerland
Student ID: 20-704-474

Supervisors: Katharina O. E. Müller, Dr. Bruno Rodrigues,
Prof. Dr. Burkhard Stiller
Date of Submission: March 25, 2024

BACHELOR THESIS – Communication Systems Group, Prof. Dr. Burkhard Stiller

University of Zurich
Department of Informatics (IFI)
Binzmühlestrasse 14, CH-8050 Zürich, Switzerland

ifi

Bachelor Thesis
Communication Systems Group (CSG)
Department of Informatics (IFI)
University of Zurich
Binzmühlestrasse 14, CH-8050 Zürich, Switzerland
URL: http://www.csg.uzh.ch/

Declaration of Independence

I hereby declare that I have composed this work independently and without the use of any aids other than those declared (including generative AI such as ChatGPT). I am aware that I take full responsibility for the scientific character of the submitted text myself, even if AI aids were used and declared (after written confirmation by the supervising professor). All passages taken verbatim or in sense from published or unpublished writings are identified as such. The work has not yet been submitted in the same or similar form or in excerpts as part of another examination.

Zürich, 25 03 2024

Dominic Vogel

Signature of student

Zusammenfassung

Im April 2021 hat Apple den AirTag auf den Bluetooth Low-Energy (BLE) Tracker Markt gebracht. Der AirTag ist ein kleiner, handlicher Tracker, der für seine beispiellose Genauigkeit bei der Ortung, einfache Handhabung und erschwinglichen Preis gelobt wird. Diese Qualitäten haben ihn jedoch zu einem zweischneidigen Schwert gemacht, da er auch für Stalking missbraucht werden kann. Zum Zeitpunkt der Produktvorstellung waren die Massnahmen zur Verhinderung und Erkennung von Stalking mit AirTags für iOS-Nutzer unzureichend und für Android-Nutzer nicht vorhanden. Trotz Verbesserungen im Laufe der Zeit sieht sich das AirTag-System immer noch zahlreichen Bedrohungen gegenüber. Diese Arbeit untersucht die komplexe Landschaft des Datenschutzes, indem sie verschiedene Frameworks analysiert und deren Gemeinsamkeiten und Unterschiede bei der Definition von Datenschutzanforderungen hervorhebt. Durch die Untersuchung bestehender Forschung zum Apple AirTag, einschliesslich einem laufenden gerichtlichen Verfahren, wird eine Reihe von Bedrohungen am AirTag-System in verschiedenen Anwendungsfällen identifiziert. Diese Bedrohungen werden methodisch kategorisiert und einer DREAD-Risikobewertung unterzogen, um ihnen eine Prioritätsstufe zuzuweisen. Abschliessend präsentiert diese Arbeit ein einzigartiges Framework mit einer umfassenden Liste von Datenschutzanforderungen und einem detaillierten Katalog von Bedrohungen auf das AirTag-System. Diese Bedrohungen werden basierend auf den Datenschutzanforderungen, die sie verletzen, kategorisiert. Weiter wurde ein Prototyp eines Bedrohungs-Klassifikationsbaum für eine verbesserte Visualisierung und Verständnis erstellt.

Abstract

Apple introduced the AirTag to the Bluetooth Low-Energy (BLE) tracker market in April 2021. The AirTag is a small handheld tracker highly praised for its unparalleled tracking accuracy, ease of use, and affordability. However, these qualities have made it a double-edged sword, as it has also become exploited for stalking purposes. At the time of the release, measures to prevent and detect stalking were insufficient for iOS users and non-existent for those on Android devices. Despite improvements over time, the AirTag system still faces numerous threats. This thesis delves into the complex landscape of privacy concerns by analyzing various privacy frameworks, highlighting their commonalities and differences in defining privacy requirements. By examining existing research on the Apple AirTag, including ongoing legal proceedings, a range of AirTag-related threats across different use cases are identified. These threats are methodically categorized and subjected to a DREAD risk assessment, determining their priority levels. In conclusion, this thesis presents a unique framework featuring a comprehensive list of privacy requirements along with a detailed catalog of threats specific to the AirTag system. These threats are categorized based on the privacy requirements they violate, resulting in a prototype threat classification tree for enhanced visualization and comprehension.

Acknowledgments

First and foremost, I would like to thank my supervisor, Katharina O. E. Müller for her invaluable guidance, support, and insightful feedback throughout the completion of this thesis. Additionally, I would to express my gratitude towards Prof. Dr. Stiller and the Communication Systems Group at the University of Zurich for providing me with a highly intriguing topic for my Bachelor Thesis.

Contents

Chapter 1

Introduction

In the early 1980s, telecommunications engineer, entrepreneur, and aviator *Ed Tuck* often flew along the California coast. Frustrated by navigating small, hard-to-find airports without control towers in thick fog, he sought a solution [1], [2]. With the *global positioning systemm* (GPS) already being declassified by United States President Ronald Reagan, Tuck wanted to create a mobile, handheld, and affordable GPS device that could tell him where he was while flying. Through his venture capital firm, *Boundary Fund*, he invested and founded a company called *Magellan* in 1986. Two years later, *Magellan* introduced the first handheld GPS tracker, the *NAV 1000* [3]. It was about the size of a large calculator and weighed 850 grams. Although the initial price of $2500 per unit made it quite expensive to the average consumer, with the introduction of the *NAV 1000* to the consumer-based market, Tuck was able to spark widespread interest in GPS among the general population [1], [3].

Over the years, consumer-based, handheld trackers evolved and started making use of different technologies such as GPS, radio waves, *Wireless Fidelity* (WiFi), Bluetooth, and BLE [4], [5]. BLE has asserted itself as the prevalent choice among the most modern trackers, as it has a reduced power consumption while maintaining a similar communication range compared with other technologies like Bluetooth [6]. Initially, the first generation of BLE trackers were typically linked directly to a user's smartphone. This led to limitations in range and capabilities, as they had to stay within BLE range of each other, rendering them useless if they got lost, misplaced, or stolen [7].

Today's generation of trackers has overcome these limitations by embedding trackers into an ecosystem called *crowdsourced offline finding networks* (COFN). By pairing a tracker to a user's smartphone and then using an app as an intermediary, the tracker becomes connected to the internet. Being inside a COFN allows a tracker to report its location to the owner without being connected to the internet. This works as the tracker emits BLE advertisements to nearby devices inside the same network. In contrast, these nearby devices have an internet connection, thereby allowing them to collect and report the tracker's location, letting its owner know of its whereabouts [8].

In April 2021, Apple announced its market entry into handheld BLE tracking devices with its newest product: The AirTag [9]. As of early 2022, it is the cheapest computing device sold by Apple for a starting price of $ 29 [9], [10]. The AirTag is the size of a large coin and weighs 11 grams [9], [11]. It is embedded into the *Find My* network, Apple's implementation of its own COFN. This demonstrates how far technology has advanced since the introduction of the *NAV 1000*, in just 33 years.

The introduction of the AirTag marked a significant advancement in tracking technology; however, it also raised concerns about the potential misuse and privacy implications of such compact and affordable tracking devices. While designing the AirTag, Apple tried to mitigate the possible abuse of their technology. Yet, due to the AirTag's low cost, size, and tracking capabilities without internet connectivity, they could be surreptitiously placed in someone's handbag, car, jacket, or backpack. Such actions could lead to threats like stalking, where individuals may unknowingly be tracked by malicious actors [12].

Subsequent updates to the AirTag improved privacy measures, such as the *Item Safety Alert* (ISA), as it was designed to notify iPhone users if a tracker has been following them. The ISA feature allows users to connect to the unknown AirTag, make it emit a sound, or utilize *ultra-wideband* (UWB) technology to find it [12]. However, popular media reports have questioned the reliability of the ISA feature [13], [14]. It initially aired in iOS 14.3, yet was removed until its reappearance in iOS 14.5 in April 2021 [15]. Additionally, the limitation of ISA to Apple devices prompted the release of the *Tracker Detect* app on the Google Play Store in December 2021. This app aimed to address concerns of location tracking attacks on Android users by enabling manual scanning for BLE devices within the *Find My* network [7]. Despite these efforts, recent studies have highlighted ongoing flaws in the detection of tracking devices within the *Find My* network, underscoring the persistent security risks associated with the AirTag and similar tracking technologies [7], [16], [17].

Another critique of the feature is that it is only available for Apple devices. Therefore, in December 2021, Apple released the *Tracker Detect* app on the Google Play Store. This was Apple's "*answer to numerous concerns about location tracking attacks on Android users by AirTags*" as described by [7]. It finally allowed Android users to manually scan for BLE devices inside the *Find My* network without requiring an Apple device. However, this approach was flawed, as Android users have to actively scan and be suspicious of devices inside the *Find My* network, which does not apply to real-life scenarios [16]. Recent studies have shown that the detection of tracking devices inside the *Find My* network still has flaws, and the AirTag, in general, poses several security risks [7], [16], [17].

1.1 Motivation

Apple created the AirTag to allow for "*a supereasy way to keep track of your stuff*", as stated on their official website [9]. It has received widespread adoption for different uses in tracking personal items. Next to typical items like keys or wallets, users even started keeping tabs on their luggage while traveling [18]. In the United States, the New York City municipal government even began using AirTags to combat the rising number of car thefts by temporarily handing out AirTags for free [19]. As the New York Police Department Chief of Department Jeffrey B. Maddrey concluded in a tweet: "*The 21st century calls for 21st century policing. AirTags in your car will help us recover your vehicle if it's stolen. [...] Help us help you, get an AirTag*" [19]. Positive reports on the usefulness and reliability of AirTags helped Apple sell over $1 billion worth of AirTags in less than two years from April 2021 to December 2022 [20].

While there is a lot of positive word-of-mouth on the Airtag, Apple has also received some backlash, as the small tracker is being used for malicious reasons. In June 2022, a U.S. State of Indiana woman tracked her boyfriend using an AirTag. She suspected he was cheating on her and killed him later on after finding him talking to another woman [21]. In a statement in February 2022, Apple emphasized: "*[The] AirTag was designed to help people locate their personal belongings, not to track people or another person's property, and we condemn in the strongest possible terms any malicious use of our products*" [22].

Nonetheless, stalking with AirTags has not stopped since. In another example, two women sued Apple for aiding their former partners to help track them down. One of the plaintiffs, Lauren Hughes, claimed that her former boyfriend learned where she had moved to avoid him by placing an AirTag inside her car's wheel well. The other plaintiff, who preferred to remain anonymous, said her estranged husband could track her whereabouts by placing an AirTag inside their child's backpack [23]. Several privacy measures, like the ISA feature or the triggering of a sound on the Airtag, were useless. For Hughes, she was notified that an unknown Apple AirTag was following her only after returning to the hotel where she had moved. Additionally, she could only hear it once after trying to find it by making it play a sound. In their lawsuit, they accused Apple of negligence, design defects, and privacy violations, among other allegations [24].

Two years after the release of the AirTag in 2021, Apple is still developing viable solutions to mitigate the risk of privacy breaches targeting potential victims. In May 2023, they announced a collaboration with one of their largest competitors, Google. Together, they are developing a specification to help detect unauthorized tracking and create alerts across iOS and Android platforms. They planned on releasing this by the end of 2023, yet its launch was delayed for several reasons [25].

The motivation for this thesis stems from the pressing and evolving nature of privacy threats associated with BLE trackers. In an increasingly interconnected world where Bluetooth trackers have become ubiquitous, understanding their functionality, evaluating their privacy measures, and uncovering potential risks have gained paramount importance. Investigating and addressing these concerns is imperative to safeguard individuals' privacy and security.

1.1.1 Thesis Goals

This thesis investigates the potential risks associated with the AirTag system and provides new insights into them. The AirTag is a relatively new and contentious product, so extensive research has been conducted to understand its inner workings, such as the protective measures implemented. This study aims to gather AirTag-related threats from various sources, such as scientific literature, news articles, and the aforementioned ongoing lawsuit, to further classify them systematically by analyzing the impact of Apple's privacy measures. Furthermore, the following *research questions* (RQ) are explored and answered:

- RQ1: What are the differences and commonalities between privacy requirements among various legal frameworks and industry standards?
- RQ2: To what extent do Apple's privacy measures in the AirTag align with the broad spectrum of privacy requirements and standards across different use cases, including ones with malicious intent?
- RQ3: What are possible solutions to improve the detection mechanisms of BLE trackers to improve user privacy?

This paper provides insight into the AirTag system's privacy landscape through an in-depth analysis of these research questions. To the point of writing this thesis, no papers put AirTag threats into the perspective of the privacy requirements they breach. There has been extensive research on identifying attacks on the AirTag system, yet none categorize the identified threats by privacy requirement and severity to the system. Therefore, by collecting several AirTag threats and analyzing them quantitatively and qualitatively, a comprehensive evaluation will be delivered to compare the severity of different threats and place them in a larger context within the privacy landscape.

1.2 Thesis Outline

This thesis begins with a background section in Chapter 2, analyzing different location-tracking technologies, techniques, and devices. The background section concludes with an overview of privacy measures and concerns by analyzing privacy frameworks. Furthermore, Chapter 3 presents related work on COFNs, tracker detection methods, and privacy in the *Internet of Things* (IoT). In Chapter 4, the designs of various experiments and mappings are contained. It begins with comparing the different privacy requirements discovered by analyzing the privacy landscape. Moreover, several use cases are created and explained, followed by the design for *Received Signal Strength Indication* (RSSI) value experiments.

Chapter 5 presents the evaluation and results found in the experiments. It begins with the analysis of the privacy use cases and requirements. It continues with an analysis of the *Asia-Pacific Economic Cooperation* (APEC) privacy requirements, which are then mapped to the previous privacy requirements mapping. Further, the lawsuit *Hughes et* *al. versus Apple* is analyzed. The privacy threat modeling and risk assessment section follows this up. The privacy threats are then categorized and evaluated in a threat classification tree. Chapter 6 discusses the previously discovered findings, and lastly, Chapter 7 summarizes conclusive remarks concerning the aforementioned RQs and provides an outlook on future work.

Chapter 2

Background

This chapter introduces the necessary background knowledge on different relevant topics related to BLE Trackers. It starts off with an overview of the BLE and UWB technologies. Further, the topic of COFN is explained by using the *Find My* network as an example. It continues by analyzing different state-of-the-art trackers. Lastly, it demonstrates privacy measures taken within the AirTag system and presents an overview of the current privacy landscape in which BLE trackers are situated.

2.1 Location Tracking Technologies

Tracking a user's location using a handheld tracking device has been revolutionized over the past decades. Early innovations featured GPS technology, like the *NAV 1000* [3]. While modern trackers employ technologies like Bluetooth 5.0, GPS, or Long Range Wide Area Network (LoRaWAN) [26], BLE has become the prevalent choice among many, as it has a reduced power consumption while maintaining a similar communication range compared to other technologies [6]. The following section analyzes BLE technology by giving an in-depth overview and showing its use in today's trackers. Additionally, UWB technology is discussed, as it is used within the Apple AirTag, for showing the exact distance and direction of one nearby [10].

2.1.1 Bluetooth Low-Energy

BLE was introduced as part of the Bluetooth Core Specification version 4.0. It was developed as an alternative to the traditional Bluetooth Basic Rate/Enhanced Data Rate (BR/EDR) technology, with different capabilities and qualities. BLE was designed to cater to the requirements of new generations of products, making it a perfect alternative to Bluetooth BR/EDR. Devices that use small, coin-sized batteries were envisioned for BLE as its original design goal was to be highly efficient in its use of power [27].

Topologies

Conventional *Bluetooth BR/EDR* and BLE support Point-to-Point Topology, a one-to-one communication style. Table 2.1, adapted from [28], displays the differences from the Bluetooth technologies using this topology. Here, it becomes evident that through a smaller data rate, a smaller max payload size, and an optimization towards short burst data transmission, BLE offers a different approach than *Bluetooth BR/EDR*.

Properties	**BLE**	**Bluetooth BR/EDR**
Optimization	Short burst data transmission	Continuous data streaming
Max. # of Connections per device	Unlimited	7
Data rate	125 Kb/s to 2 Mb/s	1 Mb/s to 3 Mb/s
Max payload size	251 bytes	1021 bytes

Table 2.1: Point-to-Point Topology Comparison for *Bluetooth BR/EDR* and BLE [28]

Additionally, BLE supports two other topologies: Broadcast and Mesh. Advertising Broadcast is used for one-to-many device communication, providing a connectionless mode. Advertising packets can be received by any BLE-capable scanning device in the advertising device's range, meaning advertising can simultaneously send data to many devices. This process is called *passive scanning*. BLE advertising only supports sending data from the advertising device to the scanning devices. However, scanning devices can reply to advertising packets to request further information. If this happens, the process is defined as *active scanning*. The advertising packets are 37 bytes long, with a six-byte header and a 31-byte long payload [27]. As this thesis focuses on BLE advertising, Mesh, a many to many topology, will not be further analyzed.

Physical Layer

BLE operates in the 2.4 gigahertz (GHz) Industrial, Scientific, and Medical band, which is also used by conventional *Bluetooth BR/EDR*. This 2.4 GHz band is divided into 40 separate channels, each spaced 2 MHz apart. This separation helps mitigate interferences and provides multiple options for communication. For the transmission and decoding of data, BLE uses *Gaussian Frequency Shift Keying* (GFSK) modulation [27].

Link Layer

The *Link Layer* defines several types of packets transmitted over air and an associated air interface protocol. Its operation is subject to a state machine. A state defines how the link layer may operate [27]. Subsequently, the topics of *Packets*, *States*, *Channels*, and *Addresses* will be analyzed.

- **Packets**: Table 2.1 shows that BLE is optimized for short burst data transmission. This means the data is divided into individual packets instead of continuous data transmission. Figure 2.1 shows the different fields each package includes.

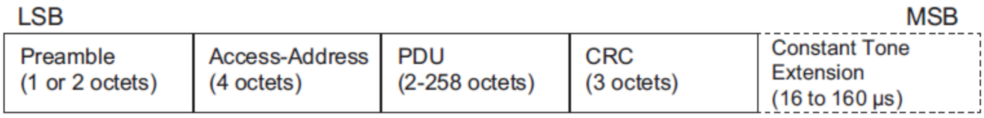

Figure 2.1: Link Layer packet format [27]

The *Preamble* helps a receiver to synchronize the signal's frequency. The *Access-Address* is a 32-bit address unique for each connection. It helps to distinguish between background noise, relevant signals, or advertisements. In the case of advertisement packets, the *Access-Address* is set to *0x08e89bed6*. The *Cyclic Redundancy Check* (CRC) helps check for errors in one or more bits to ensure data integrity [27]. Lastly, the *Protocol Data Unit* (PDU) contains the payload data used for the transmission. Figure 2.2 displays an example of a PDU of a BLE advertisement. It shows the header on the *most significant bit* (MSB) side and the payload on the *least significant bit* (LSB) side. The payload contains the *Advertising Address* (AdvA), followed by *Advertising Data* (AD) structures. The *length* field contains the length of the AD structure. *Type* specifies the nature of the data and the *data* field contains the AD itself [6].

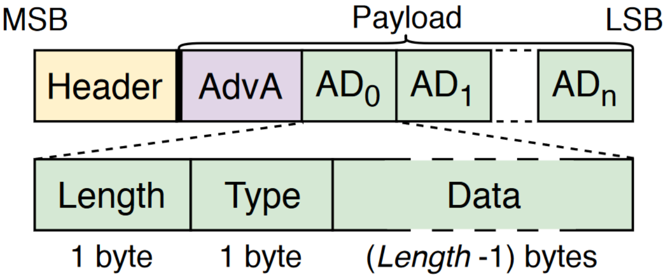

Figure 2.2: Representation of a BLE advertising PDU's header and payload [6]

- **States**: BLE devices have various states they can be in during the communication process. The different states are explained in Table 2.2.

  For this thesis, the *Advertising* and the *Scanning* states are of particular importance. A device listening for advertising packets from other devices is in the *Scanning* state. Another device emits these advertising packets in the *Advertising* state. If the other device scans the advertising packets, these two devices change their state to the *Connection* state, as they are now connected [27]. The relation of a device in an *Advertising* state and another in the *Scanning* state can also be exemplified by assigning two main roles, *Peripheral* and *Central*, following a client-server model. A *Peripheral* device emits advertisement packets and, if connectable, accepts incoming connection requests from *Central*. A *Central* device, on the other hand, scans for advertisement packets and, when applicable, initiates a connection with the *Peripheral* [27].

State	**Description**
Standby	Device neither transmits nor receives packets
Initiating	Responding to advertising packets from a particular device to request a connection
Advertising	Transmits advertising packets and potentially processes packets sent in response to advertising packets by other devices
Connection	In a connection with another device
Scanning	Listening for advertising packets from other devices
Isochronous Broadcast	Broadcasts isochronous data packets
Synchronization	Listens for periodic advertising belonging to a specific advertising train transmitted by a particular device

Table 2.2: Table of Link Layer states [27]

- **Channels**: BLE divides a 2.4 GHz frequency band into 40 separate channels. The link layer controls the use of these channels. Advertisements are broadcasted using the channels 37, 38, and 39. These are known as the *primary advertising channels* [27]. The other channels, channels 0 to 36, are used for data exchanges between two connected devices [6].

- **Addresses**: BLE supports three random types of random addresses in addition to the globally unique *media access control* (MAC) address. The *Public Device Address* is uniquely allocated to a device from its manufacturer by following the IEEE specifications of MAC addresses. This address is more commonly known as the MAC address. The *Random Static device address* is a randomly generated device address that can be renewed after each power cycle. During the use of the device, it will not change. The *Random Non-resolvable device address* is a randomly generated address that can be renewed anytime. Lastly, the *Random Resolvable device address* is an address composed of a 22-bit random number called *prand* and a 24-bit hash produced by the hashing of *prand* with a 128-bit secret Identity Resolution key [6]. The Bluetooth Core Specification recommends to renew the *Random Non-resolvable* and *Random Resolvable* addresses every 15 minutes [29].

To measure the power level of a Bluetooth signal, it is common to use the RSSI value. On a logarithmic scale, RSSI values are negative and measured in decibels referenced to one milliwatt. A lower number indicates a weaker signal, which leads to the conclusion that a BLE-emitting device is further away. Although values vary, values in the range of -20 to -30 indicate a device is close, and the value -120 indicates that the device is near the detection limit and thus quite far away [30]. RSSI only implies the relative distance from a receiving device to an emitting one. It is impossible to accurately determine the range between these two devices based on RSSI values. However, with enough data, a trend between RSSI values and distance can be determined [31].

2.1.2 Ultra-Wideband

UWB, in general, is a term for radio communication using a bandwidth of at least 500 MHz [32]. It is not a new technology. In fact, UWB systems are the oldest form of radio communication, with early adoptions dating back to the late 19th century [33]. Recently, *Impule Radio UWB* (IR-UWB), a specific implementation of UWB technology, has become a focus in research, as there are some core benefits IR-UWB provides. [32] describes IR-UWB as a technique, that "*uses radio frequency pulses with a very short time-duration (nano- or picoseconds), resulting in a large bandwidth*". They name three core benefits of this technique: The first is that IR-UWB supports a high channel capacity due to the high bandwidth, which enables low transmission power needed to avoid narrowband interference with other wireless technologies. The second benefit is the short time duration of the pulses [32]. This makes IR-UWB more robust to multipath effects, and the spatial diversity can be exploited to improve the localization accuracy [34]. The third benefit is that precise timing is achieved due to the high temporal solution. IR-UWB has a steep rising edge, allowing the receiving device to accurately determine the arrival time of an incoming signal, leading to a centimeter-level accurate ranging [32]. Since IR-UWB is a subtype of UWB, subsequent references toward IR-UWB in this thesis will be referred to as UWB to increase cohesiveness.

UWB technology, however, also has some downsides connected to it. Its low transmission power, which is required to avoid narrowband interference, limits UWB technology to relatively short distances. Additionally, comparing it with narrowband, the high bandwidth causes UWB pulses to be severely distorted, which can limit the performance of UWB receivers [32].

In the early 2020s, UWB technology gained a lot of support with the introduction of the UWB 802.15.4z amendment and a publication [33] from the Fine Ranging (FiRa) Consortium. The FiRa Consortium was founded in August 2019 by NXP, Samsung, Bosch, and HID Global [35]. They dedicate themselves to transforming how humans interact with their environment by enabling precise location awareness with their devices. They refer to UWB as the "*most available technology for delivering accurate ranging and positioning in challenging real-world environments, allowing devices to add real-time spatial context and enabling new experiences*" [36]. It aims to support UWB's broad market acceptance by offering effective industry support [36] and by solving ecosystem and interoperability challenges that occur within UWB applications [32]. The FiRa Consortium achieves this by referring UWB to its many members, some of which are market leaders in the respective fields of consumer technology, semiconductors, networking, and secure access [36].

In late 2019, Apple decided to develop its own UWB chip called the *U1* and has since added it to all iPhones 11 and newer. Samsung introduced UWB in August 2020 within its Galaxy Note 20 Ultra and Z Fold 2 and has since incorporated it into its newer devices. Apple has since extended the list of devices containing the UWB U1 chip by embedding it into devices like the Apple Watch and AirTag. Samsung has also added it to its tracker called the Galaxy SmartTag+ [35]. With the increasingly widespread adoption of UWB in different segments, it has been predicted that UWB will become increasingly embedded into our daily lives. In a study by *ABI Research*, they expect that in 2025, there will be over one billion annual device shipments of UWB technology [35]. This is depicted by the chart in Figure 2.3.

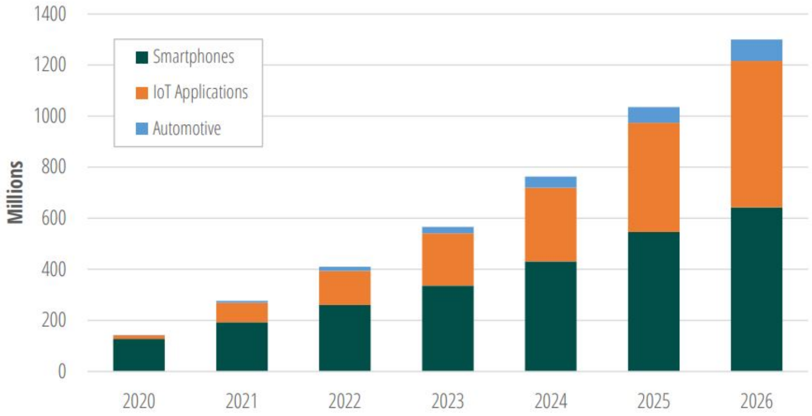

Figure 2.3: UWB enabled device shipments by market segment [35]

The automotive segment relates to a "*desire to solve the automotive industry's challenge of providing a secure, interoperable, hands-free, digital key experience*" [35]. This has sparked the interest in providing a solution using UWB technology [35]. This thesis revolves around trackers using UWB technology; this will not be covered extensively.

UWB utilizes *Time of Flight* (ToF) to measure the distance between a UWB emitting device and a receiver. ToF measures the time the signal travels between the transmitter and the receiver. Multiplying this time with the speed of light yields the distance between the two devices with high accuracy. This means that compared to BLE, which uses RSSI values to approximate the distance between an emitting device and a receiving one, UWB has a higher distance calculating accuracy with a high degree of certainty [37].

2.1.3 Interoperability between BLE and UWB

As UWB is being adapted to target secure fine-ranging applications specifically, it is mainly used complementary to many existing tracking technologies. The initial discovery is performed using alternative wireless communication technologies, like WiFi, BLE, and *Near Field Communication* (NFC). Although the combination of UWB with these technologies has its unique benefits and drawbacks, BLE has been predominantly chosen to be integrated alongside UWB. BLE is already a widely adopted technology embedded with many devices that UWB targets, like smartphones. As mentioned in Section 2.1.1, like UWB, BLE is a low-power technology. Therefore, by selectively activating UWB when necessary, the combination of UWB and BLE enables devices to achieve an extended battery life [35].

2.2 Tracking Techniques

The first generation of BLE trackers were typically linked directly to a user's device. Their functionality was limited in terms of range and capabilities, as the user's device had to stay within the BLE range of the tracker. These limitations have been conquered using COFNs [7]. The following sections analyze the idea and functionality of such a network and investigate Apple's implementation of this.

2.2.1 Crowdsourced Offline Finding Networks

Today's trackers are usually found inside an ecosystem called COFN which contains four different entities [38]:

- **Lost Device**: Either a tracking device without complicated built-in functions of communication and positioning, or a rich[1] device without online access.
- **Finder Device**: A group of users that are in BLE range of the lost device. These are rich devices with built-in functions for communication and positioning. They are connected to the internet and act as volunteers to help report the location of the offline lost device.
- **Cloud Server**: A remote server that provides storage service of reported locations. The location data sent by the finder devices is encrypted to protect the privacy of the finder device.
- **Owner Client**: A device with Internet Access that can query the location data of the lost device on the cloud server and decrypt it.

Every tracking device inside a COFN has an owner. The owner can perform privileged operations on it, like making it play a sound or flash, depending on the type of tracker. Ownership is established by registering and pairing a tracking device to the user's rich device. This is usually done using an app provided by the tracking device manufacturer [8]. Figure 2.4 depicts an overview of a COFN. The lost tracking device advertises BLE packets to nearby finder devices, which report encrypted location data to a cloud server. The owner of the lost device can access the decrypted location data by querying the cloud server using the owner client [38].

Industrial pioneers like Tile, Apple, and Samsung launched their proprietary COFN. Apple has the *Find My* app [39], Samsung has *Find My Mobile*, which later on they renamed to *SmartThings Find* [40], and Tile has its Tile app [5] ecosystem. In this thesis, Apple's *Find My* network will be analyzed extensively, as it primarily revolves around the Apple AirTag and its system.

[1] Although not defined by [38], it can be concluded that rich devices are devices with a higher computational power compared to the tracking devices, which do not have any complicated built-in functions of communication and positioning. Additionally, rich devices have online access, which the tracking devices do not [38].

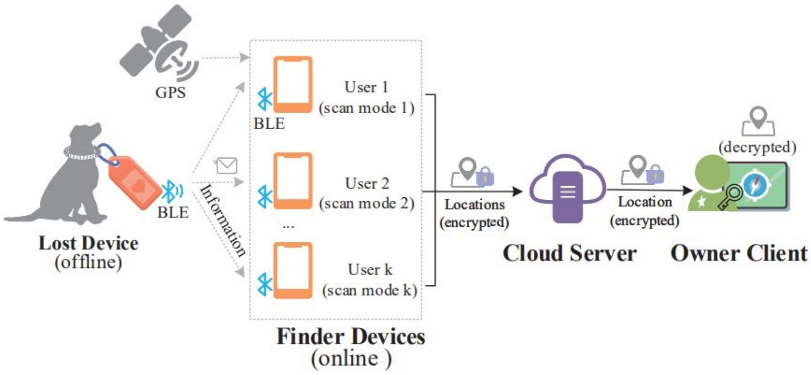

Figure 2.4: A *Crowdsourced Offline Finding Network* [38]

2.2.2 Apple's Find My Network

With over a billion active iPhones inside Apple's *Find My* network [41], Apple has a massive edge over Tile, Apple's biggest BLE tracking competitor, which only has 35 million devices inside its network [42]. A larger density of finder devices inside such an ecosystem leverages network effects since the likelihood that a bystander with a compatible device can report the location of a lost one increases with each added device inside the ecosystem [12]. At its core, Apple's *Find My* network functions like the exemplary COFN displayed in Figure 2.4. However, there are some key Apple-specific differences, which will be analyzed in the following section.

AirTag and Find My Accessories

With a COFN as big as Apple's *Find My* network, Apple elected to open up the network to allow third-party manufacturers to integrate their accessories inside of it [43]. Next to Apple's AirTag, which will be further analyzed in Section 2.3, other trackers by third-party manufacturers have been introduced into the *Find My* ecosystem. Examples include the *One Spot* by *Chipolo*, the *Sky Tag* by *4 Smarts*, the *Keyfinder* by *Atuvos*, and the *BT Tag 10 WT* by *Nedis* [26]. To be eligible for inclusion in the *Find My* network, the third-party manufacturer must adhere to all of Apple's privacy protections of the network and pay a membership fee. Approved accessories then receive a *Works with Apple Find My* badge to verify its compatibility with the *Find My* network [43].

Pairing of AirTags

Upon initial pairing, an AirTag connects to an iCloud account [10]. iCloud is a term containing all Apple services handling online data storage and synchronization via Apple's servers [44]. This allows the account owner to view the location of its devices connected to it on any device eligible for running the *Find My* app [45]. Generally, the AirTag will remain paired to the iCloud account unless the owner actively removes it. Not even hard resetting the AirTag[2] will unbind it from the paired account [46]. This mechanism prevents the stealing of AirTags, as they are rendered useless and can not be paired to a new iCloud account while being paired with an existing one [10].

States

Upon pairing, the AirTag will always be in one of three states [12]:

- *Connected*: After the initial pairing, it starts in this state. As long as the owner device with which the AirTag was paired is within BLE range of the AirTag, it will stay in this state.
- *Nearby*: If the AirTag loses connection to the connected owner device, it will transition to the *Nearby* state. While in the *Connect* or *Nearby* states, the AirTag broadcasts a short advertising message with a public key derived from the *Secret Key Nearby* and the *Master Public Key.* This key and the public MAC address used in the broadcasting message are rotated every 15 minutes to prevent tracking by using a static identifier.
- *Lost*: If the AirTag is in the *Nearby* state for longer than 15 minutes, it transitions into the *Lost* state. An owner of an AirTag can also manually set the tracker to *Lost* mode by activating it inside the *Find My* application [38]. At this point, the AirTag becomes active in the *Find My* network, emitting *Lost* BLE advertisements. Any finder iPhone will report its location to Apple's servers given they are within BLE proximity. These *Lost* BLE advertisements are emitted every two seconds. If the AirTag is reunited with its owner by getting in BLE proximity of its device, it transitions back to the *Connected* state [47].

BLE Advertisement Format

The size of a standard BLE advertising message is extremely limited, with a maximum size of 37 bytes, six of which are reserved for the MAC address itself. Therefore, the structure of a *Lost* BLE advertisement must be very compact. The entire public key has a length of 28 bytes. To fit this into the BLE advertisement, the first six bytes of the public key are stored inside the MAC address field. The remaining 22 bytes are fitted in the manufacturer data field. This is illustrated in Figure 2.5. The first two bits of the MAC address field are set to `0b11` to identify the MAC address as a random static address. This results in the first two bits of the MAC address being stored in the 29th byte of the BLE payload [47].

[2] A hard reset can be achieved by manually removing and replacing the battery of the AirTag five times [46]

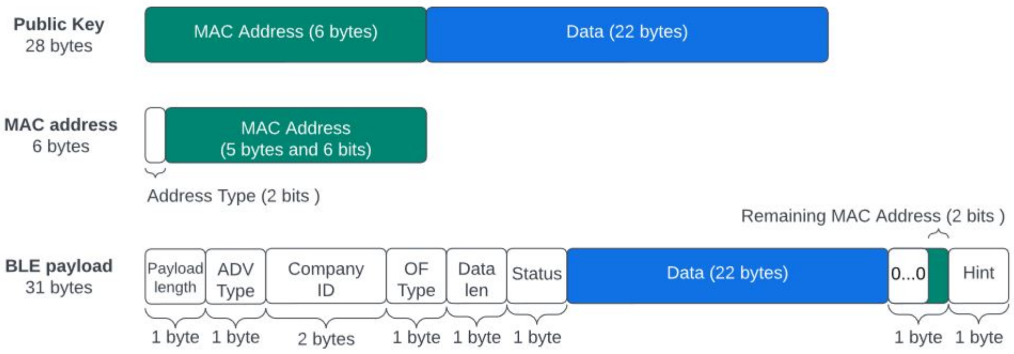

Figure 2.5: An AirTags BLE advertisement [47]

The BLE advertisement payload depicted in Figure 2.5 includes the following fields [47]:

- **Payload Length**: Refers to the payload length of 30 bytes.
- **Advetisement type** (*ADV type* in Figure 2.5): For AirTags the data type is `0xFF`. This is manufacturer-specific.
- **Company ID**: The Company ID is a unique company identifier the Bluetooth Special Interest Group assigns. Apple's Company ID is `0x004C`. A complete Company ID list can be found at [48].
- **OF type**: The OF type indicates the service offline devices request. Examples are `0x12` for the *Find My* services or `0x07` when the AirTag is unpaired.
- **Data length**: The Data Length refers to the length of the data in bytes.
- **Status Code**: The Status Code contains the device type and battery level. The device type is divided into the categories *Apple Device*, *Find My device*, AirTag, or AirPod. An *Apple device* is considered any device by Apple with a screen, such as iPhones, MacBooks, and iPads. A *Find My Device* is any third-party *Find My* compatible device, like the Chipolo Spot ONE. The battery status is divided into four categories: Full, medium, low, or critically low. A *Status Code* with the value `0x10` represents a fully charged AirTag, `0x50` is an example of an AirTag with a medium battery level.
- **Hint byte**: The Hint byte changes every 15 minutes, yet is not part of the public key. [12] mentions that the purpose of the Hint byte is not given and that setting it to arbitrary values does not affect the protocol.

Location Encryption Methodology

Upon pairing an AirTag with an owner's device, a private-public key pair and a random secret are initialized using an elliptic curve P-224. These initial keys are referred to as the master beacon keys. By beginning with the private-public key pair, an infinite number of rotating key pairs can be created using a key derivation function and the random secret. If a tracker loses connectivity with its owner device, it emits the current public key using BLE advertisements. Finder devices inside BLE proximity of the tracker extract the public key, generate a temporary private-public key pair, and perform a one-sided key exchange using Elliptic-curve Diffie-Hellman. The resulting shared secret is then used to encrypt the tracker's geolocation coordinates. The encrypted location and the finder's public key are uploaded to Apple's servers. The owner of the lost tracker can access the location reports using the *Find My* application. To decrypt the location data, the owner device performs the other side of the key exchange by using the temporary public key from the finder device and the private key from the tracker. This results in the same shared secret and the decryption of the data, revealing the tracker's most recent location in the *Find My* application [7].

iCloud plays a big part in the encryption methodology, as it stores and synchronizes the master beacon keys of all devices using the same iCloud account [7]. Any device inside the same iCloud account can access the iCloud keychain, decrypt the master beacon keys, and generate the same private-public keypair that was emitted in the BLE advertisements. Apple promises not to access a user's master beacon keys, and there has not yet been a case in which they have violated this promise. In conclusion, Apple cannot decrypt the location reports of a lost device gathered by finder devices [44].

2.3 Tracking Devices

Over the years, trackers have gotten smaller, lighter, and more affordable. Today's BLE trackers are the size of a coin, weigh about 10 grams, and cost around 30 Swiss francs [26]. Despite being so small, their batteries can last over a year since they use technologies developed for low battery usage. The following section analyzes the existing state-of-the-art handheld BLE trackers' functionality and specifications.

2.3.1 Apple AirTag

The Apple AirTag was released in April 2021 [9]. Designed as a tracking device to keep track of personal items, the AirTag is Apple's cheapest computing device, sold for as little as $ 29 [9], [10]. With a diameter of 31.9 mm and a weight of 11 grams, the Apple AirTag resembles the size and weight of a large coin. The AirTag runs on a replaceable CR2032 coin cell battery. Apple claims that an AirTag can run for over a year using the same battery, although it may vary depending on usage, environmental conditions, and replacement battery manufacturer, among other factors. A built-in speaker allows the AirTag to play a sound, facilitating the search if it is nearby [9].

The AirTag's hardware contains an *nRF52832* chip, which supports wireless connectivity with BLE and NFC. While BLE is needed for its main tracking technology, the *Find My* network and communication with the owner's iPhone, NFC creates a link to contact the AirTag's owner and enables a more simplified pairing process. *NFC tap* is a functionality where a finder of an Apple AirTag can tap it with the finder device and, through NFC, receive the contact information of the owner of the AirTag [9]. A separate chip, called *U1*, handles UWB. This is used either for fine ranging, guiding users to their AirTag similarly like a compass showing direction [10] and distance while extending the range of the AirTag [45]. BLE protocol has a range of up to 100 meters, however, with UWB this range is elevated [45]. However, to use the UWB technology inside the AirTag, the owner device must also have a UWB chip inside. As mentioned in Section 2.1.2, a UWB chip has been installed in all iPhones 11 or newer.

2.3.2 Chipolo One Spot

In 2013, Chipolo was founded as a *Kickstarter* project, with the main goal of bringing one of the first BLE trackers into the market. Their original Chipolo tracker was designed to be colorful, slim, and tight, being one of the thinnest tracker gadgets to date [49]. Over the years, Chipolo developed different trackers, and today, they offer three different product lines: The *ONE & Card*, the *ONE Point & CARD Point*, and the *ONE Spot & CARD Spot*. These product lines vary mainly in the ecosystem in which they are allocated. The *ONE & Card* uses the Chipolo app, Chipolo's own COFN. The *ONE Point & CARD Point* are finders embedded into Google's *Find My Device* ecosystem. And the *One Spot & CARD Spot* leverages Apple's *Find My* network [4]. The Chipolo One Spot was one of the first third-party devices verified and added to the *Find My* network [43]. Chipolo is aware of the sheer size difference between the ecosystems. Although these are only estimates by Chipolo themselves, the OF Network size of the Chipolo app is around 5 million. Competitors, like Apple's *Find My* network and Google's *Find My Device* network, each accumulate hundreds of millions of devices or more. Knowing this, they recommend the *ONE & Card* for finding misplaced items and left-behind alerts, as users can customize alerts, sounds, and ringtones. The *One Spot & CARD Spot* and the *ONE Point & CARD Point* are recommended for the global locating of items [4].

The Chipolo One Spot is, in many aspects, very similar to the AirTag. It also runs on a CR2032 coin cell replaceable battery, which lasts up to a year, has a built-in speaker, is priced at around 34 euros, and has the size and weight of a coin. They claim that the speaker can play a sound up to 120 dB. In contrast to the AirTag, the One Spot only has BLE, its main tracking technology, and its range is limited to 60 meters [4].

2.3.3 Other State-of-the-Art Trackers

Many handheld trackers in the market have capabilities and specifications similar to those previously mentioned. However, there are trackers with technologies that not only stand out but also offer a competitive edge in certain aspects over most trackers. Some of these will be discussed briefly in the following sections.

Samsung SmartTag, SmartTag+ and SmartTag 2

Samsung announced their first location trackers in January 2021 with the SmartTag and SmartTag+ at their Samsung Galaxy Unpacked event [50]. The SmartTag was released in late January 2021, with BLE tracking capabilities inside Samsung's *SmartThings Find* network. Following the early success of the SmartTag, in April 2021, Samsung followed up with the release of the SmartTag+. The SmartTag+ added new capabilities by incorporating UWB technology. With the help of UWB, pinpointing a location can be simplified, guiding the owner to the SmartTag+ with spatial accuracy and directional capabilities. Another capability added with the SmartTag+ is the *AR Finder* technology, which combines UWB with augmented reality technology to visually lead a user to its SmartTag+ using the smartphone's camera. However, for an owner to use UWB technology, it must have a UWB-compatible smartphone. Next to their tracking abilities, Samsung's SmartTag series also acts as a remote control with a programmable button that can be used to control SmartThings-compatible smart-home products [40].

In October 2023, Samsung released its third version of the SmartTag, the SmartTag 2. With the new *Power Saving Mode*, the battery is longer-lasting and lasts up to 700 days, a 100% increase compared to the older models. Next to this, there are many small upgrades to point out, like an improved User-Interface design inside the *SmartThings* app, or a *Lost Mode*, allowing finder devices to tap the SmartTag 2 with their device, and through NFC they receive the contact information of the owner [51]. As of May 2023, Samsung's *SmartThings Find* network accumulated over 300 million devices inside of its ecosystem [52]. Although its size is relatively small to Apple's >1 billion device network [41], Samsung achieved rapid growth as in the time-frame between July 2022 and May 2023, an additional 100 million devices were registered, signifying a 150% expansion in just 10 months [52].

Pebblebee clip

The Pebblebee Clip is their fourth-generation item tracker, which uses Bluetooth 4.0 or newer and has a connectivity range of 150 meters. Its main benefits are that it is rechargeable with a USB-C port, has a battery that lasts up to six months on a single charge, and supports visual searching of the tracker with a bright LED built inside of it. An owner can either use it within Apple's *Find My* network if the owner device is an iPhone or use it within Pebblebee's for COFN Android devices [53].

Jiobit Smart Tag

The Jiobit Smart Tag is quite different than the previously analyzed trackers as it was designed to keep track of pets, children, seniors, or adults. It uses network technologies like LTE 5G Cellular, next-gen GPS, WiFi, and BLE to ensure the most accurate, real-time location tracking. For this, a subscription data plan costs 8$ per month on a yearly prepaid schedule. Compared to the other trackers analyzed in previous parts, the Jiobit Smart Tag's price is set relatively high at 130$. But with this price tag, a lot of additional functionalities are guaranteed. *Geo-fencing* allows a user to set up a virtual perimeter for the Smart Tag, and if it leaves that perimeter, the user gets notified in a matter of minutes. In the Jiobit Smart Tag, the *Geo-Fencing* functionality is used within *Trusted Places*. *Trusted Places* lets users add virtual perimeters around specified areas known and trusted to the user. The user gets notified if the Jiobit Smart Tag arrives in said area. Figure 2.6 depicts an example of the notification and the perimeter. A built-in *SOS Mode* enables location-aware 911 emergency dispatch, which the user can activate remotely through its device. This triggers a *Care Team* to contact a professional dispatcher, handing over all emergency details with the tag's current location. Lastly, *Timeline* lets a user see a 30-day location history, showing all the movements in that time frame. The Jiobit Smart Tag offers additional features, but these are the ones that stand out.

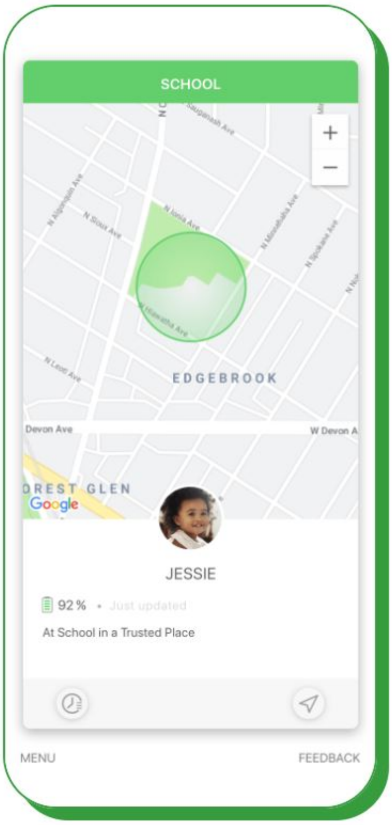

At School in a Trusted Place

Figure 2.6: The *Trusted Places* functionality using *Geo-Fencing* [54]

Summary

Table 2.3 compares various previously analyzed trackers and includes the Tile Pro tracker. It analyzes these trackers based on different aspects. It shows that there is no absolute *best* tracker. Different trackers cater to different needs and come at different prices. Especially with newer models, more modern technologies and functionalities are employed, yet it usually also depends on the user device to make use of these (e.g., *AR Finding* requires a device with a UWB chip installed).

Brand	Apple [9]	Chipolo [4]	Samsung [51]	Pebblebee [53]	Jiobit by Life360 [54]	Tile by Life360 [5]
Product Name	AirTag	One Spot	SmartTag 2	Clip	Jiobit Smart Tag	Pro
Release Date	April 2021	August 2021	October 2023	June 2023	2018	June 2022
Technology	BLE, UWB, NFC	BLE	BLE, UWB, NFC	BLE	LTE 5G Cellular, GPS, WiFi, BLE	BLE
Ecosystem	Find My	Find My	SmartThings Find	Find My (iOS), Pebblebee App (Android and iOS)	N/A	Tile App (iOS and Android)
Size of Ecosystem	>1 billion devices [41]	>1 billion devices [41]	>300 million devices [52]	>1 billion devices [41]	N/A	No data available
Battery Life	More than a year, replaceable battery	Up to a year, replaceable battery	Power Saving Mode: 700 days, Normal Mode: 500 days	Up to 6 months, rechargeable battery	Up to 30 days, rechargeable battery	Up to one year, replaceable battery
Bluetooth Range	Over 100 meters	Up to 60 meters	Up to 120 meters	Up to 150 meters	No data available	Over 120 meters
Price	33 CHF	33 CHF	28 CHF	37 CHF	120 CHF + 8 CHF Monthly subscription	40 CHF
Special Features	NFC Tap to retrieve contact information of the owner	N/A	AR Finding, NFC Tap to retrieve contact information of the owner	Rechargeable battery, LED flashlight	Geo-Fencing, SOS Mode, Timeline	QR Code in the back to retrieve contact information of the owner

Table 2.3: Summary of the different state-of-the-art trackers

2.4 Privacy Measures and Concerns

Unfortunately, modern trackers such as AirTags have gotten a lot of negative publicity, as reports of unwanted tracking, stalking, and theft of valuables like cars using AirTags have accumulated since its release in 2021 [19], [21], [23]. This made Apple release a public statement emphasizing the intended use of the AirTag and announcing that they worked with law enforcement to assist them in investigating misuse of AirTags [22]. This section focuses on the countermeasures Apple has taken against the misuse of its AirTag product and investigates the current privacy landscape in BLE and UWB tracking.

2.4.1 Apple's Privacy Measures

Apple anticipated the potential for stalking abuse of the AirTags before their release and implemented an anti-tracking technology in iPhones [12]. These ISAs are time-sensitive alerts sent to the victim if an unknown AirTag has been moving with them for an extended amount of time [47]. The iPhone can detect if an unknown AirTag in *Lost* mode consistently emits lost messages. There were several updates to improve the timeliness of the reports, as in the beginning, there was quite a big delay of a couple of hours until the alert was shown. This renders the feature useless, as a victim may already have reached their home, and a stalker would have already become aware of the victim's domicile address. Apple improved the feature in the iOS 15 beta and reduced the delay to 30 minutes [12].

However, the ISA feature is still flawed in many aspects. On one hand, as it is a built-in iOS feature, only Apple devices receive these alerts. To counteract stalking victims of Android devices using AirTags, Apple developed and released an Android application called *Tracker Detect* for Android devices in December 2021. This allowed Android users to actively scan for BLE devices inside Apple's *Find My* network. The problem with the application is that the owner has to manually scan it repeatedly to find a tracking device. A user would have to be suspicious of potential tracking and then use it, which renders the app useless since nobody would regularly conduct active scanning unless they suspect something is up [7]. Studies have demonstrated solutions to this issue. These are further discussed in Section 3.

Google presented its COFN called *Find My Device* at the Google I/O 2023 [55]. Simultaneously, they also announced a tracker detection system called *unwanted tracker alerts* (UTA), allowing Android users to scan for Apple *Find My* devices [56]. The rollout of these features was planned for the summer of 2023, yet these have been both delayed. In an update to the blog post [56], Google announced that they wanted to give Apple time to develop an adequate mechanism for devices within the Google *Find My Device* network. This is the reason for the delay. At the beginning of 2024, Google's UTA was released, which allows Android users to passively scan for *Find Me* devices. Additionally, Android users are also capable of active scanning [56], which is an improvement to Apple's ISA.

Apple has not implemented any detection mechanisms for Google's *Find My Devices* network, as in a separate, official statement in May 2023, they announced their partnership with Google to submit a: "*industry specification to combat the misuse of Bluetooth* *location-tracking devices for unwanted tracking*" [25]. The goal is a universal, operating-system-solution that helps any device in daily use detect trackers manufactured by different companies. Apple submitted the specifications on the day of the statement's release. The next steps are the review of the specification, then Apple and Google jointly working on implementing the feedback of the reviews, and lastly, the implementation of the production version of the specification by the end of 2023, which will be supported by future versions of iOS and Android [25].

There are other technologies that all account for their Anti-Stalking protections. An example is *Precision Finding*, where a victim with an iPhone 11 or newer is guided to the location of an unknown AirTag utilizing UWB technology. Additionally, the victim can make the AirTag play a sound to facilitate its search [47]. Popular media have criticized the effectiveness of this feature, as the victim could only make the AirTag play a sound once, and if you do not find it, you can't play another sound [57]. Other protection measures include *NFC Scan*, where a victim with an NFC-capable device can tap the AirTag to get specific information such as the serial number of the AirTag or the last four digits of the linked phone number registered with the owner's Apple ID. Apple hereby supports victims by disclosing the stalker's account details if a valid subpoena or law enforcement request is given. Apple also included the functionality that the AirTag repeatedly plays a sound if it gets disconnected from its owner. As a last resort, if a victim finds the unknown AirTag nearby, it can disable the AirTag's tracking capability by removing the CR2032 battery inside it [47].

2.4.2 Privacy Landscape

With wirelessly connected devices becoming ubiquitous in our everyday lives, privacy risks have reached unprecedented levels for modern consumers. The wireless information transferred by communication devices in both active and inactive modes often prioritizes factors such as extending battery life over privacy concerns. Therefore, certain privacy measures have been developed and are widely adopted among manufacturers and specifications. An example of a privacy enhancement is the MAC address randomization, where the hardware identifying MAC address is randomized periodically. Although its general idea is sensible, MAC randomization is flawed, as user-sensitive information is leaked on different levels [58]. [59] conducted a study on BLE mac randomization and its effectiveness and came to the following conclusion: "*[MAC] address randomization [...] fails to provide the promised privacy protection. Various developers and manufacturers do not implement it properly; they rely on public Bluetooth addresses, apply weak randomization, or keep a consistent address for a long time. On the other hand, even if address randomization is properly implemented, other information in the advertisement or in the device might contain unique information that allows for its tracking*". They add that an adversary can track, profile, and potentially even harm the owner of a BLE device, which alerts other devices of its presence by emitting BLE advertisements [59].

Privacy by definition

To this day, there is no universal definition of the term privacy. Although it has been recognized as both a human need and right among many countries [60], the comprehension of the term privacy differs according to prevailing societal characteristics and the economic and cultural environment [61]. This means that the definition of privacy must always be determined based on the current situation and point in time [62]. In 1967, Alan F. Westin, a law professor at Colombia University, defined privacy as: "*The claim of individuals, groups, or institutions to determine for themselves when, how and to what extent information about them is communicated to others*" [63]. In today's context, protecting user privacy has become increasingly difficult, as gathering personal information has become a passive, pervasive, and less intrusive process, making users unaware if their data is being collected [64].

With the ubiquity of data gathering and transfer, the necessity for fundamental groundwork on privacy rights and protection arose. Different governments and researchers have developed proprietary privacy frameworks to increase the protection of end users. These differ in focus and depth as they stem from various economic and cultural environments. The following sections will analyze, discuss, and compare selected frameworks.

European Union: General Data Protection Regulation

In 2016, the *General Data Protection Regulation* (GDPR) was accepted by the European Union, and they claim for it to be "*the toughest privacy and security law in the world*" [65]. Two years later, in May 2018 the legal framework was put into effect. The GDPR aims to signal a firm stance on data privacy and security and applies to protecting the personal data of European citizens or residents. Anyone who processes personal data must adhere to the following seven protection and accountability principles [65]:

- **Lawfulness, fairness and transparency**: Processing must be lawful, fair, and transparent to the data subject.
- **Purpose limitation**: You must process data for the legitimate purposes specified explicitly to the data subject when you collected it.
- **Data minimization**: You should collect and process only as much data as necessary for specified purposes.
- **Accuracy**: You must keep personal data accurate and up to date.
- **Storage limitation**: You may only store personally identifying data for as long as necessary for the specified purpose.
- **Integrity and confidentiality**: Processing must be done in such a way as to ensure appropriate security, integrity, and confidentiality (e.g. by using encryption).
- **Accountability**: The data controller is responsible for being able to demonstrate GDPR compliance with all of these principles.

The GDPR draws attention to factors like consent, transparency, and rights for the data subject. There are many rights assigned to the data subject, yet some of the most important include the *Right to Access* (Article 15), *Right to Erasure* (Article 17), *Right to Restrict Processing* (Article 18), and the *Right to Withdraw Consent* (Article 7) [66]. A data processor and/or data controller must adhere to secure data handling by implementing appropriate technical and organizational measures. Technical measures can range from two-factor authentication on employees' accounts to end-to-end encryption when leveraging cloud services. Organizational measures refer to staff training, company-wide data privacy policies, or the limitations of access to personal data. To conclude, the GDPR grants European citizens or residents rights over their personal data by enforcing obligations on organizations that handle their personal data. These obligations are designed to enhance transparency and ensure lawful data processing, imposing significant fines in cases of non-compliance [65].

NIST: Privacy Framework

The *National Institute of Standards and Technology* (NIST) is part of the U.S. Department of Commerce. Its main goal is to "*promote U.S. innovation and industrial competitiveness by advancing measurement science, standards, and technology in ways that enhance economics security and improve our quality of life* [67]. To achieve this, the NIST publishes many frameworks on various industrial topics. The NIST Privacy framework builds on the premise that managing cybersecurity risks contributes to managing privacy risks, yet is insufficient, as privacy risks can be unrelated to cybersecurity risks [68]. This is shown in the Venn diagram in Figure 2.7. Unrelated privacy risks include privacy events arising from data processing in digital or non-digital form and entail its complete lifecycle from the data collection to its disposal [69].

It is worth mentioning that the NIST Privacy Framework is solely a voluntary tool designed to help organizations identify and manage their privacy risks to build more innovative products and services while simultaneously protecting an individual's privacy. These are important for an organization to follow, as problems resulting from data processing can cause an individual to experience a direct impact, like embarrassment, discrimination, or economic loss. As a consequence, the organization experiences "*impacts such as noncompliance costs, revenue loss arising from customer abandonment of products and services, or harm to its external brand reputation or internal culture*" [70]. These types of impacts are usually managed at the enterprise risk management level.

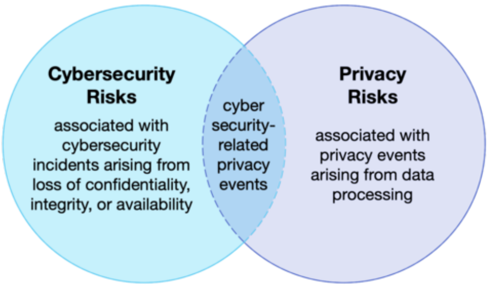
Cybersecurity Risks
associated with cybersecurity incidents arising from loss of confidentiality, integrity, or availability
cyber security-related privacy events
Privacy Risks
associated with privacy events arising from data processing

Figure 2.7: Cybersecurity and Privacy Risk Relationship [68]

The NIST Privacy Framework follows the structure of the *Framework for Improving Critical Infrastructure Cybersecurity* (NIST Cybersecurity Framework). Both frameworks can be used together and have three main parts [68]:

- The **Core** is a set of increasingly granular activities and outcomes that enable an organizational dialogue about the management of privacy risks. It is further divided into key categories and subcategories, which are discrete outcomes of the following five functions:
  - Identify-P: Develop the organizational understanding to manage privacy risk for individuals arising from data processing.
  - Govern-P: Develop and implement the organizational governance structure to enable an ongoing understanding of the organization's risk management priorities that are informed by privacy risk.
  - Control-P: Develop and implement appropriate activities to enable organizations or individuals to manage data with sufficient granularity to manage privacy risks.
  - Communicate-P: Develop and implement appropriate activities to enable organizations and individuals to understand how data are processed and associated privacy risks.
  - Protect-P: Develop and implement appropriate data processing safeguards.
- **Profiles** are a selection of specific Functions, Categories, and Subcategories from the Core that an organization has prioritized to help it manage privacy risk.
- **Implementation Tiers** support communication about whether an organization has sufficient processes and resources in place to manage privacy risk and achieve its target Profile. These are split into four different tiers: (1) Partial, (2) Risk Informed, (3) Repeatable, and (4) Adaptive.

There are no prescribed *Profile* templates, as the framework's goal is to allow for flexibility in its implementation. An organization tailors a Profile to its specific needs and can develop its own additional *Functions*, *Categories*, and *Subcategories* if necessary to account for unique organizational risks. These needs are determined by analyzing the organization's business objectives, privacy values, risk tolerance, and the privacy needs of the individuals who are (in-)directly served or affected by an organization's systems, products, or services. Figure 2.8 shows that there is no specified order of development of *Profiles.* An organization can first develop a *Target Profile* and then develop the current *Profile* to identify gaps. Alternatively, an organization can do it the other way by assessing the current *Profile* and then developing the target *Profile* [68].

The NIST Privacy Framework is a comprehensive and flexible tool for helping organizations manage and prioritize privacy risks in their operations, including data processing. The framework provides a structured approach to privacy risk management, enabling organizations to align their privacy practices with their overall risk management strategies and ensure the necessary resources are allocated. Ultimately, the NIST Privacy Framework supports organizations in adapting to changing privacy challenges in an evolving world by promoting responsible data control.

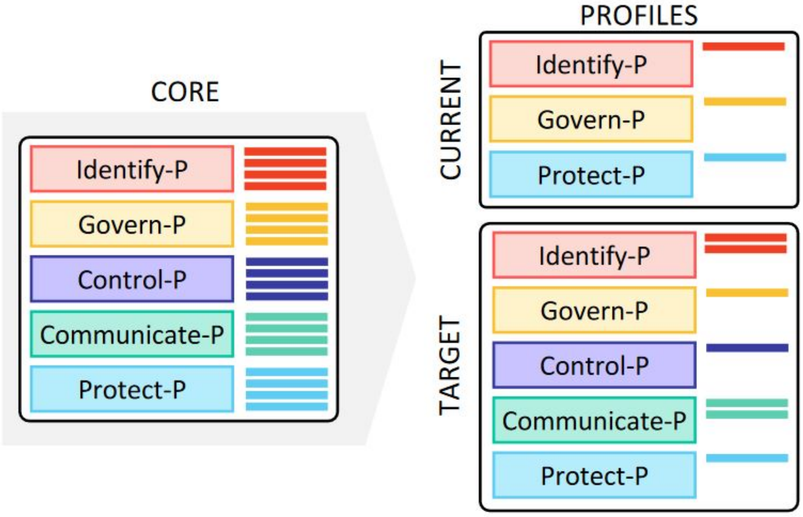

Figure 2.8: Relationship between Core and Profiles [68]

Florence and Trento University Collaboration: Core Ontology for Privacy requirements engineering V.2

Various studies have shown that most privacy concerns can be tackled by the requirements engineers in the design phase of a system under development if the privacy requirements are considered and addressed properly [71], [72]. However, it has also become evident that most requirement engineers are unfamiliar with privacy requirements and how they differentiate from other requirements, like security [73]. To counteract this problem, [74] proposes, implements, validates, and evaluates an ontology that captures key privacy-related concepts and conceptualizes privacy requirements in their social and organizational setting. Under the name *Core Ontology for Privacy requirements engineering* (COPri) [75] proposed an initial ontology in June 2020 [75], and then followed up in 2021 with an improved COPri V.2 version [74] having made improvements to it according to feedback received from privacy and security experts in [75]. COPri defines the relationships between threats and vulnerabilities and their impact on privacy goals at its core. Its concepts are organized in the following four dimensions [74]:

1. **Organizational dimension**: Includes concepts for capturing the social and organizational aspects of the system. These are again divided into several categories:
   - Agentive entities: Capture active entities that are intentional, have goals, and carry out actions toward their fulfillment. A role is further divided into three entities:
     - Data Subject: An identifiable natural person who can identified directly or indirectly by reference to an identifier such as a name, location data, etc.
     - Data Controller: A natural or legal person, public authority, agency, or other body which, alone or jointly with others, determines the purposes and means of the processing of personal information.

    - Data Processor: A natural or legal person, public authority, agency, or any other body, that processes personal information on behalf of the Data Controller.

  - Intentional entities: Capture goals that active entities aim to achieve.
  - Informational entities: Capture informational assets like a user's personal data
  - Ownership, Permission & Consent: Capture who and how can control the use of personal information. They determine who owns it or who has which permissions or consent.

2. **Risk dimension**: Includes risk-related concepts that might endanger privacy needs at the social and organizational levels. Vulnerabilities, threats, and impacts are all part of the risk dimension.

3. **Treatment dimension**: Includes concepts to mitigate risks. Examples of concepts like these are privacy goals, privacy constraints, privacy policies, or privacy mechanisms.

4. **Privacy dimension**: Includes concepts to capture the data subjects' privacy requirements/needs concerning their personal information. Following is a listing of all privacy requirements:

   - Confidentiality: Personal information should remain inaccessible to incidental or intentional threats.
   - Notice: A data subject should be notified when its information is collected. This can be achieved if the data subject gives permission to collect his data.
   - Anonymity: Personal information should be used without disclosing the identity of its data subject. It can be anonymized depending on some privacy mechanism. This can be achieved by removing or substituting primary/secondary identifiers of a data subject (e.g. name, social security number, address, etc.).
   - Unlinkability: It should not be possible to link personal information back to its data subject. A privacy mechanism can be used to remove any linkage between personal information and its data subject. Although unlinkability may appear similar to anonymity, anonymity cannot guarantee unlinkability; each one does not imply the other. An example hereby is, that an attacker might be able to link data to a specific data subject, yet can not determine the identity of the data subject.
   - Unobservability: The aim of unobservability is to hide activities that are performed by a data subject. It should be impossible for others to know whether a data has performed an activity (e.g., use a resource of service) or not.
   - Transparency: A data subject should know who uses its information for what purposes and to what extent.
   - Accountability: A data subject should be able to hold information users accountable for their actions concerning its information.
   - Minimization: The collection of *Personally Identifiable Information* should be kept to a strict minimum.

The COPri ontology shows the relationships between these dimensions and how a data subject's personal information can be exploited by vulnerabilities that different privacy goals can mitigate. These privacy goals interpret the aforementioned privacy requirements, catering to data subjects' privacy needs. COPri was designed to assist requirement engineers while dealing with privacy requirements for systems that handle personal data. It provides a comprehensive set of necessary and sufficient concepts to analyze privacy requirements in their social and organizational context [74].

Switzerland: New Federal Act on Data Protection

The first Swiss Federal Data Protection Act dates back to 1992. With many changes and technological advances in the past 30 years, the Swiss Government decided to release a completely overhauled version of its initial Federal Data Protection Act and released the New Federal Act on Data Protection (nFADP) in September 2023. Its main cause is to protect "*personality and fundamental rights of natural persons, whose personal data is processed*" [76]. As Switzerland is a singled-out country inside of Europe and, therefore, the European Union, it makes sense that there are many similarities and few differences between the GDPR and the nFADP. Table 2.4 shows the few differences.

Topic	**nFADP**	**GDPR**
Sanctions	Up to 250'000 against responsible private persons	Up to EUR 20 million or 4% of the company's worldwide annual revenue
Designation of a Data Protection Officer	Not mandatory but recommended	Mandatory according to article 37 of the GDPR
Data Breach Notifications	Mandatory reporting as soon as possible	Mandatory reporting within 72 hours
Data Exports	Adequacy is determined by Swiss Federal Council	Adequacy is determined by the European Commission
Data Protection Impact Assessment	Consultation of a Data Protection Officer instead of the FDPIC is possible in case of high risk despite measures taken	Duty to consult the supervisory authority in case of high risk despite measures taken
Profiling	General obligation to obtain consent is only imposed for high-risk profiling	General obligation to obtain consent
Sensitive Data	Includes the two additional categories "data on administrative or criminal proceedings and sanctions" and "data on social security measures	According to article 9 of the GPDR

Table 2.4: Differences between the GDPR and the nFADP [77]

By looking at Table 2.4, it becomes evident that the existing differences are on a more specific scale. In both, there is a primary focus on data privacy and protection. It is interesting to point out that the sanctions in the nFADP are specified against responsible private persons, and the GDPR sanctions are directed toward organizations. The concept of profiling is a new addition to the nFADP and is described as the automated processing of personal data. Additionally, the concepts of *Privacy by Design* and *Privacy by Default* are introduced in the nFADP. These are both common principles inside the GDPR (see Article 25 [66]) [65]. *Privacy by Design* implies that developers "*integrate the protection and respect of users' privacy into the very structure of the products of the products or services that collects personal data*" [78]. *Privacy by default*, on the other hand, ensures that the highest possible level of privacy is active by default as soon as the products or services are released [78].

The nFADP defines seven principles for the handling and processing of personal data [76]:

1. Personal data must be discussed **lawfully**.
2. The processing must be carried out in **good faith** and be **proportionate**.
3. Personal data may only be collected for a **specific purpose** that the data subject can **recognise**; personal data may only be further processed in a manner that is **compatible with this purpose**.
4. They shall be **destroyed** or **anonymized** as soon as they are **no longer required** for the purpose of processing.
5. Any person who processes personal data must satisfy themselves that the data are **accurate**. They must take all appropriate measures to correct, delete, or destroy incorrect or incomplete data insofar as the purpose for which they are collected or processed is concerned. The appropriateness of the measures depends, in particular, on the form and the extent of the processing and on the risk that the processing poses to the data subject's personality or fundamental rights.
6. If the **consent** of the data subject is required, such consent is only valid if given **voluntarily** for one or more specific instances of processing based on appropriate information
7. The consent must be explicitly given for:
   (a) processing **sensitive personal data**;
   (b) high-risk profiling by a private person; or
   (c) profiling by a federal body.

As per article 5, paragraph c, of the nFADP, sensitive personal data is partially defined as "*[...] data relating to health, the private sphere or affiliation to a race or ethnicity*" [76]. Therefore, BLE trackers collect location-tracking data, which can be considered part of the private sphere. This leads to the deduction that location tracking data is considered sensitive personal data for which, according to the seventh principle, explicit consent must be given by the data subject. Conclusively, the nFADP is a legal framework that has recently undergone a large overhaul. It shares many similarities to the GDPR, yet some minor differences exist in some categories.

Chapter 3

Related Work

Modern BLE trackers have become the focus of several recent studies. These range from taking a deep dive at COFNs, including Apple's *Find My* network, to bypassing or triggering tracker detection methods manufacturers employ to hinder misuse, all the way to analyzing privacy in IoT. This chapter briefly presents the most important studies.

3.1 Crowdsourced Offline Finding Networks

With the rising importance and implementation of COFNs, many studies have been conducted analyzing the functionality and behavior of these networks. Generally, a higher density of devices inside a COFN signifies a higher likelihood that the BLE packets emitted by a lost device are detected by a finder device [12]. As Apple's *Find My* network is one of the largest and densest OF networks to date, it has become the primary focus of many recent studies. One of these is *OpenHaystack* by [79]. *OpenHaystack* is an open-source framework developed to allow users to integrate custom proprietary Bluetooth-capable devices into Apple's *Find My* network and exploit its services. It relies on the fact that finder iPhones can not distinguish between genuine and fake *Find My* accessories, therefore uploading location reports of both fake and genuine AirTags to Apple's servers. Further studies have used this framework, as it allows mimicking genuine Apple AirTag behaviors with custom-built trackers. Additionally, several papers have analyzed the security and privacy flaws of COFNs. [8] proposes a new design for a secure COFN called *SE-Crow*, which still allows leveraging these networks' benefits without sacrificing security and privacy. In *Blind My*, [17] present the first formal definitions for a privacy-preserving crowdsourced tracking protocol, which is secure against malicious trackers. They add a property called *Beacon Unforgeability* to Apple's *Find My* protocol, which addresses an issue allowing malicious tracking devices to be used covertly to track unsuspecting victims.

3.2 Tracker Detection

Within the *Find My* COFN, Apple has implemented several tracker detection methods to impede the malicious use of AirTags. An example is the ISAs, which is a time-sensitive alert sent to the victim if an unknown AirTag has been moving with them for an extended amount of time [47]. However, several studies have circumvented the triggering of these alerts with various methods. In *Who Tracks the Trackers?*, [12] shows that no alerts will be triggered by creating custom tracking devices and modifying specific bytes in the battery status part of the BLE payload. With a modification to the *OpenHaystack* framework, *Find You*[1] avoids detection by periodically broadcasting new public keys [80]. As Apple's ISAs are solely for iOS and macOS devices, Apple released the App *Tracker Detect*, which enables Android users to scan for Apple AirTags nearby manually. As automatic detection is not built-in, a victim must suspect a potential threat and manually scan for a nearby AirTag. *AirGuard* proposed an initial improvement to this design by allowing the automatic detection and alerting of Android users if they encounter the same AirTag in three separate locations within 24 hours [7]. With *BLE-Doubt*, [81] improved this design by extending it to generic tags, not just AirTags. [16] takes the same approach with *HomeScout* and extends the functionality to automatically detect generic trackers. Additionally, *HomeScout* improves detection time compared to *AirGuard* [16].

3.3 Attacks on BLE Trackers

While many studies have focused on the protective side of detecting nearby trackers, there have been investigative studies focusing on attacks on BLE trackers by exploiting persisting hard- and firmware issues. Several studies have extensively analyzed Apple's Continuity Protocol [58], [82], [83]. Therefore, newer papers have shown that other vulnerabilities such as hardware attacks [10], or attacks on UWB technology [84] can be exposed. In *AirTag of the Clones*, [10] analyzed the hard- and firmware security of AirTags, conducted voltage glitching attacks on the AirTag's nRF chip, and was then able to change the AirTag's configuration data. This allowed him to modify the internal behavior, such as enabling the cloning of an AirTag, customizing its soundset, and using its accelerometer as a microphone. Additionally, the voltage glitching attacks also allowed him to change the BLE and NFC behavior of the AirTag, which could potentially exploited maliciously. [84] takes a different and novel approach by demonstrating the first practical distance reduction attacks against UWB implemented by Apple's U1 chip. He points out that UWB, by design, promises a high-security level based on cryptographic encryption. Still, the actual security level would depend on the obscure design choices made by the UWB signal-receiving device.

[1] https://github.com/positive-security/find-you

3.4 Privacy in IoT

Privacy concerns, including attacks on BLE trackers and insufficient tracker detection methods, hinder the widespread adoption of IoT products. Respecting user privacy is required to ensure confidence and self-assurance in IoT products and related services [85]. Government standards and regulations regarding data handling are frequently updated and revised to ensure the protection of privacy and of end users in the rising widespread adoption of IoT products [65], [68], [76]. Simultaneously, it has become the focus of researchers, as [86] show, that alone in the years 2020 and 2021, there have been close to 40'000 publications on the topic *IoT Privacy* on *Google Scholar* alone. [86] provides a high-level introduction to the current privacy-preserving solution in IoT systems within data collection, transmission, and storage phases. *Privacy Preference Management* (PPM) tools have become popular as they focus on providing rules for transferring data generated by IoT devices to applications. They also include tools that provide transparency to the data flow process, which ameliorates the system's trustworthiness [87]. [87] conclude that although PPM components can improve data privacy in IoT data handling platforms, other challenges remain. These challenges depend on several factors, such as the type of data generated, the context for each type of data, the different applications that want to access the data, and the different purposes for its use. Ultimately, it is up to the data subject to decide on consent and privacy preference options.

Chapter 4

Design

This chapter contains the designs for this thesis's different experiments and mappings. Section 4.1 displays a mapping between the different privacy requirements found in various privacy frameworks and laws thoroughly analyzed in the background, Section 2.4.2. Four different use cases of the Apple AirTag are presented in Section 4.2. Additionally, the design for an experiment on gathering RSSI values is explained in Section 4.3.

4.1 Privacy Requirements Mapping

The design choices outlined below stem from assessing the existing privacy landscape, as discussed in Section 2.4.2. The requirements and specifications were taken from the following regulations and frameworks:

- NIST's Privacy Framework: The subcategories from the Core [68]
- European Union's GDPR: The seven protection and accountability principles for the processing of personal data [65]
- COPri V.2: The Privacy Requirements from the privacy dimension [74]
- Switzerland's nFADP: Seven Principles for the handling and processing of personal data [76]

The following mapping shows the differences and similarities between these privacy-respecting principles. The principles with the highest commonalities are mentioned first, continuing in descending order. As the depth of the description of these requirements differs, some assumptions must be made based on the information and definitions given. These will be thoroughly discussed. Lastly, a table summarizing the findings will be added and discussed briefly.

- **(1) Awareness**: The Data Subject is aware that its personal data is being collected and processed.
  - COPri, Notice: *The Data Subject should be notified when its information is being collected.*
  - GDPR, Lawfulness, Fairness, and Transparency: *Personal data shall be processed [...] in a transparent manner in relation to the data subject.*
  - NIST, Control: Data Processing Policies, Processes and Procedures: *CT.PO-P1: Policies, processes, and procedures for authorizing data processing (e.g., [...] individual consent) [...] are established and in place.*
  - nFADP, Consent: *If the consent of the data subject is required*[1]*, such consent is only valid if given voluntarily for one or more specific instances of processing based on appropriate information.*
- **(2) Transparency**: The Data Subject is aware of the specific data processing purposes.
  - COPri, Transparency: *A Data Subject should be able to know who is using its information and for what purposes.*
  - GDPR, Purpose Limitation: *Personal data shall be collected for specified, explicit and legitimate purposes.*
  - NIST, Communicate: Communication Policies, Processes, and Procedures: *CM.PO-P1: Transparency policies, processes, and procedures for communication data processing purposes [...] are established and in place.*
  - nFADP, Data Collection for specific purposes: *Personal data may only be collected for a specific purpose that the data subject can recognize [and] [...] only be processed in a manner that is compatible with this purpose.*
- **(3) Confidentiality**: Personal information is stored with appropriate security measures.
  - COPri, Confidentiality: *Personal information should remain inaccessible to incidental or intentional threats*
  - GDPR, Integrity and Confidentiality: *Personal data shall be processed in a manner that ensures appropriate security of the personal data, including protection against unauthorised or unlawful processing and against accidental loss, destruction or damage, using appropriate technical or organisational measures*
  - NIST, Protect: Protective Technology: *PR.PT-P4: Mechanisms [...] are implemented to achieve resilience requirements in normal and adverse situations*
- **(4) Accountability**: The Data Subject can hold the controllers for their actions accountable.
  - COPri, Accountability: *A data subject should be able to hold information users accountable for their actions concerning its information.*

[1] According to principle number 7, consent is required for the processing of sensitive personal data (7a), so consent is indeed required.

  - GDPR, Accountability: *The controller shall be responsible for, and be able to demonstrate compliance with, paragraph 1*[2]
  - NIST, Govern: Governance Policies, Processes, and Procedures:
    1. *GV.PO-P3: Roles and responsibilities for the workforce are established with respect to privacy*
    2. *GV.PO-P4: Privacy roles and responsibilities are coordinated and aligned with third-party stakeholders*

- **(5) Data Minimization**: The personal data acquired is kept to a (a) bare minimum and (b) proportionate[3].
  - COPri, Data Minimization: *The collection of Personally Identifiable Information [...] should be kept to a strict minimum* (a).
  - GDPR, Data Minimization: *Personal data shall be adequate, relevant and limited to what is necessary in relation to the purposes for which they are processed* (b).
  - NIST, Control[4]:
    1. Data Processing Management: *CT.DM-P: Data are managed consistent with the organizational's risk strategy to protect individuals' privacy, increase manageability, and enable the implementation of privacy principles*
    2. Disassociated Processing: *T.DP-P: Data processing solutions increase disassociability consistent with the organization's risk strategy to protect individuals' privacy and enable implementation of privacy principles*
  - nFADP, Proportionate: *The processing [of personal data] must be carried out in good faith and be proportionate* (b).

- **(6) Accuracy**: The personal data stored must be accurate and correct.
  - GDPR, Accuracy: *Personal data shall be accurate and, when necessary, kept up to date.*
  - NIST, Control-P: Data Processing Policies, Processes, and Procedures: *CT.PO-P2: Policies, processes, and procedures for enabling data review, transfer, sharing or disclosure, alteration, and deletion are established and in place (e.g., to maintain data quality, [and] manage data retention).*
  - nFADP, Accuracy: *Any person who processes personal data must satisfy themselves that the data are accurate.*

[2] Paragraph 1 refers to the handling guidelines of personal data which are all the other privacy requirements from the GDPR in this section.

[3] Although similar, (a) is on a stricter basis with a higher level of minimization, while (b) concerns the data to be proportionate to its use.

[4] NIST does not define a subcategory dedicated to data minimization itself. Instead, the categories *Data Processing Management* and *Disassociated Processing* both contain principles and guidelines which in turn can lead to data minimization.

- **(7) Storage Limitation**: The data may only be stored for as long as necessary.
    - GDPR, Storage Limitation: *Personal Data shall be kept in a form which permits identification of data subjects for no longer than is necessary for the purposes for which the personal data are processed.*
    - NIST, Control: Data Processing Policies, Processes, and Procedures: *A data life cycle to manage data is aligned and implemented with the system development life cycle to manage systems.*
    - nFADP, Data Removal: *[Personal data] shall be destroyed or anonymised as soon as they are no longer required for the purpose of processing.*
- **(8) Lawfulness**: Data processing is done lawfully.
    - GDPR, Lawfulness, Fairness and Transparency: *Personal data shall be processed lawfully [...] in relation to the data subject.*
    - NIST, Govern: Governance Policies, Processes, and Procedures: *GV.PO-P5: Legal, regulatory, and contractual requirements regarding privacy are understood and managed.*
    - nFADP, Lawful data processing: *Personal data must be processed lawfully.*
- **(9) Anonymity**: Personal data should limit the possible identification of the subject[5].
    - COPri, Anonymity: *Personal information should be used without disclosing the identity of its data subject.*
    - NIST, Control: Disassociated Processing: *CT.DP-P2: Data are processed to limit the identification of individuals.*
- **(10) Unlinkability**: It should be impossible to connect personal information back to the data subject.
    - COPri, Unlinkability: *It should not be possible to link personal information back to its data subject.*
    - NIST, Control: Disassociated Processing: *CT.DP-P1: Data are processed to limit observability and linkability.*
- **(11) Unobservability**: The use of services by the data subject is not visible to third parties.
    - COPri, Unobservability: *Unobservability aims at hiding activities that are performed by a data subject.*
    - NIST, Control: Disassociated Processing: *CT.DP-P1: Data are processed to limit observability and linkability.*
- **(12) Good Faith**: Personal Data must be processed in good faith.
    - nFADP, Good Faith: *The processing must be carried out in good faith*

[5] This can be avoided by employing de-identification privacy techniques such as removing identifying data of a data subject (e.g., social security number, name, address, etc.).

By analyzing the four different frameworks and regulations, similarities and differences become evident. Starting with the *Awareness*, the data subject must be aware of its data being processed. The GDPR (*article 6.1 (a)*), nFADP (*article 6.6*), the NIST (*CT.PO-P1*), and COPri (*Notice*) all require the data subject to give consent before the collection and processing of its data. Although COPri solely requires notification of the data subject, a notice violation is raised in the case of unpermitted data collection.

Another requirement, *Transparency*, necessitates informing the data subject about the precise purpose of the data processing. This is specified in all four frameworks. The nFADP defines the requirement with "*Personal data may only be collected for a specific purpose that the data subject can recognize*" [76].

The requirement *Confidentiality* is mentioned by COPri (*Confidentiality*), the GDPR (*article 5.1 (f)*), and the NIST (*PR.PT-P4*). All imply that appropriate security measures are necessary when storing and handling the personal data of a data subject. The keyword *appropriate* is challenging to define as modern technology is constantly changing and evolving, allowing for more modern technologies on one side and exposing new threats on the other. One must analyze the current circumstances and technologies involved in collecting and processing personal data to define *appropriate* security measures.

The next aspect is *Accountability.* In this context, it means that the data subject can hold the controllers of its data accountable for their actions. This is clearly stated in both COPri (*Accountability)* and the GDPR (*article 5.2*). NIST takes an indirect approach by saying that roles and responsibilities concerning privacy should be established and that these are coordinated and aligned with third-party stakeholders. It is not clearly stated that the data subject can hold the controllers of its data accountable. However, it is implied that the controllers' specific responsibilities to the data subject are established through these roles. These, in turn, can allow the data subject to hold the controllers accountable for their actions. The outlined reasoning leads to the plausible conclusion that the NIST can be seen as a partial match on the requirement of *Accountability.*

Following, *Data Minimization* handles the extent to which the personal data is stored. Hereby, it is important to differentiate between (a) keeping personal information to a bare minimum necessary and (b) a proportionate storing of it. These may seem similar, yet (a) refers to the principle of *Data Minimization* on a stricter basis and (b) allows for more leeway. The prior (a) forms part of COPri's privacy requirements, as "*the collection of Personally Identifiable Information [...] should be kept to a strict minimum*" [74]. The GDPR (*Article 5.1 (c)*) and nFADP (*Article 6.2)* can be situated into the latter (b). NIST does not declare a subcategory for defining data minimization itself. However, in two separate categories, *Data Processing Management* and *Disassociated Processing*, which are both allocated inside the *Control* function, data minimization is mentioned as a desired outcome. As no definitive subcategory handles this specification, the NIST can only partially be mapped on the aspect of *Data Minimization.*

The privacy requirement *Accuracy* states that personal data stored must be accurate and correct. Both the GDPR (*Article 5.1 (d)* and the nFADP (*Article 6.5*) mention this, with the GDPR additionally stating that it must be up to date. *Accuracy* is also indirectly contained inside the NIST, as *CT.PO-P2* states that policies, processes, and procedures for enabling data review, alteration, or deletion are in place. This allows for maintaining data quality and deleting incorrect or outdated data. Once again, as it is not implicitly stated inside the NIST that personal data stored must be accurate, the NIST can only be given a partial match to the requirement of *Accuracy*.

Storage Limitation is a privacy concern regarding the length of the storage duration of personal data. It is defined as that data may only be stored for as long as necessary. The GDPR (*article 5.1 (e)* and nFADP (*article 6.4*) contain this principle. In the NIST *Control: Data Processing Policies, Processes, and Procedures* subcategories, a data life cycle is mentioned, aligned, and implemented with the system development life cycle. This life cycle implies that at one point in time, the data will either be destroyed or retained if it is still required. Therefore, to a certain extent, the aspect of *Storage Limitation* can be found inside the NIST privacy framework, making it a partial match in the mapping.

The *Lawfulness* processing of personal data is part of the GDPR (*article 5.1 (a)*), NIST (*GV.PO-P5*) and the nFADP (*article 6.1*). The GDPR applies to all EU citizens or residents. Therefore, if a company processes their data or offers its goods or services to them, the GDPR is applicable. Similarly, the nFADP applies to all organizations that process the personal data of Swiss citizens. Hereby, it does not matter whether the organization is based in Switzerland or not. This shows that different laws and regulations apply to different countries and regions to which a data collector and processor must adhere. Several examples show [76], [88] that these are getting updated and reinstated regularly. Therefore, conducting lawful behavior entails always staying up-to-date with current regulations and guidelines in the regions where the company collects and processes personal data.

Although similar, *Anonymity* and *Unlinkability* employ different privacy-preserving principles. As per COPri (*Anonymity*), *Anonymity* means that while processing personal information, the identity of the data subject should be disclosed. This can be achieved by removing or substituting identifiers such as name, social security number, or address from the data. This can also be found inside NIST's *Disasossiacted Processing* category (*VT.DP-P2*). *Unlinkability* refers to the impossibility of linking personal information back to the data subject. This is defined in COPri (*Unlinkability*) and NIST (*CT.DP-P1*). These principles do not imply each other. For example, if an attacker can link personal data back to the data subject, yet the subject's identity is not revealed, unlinkability is breached while anonymity is still granted.

Lastly, *Unobservability* is a principle that aims at hiding actions from the data subject to third parties. This is included inside COPri (*Unobservability*) and the NIST (*CT.DP-P1*). The principle of handling personal data in *Good Faith* is only mentioned inside the nFADP (*article 6.2*).

This concludes the privacy requirements contained inside the four frameworks. Table 4.1 shows their differences and commonalities. It is important to note that the nFADP and the GDPR act as proprietary laws inside their respective regions. On the other hand, COPri and NIST's privacy framework solely act as guidance by providing recommendations on the different possibilities and necessities regarding handling personal data. Four of five functions from NIST's privacy framework have been mentioned. Only *Identify-P* has been omitted from the selection above, as it handles the development of organizational understanding to manage privacy risks company-wide. Therefore, as it solely focuses on

Index	Privacy Requirement	NIST [68]	COPri V.2 [74]	GDPR [66]	nFADP [76]
1	**Awareness**	✓	✓	✓	✓
2	**Transparency**	✓	✓	✓	✓
3	**Confidentiality**	✓	✓	✓	
4	**Accountability**	(✓)	✓	✓	
5	**Data Minimization** **(a) minimum** **(b) proportionate**	(✓) (b)	✓(a)	✓(b)	
6	**Accuracy**	(✓)		✓	✓
7	**Storage Limitation**	(✓)		✓	✓
8	**Lawfulness**	✓		✓	✓
9	**Anonymity**	✓	✓		
10	**Unlinkability**	✓	✓		
11	**Unobservability**	✓	✓		
12	**Good Faith**				✓

Table 4.1: Overview of the different privacy requirements

identifying roles and responsibilities inside the different tiers of an organization, there are no contributing factors to the privacy requirements listed above.

4.2 Privacy Use Cases

In the following section, various privacy use cases are created and explained. These range from the intended use of Apple's AirTags to malicious use, where the data of unknown victims is collected. These four use cases all follow the assumption that the owner of the AirTag is a Swiss citizen. In the subsequent part, these four privacy use cases will be evaluated on the discussed privacy requirements from Section 4.1.

- **Normal Use**:

  1. An AirTag is attached to the user's personal backpack. The user walks around Zurich with it yet never takes it off.
  2. An AirTag is attached to a user's suitcase, which the user checks in at the Zurich airport check-in counter. Sight of it is lost, yet the user can track it through the *Find My* app. After a short flight to Frankfurt Airport, the user collects the suitcase from the baggage retrieval area.

- **Malicious Use**:

  3. An AirTag is used to stalk a stranger. A stalker places the AirTag in an unknowing victim's jacket while riding in public transportation. As the victim exits the train, the stalker tracks its whereabouts using the *Find My* app. The victim has an Apple iPhone but only gets a notification when it reaches its

home. It taps on the notification, and with the help of a sound played by the AirTag, it can locate the AirTag and bring it to the police.

4. In another stalking case, an AirTag is slid into a victim's backpack at the public library. Unknowingly, the victim grabs the backpack as it is about to leave to go home. As the victim has a non-Apple smartphone and does not have Apple's *Tracker Detect* App installed, the victim is not notified of the unwanted tracker and continues with its life. Meanwhile, the stalker gathers more information about the victim, its domicile address, workplace, and friends' addresses.

4.3 RSSI Values Experiments

In a related matter, a design for experiments on collecting RSSI values from AirTag's BLE signals is presented in the following section. The goal of these experiments is to eventually predict the distance between a signal-receiving device (smartphone) and a BLE signal-emitting device (AirTag). RSSI is used to measure the relative distance between the two, yet as research suggests, it is inaccurate when trying to measure the absolute distance [16]. Only with enough data in different environments and circumstances, the absolute distance can be determined [31]. With the presented experiments, many RSSI values are collected within a controlled environment under changing conditions and distances. The data will be collected using the HomeScout[6] application by [16], which is run on an Android smartphone.

These experiments demonstrate a possible solution in an ameliorated detection mechanism for unknown trackers. With enough data, tracker detection mechanisms could increase accuracy with a faster detection time and result in fewer false positives. The experiments with the different conditions and distances are presented subsequently:

1. Placement of a single AirTag on top of the smartphone for 15 minutes.

2. Placement of a single AirTag at the following distances between the smartphone and the AirTag for 15 minutes: [5cm, 10cm, 15cm, 20cm, 25cm, 50cm, 100cm]

3. Placement of a single AirTag at the following distances between the smartphone and the AirTag for 15 minutes: [1.5m, 2m, 2.5m, 3m, 3.5m, 4m, 4.5m, 5m]

4. Placement of a single AirTag at the following distances between the smartphone and the AirTag for 15 minutes: [10m, 15m, 20m, 25m, 30m, 35m, 40m, 45m, 50m, 60m, 70m, 80m, 90m, 100m]

5. Placement of an AirTag at varying distances to the smartphone, yet with other BLE-emitting devices (AirTags) in between them.

6. Placement of an AirTag with various obstructions in between, such as walls or human bodies, for example.

[6] https://github.com/LouisBienz/HomeScout

Chapter 5

Results and Evaluation

The following chapter shows the evaluation of the different designs from Chapter 4. It starts by analyzing different use cases for Apple's AirTag, which are evaluated against various privacy requirements. Further, the APEC *Cross-Border Privacy Rules* (CBPR) and *Privacy Recognition for Processors* (PRP) systems are analyzed, and their privacy requirements are mapped to current privacy requirements mapping from Chapter 4. The following section analyzes the ongoing lawsuit *Hughes et al. versus Apple*, and further threats to the AirTag system are identified. Lastly, the different threats discovered in this thesis undergo a DREAD risk assessment to assess which level of priority and severity they should be given.

5.1 Privacy Use Cases and Requirements

Privacy requirements and regulations have been extensively analyzed in Section 4.1, where various requirements from different frameworks have been mapped on top of each other, displaying commonalities and differences. Simultaneously, four use cases of Apple's AirTag have been created, differing between the everyday, intended use of AirTags and the malicious use involving the stalking of victims. This section analyzes these scenarios by reviewing if and how the different privacy requirements are applied. With this, it is essential to define the various roles and relationships. The data subject is the person whose data is collected and processed by Apple. Therefore, Apple acts as the data processor. The owner of an AirTag can be considered the data subject as long as its data is being collected. This means that as long as the AirTag is attached to himself or its personal belongings, the owner of the AirTag is the data subject. If the AirTag starts collecting data from another individual, the data subject role is transferred to that person. Apple's *Find My* network acts as the data processing environment, which will be the main focus point of the following analysis. As AirTags emit BLE advertisements and use UWB technology, these technologies will also be analyzed regarding their privacy measures.

Index	Privacy Requirement	User-dependent?
1	Awareness	yes
2	Transparency	yes
3	Confidentiality	no
4	Accountability	yes
5	Data Minimization	no
6	Accuracy	no
7	Storage Limitation	no
8	Lawfulness	yes
9	Anonymity	no
10	Unlinkability	no
11	Unobservability	no
12	Good Faith	yes

Table 5.1: Overview of user-dependency in privacy requirements.

5.1.1 Privacy Requirements Classification

From the twelve privacy requirements defined in Section 4.1, some have to be discussed on a more general level, outside of the declared use cases. Implementing these is unrelated to how an AirTag owner uses it. They define a broader, system-related scope, which is user-independent. The characteristics of user behavior do not come into play, as these requirements pertain to the overall functioning and operation of the environment, irrespective of individual user actions or behaviors. Consequently, these will be discussed separately concerning how Apple implements and manages measures to enforce them.

Table 5.1 shows an overview of the classification. The user-independent privacy requirements are thoroughly analyzed in Subsection 5.1.2. Section 5.1.3 shows the user-dependent privacy requirements and evaluates Apple's implementation on different use cases.

The privacy requirements analysis will be primarily based on Apple's documentation of installed privacy measures. The following literature is referred to:

- Apple: Platform Security protocol [89] (last updated: 2022).
- Apple: Location Services [90] (last updated: 2019).
- Apple: A day in the life of your data [91] (last updated: 2021).
- Apple: Legal - Privacy Governance [92] (Website accessed: 10.12.2023).
- Apple: Privacy overview [93] (Website accessed: 10.12.2023).
- Additionally, various studies are also consulted and cited accordingly.

Studies are the main point of criticism for Apple. While Apple's documentation of the privacy measures shows the extent to which privacy measures were implemented, various recent studies have dissected them using methods like reverse-engineering to learn how to emulate the behavior of Apple's AirTags and the *Find My* network.

5.1.2 User-Independent Privacy Requirements Analysis

For the following user-independent privacy requirements analysis, each privacy requirement will receive a dedicated subsection, discussing its implementation and whether it is sufficient to fulfill the requirement.

(3) Confidentiality: *Confidentiality* requires adequate security measures to protect the data subjects' data. With *Privacy by design*, Apple promises to have implemented security best practices to protect user data [90]. Inside their *Find My* ecosystem, Apple employs *end-to-end* encryption in the form of advanced public key cryptography. An elliptic curve P-224 private encryption key pair containing private and public keys is generated on the user's device. Using an iCloud keychain, Apple synchronizes the private key pair and a secret among a user's devices. Importantly, Apple does not have access to the private key pair and secret [89], and there have been no reports of Apple violating this promise [44].

Apple uses the elliptic curve P-224, as the entire public key representation can be fit into a single BLE payload [89]. P-224 has been recommended by the NIST and even approved for use by the U.S. Federal Government [94]. Some cryptographers discourage the use of the NIST P-224 curve [95], yet [44] state that there have not been any practical attacks against it when used exclusively for an Elliptic-curve Diffie-Hellman key exchange, which is the case here. For this reason, it can be assumed that the elliptic curve P-224 employed to grant *end-to-end* encryption inside Apple's *Find My* ecosystem is secure and, therefore, *confidentiality* is given on the aspect of *end-to-end* encryption.

Conversely, [44] has demonstrated a somewhat concerning aspect. The advertisement keys are exchanged every 15 minutes, and OF can store the location reports from the last seven days; a total of 672[1] advertisement keys per device exist. They could all be generated from a master beacon key, yet Apple decided to cache the advertisement keys. This is most likely for performance reasons. These keys are cached on macOS in a directory readable by any application with user privileges. Therefore, any third-party application with user privileges can exploit this to access the historical geolocation data, threatening the *Confidentiality* requirement.

(5) Data Minimization: *Data Minimization* is one of Apple's core privacy principles. They promise to limit the specific personal information gathered to the bare minimum [91]. Concerning the data gathered by finder devices, the reports they upload to Apple's servers contain the current location of the finder device, an estimate of location accuracy, the time the advertisement was received, and the attempted upload time [47]. A lot of information is left out. Some examples of additional data that could be included are the BLE signal strength, the device identifier, or the battery information of the AirTag. Apple's *Find My* app combines multiple location reports from finder devices to generate a higher precision rate of the AirTag's exact location [89]. Therefore, to evaluate and combine the different location reports, additional information such as an estimation of the location accuracy and a timestamp of the received BLE advertisement is necessary and kept to a bare minimum. Therefore, Apple adheres to the *Data Minimization* privacy requirement on a strict basis (a) (see Section 4.1).

[1] The calculation works as follows: 7 days * 24 hours * 4 keys per hour = 672 keys generated over a week

Another example of how Apple adheres to *Data Minimization* becomes evident by analyzing when AirTags emit BLE packets. AirTags do not constantly emit BLE advertisement packets. Only when they are disconnected from their owner devices for more than 15 minutes does their state change to *lost*, which is when it starts broadcasting BLE advertisement packets.

(6) Accuracy: The privacy requirement *Accuracy* ensures that the personal data stored is accurate and correct. As stated in [89], a higher *accuracy* is achieved by aggregating the location reports of different finder devices and considering aspects like location accuracy estimates. [44] demonstrate that compared to previous solutions, Apple's *Find My* network provides a high level of accuracy, especially when using a slower transportation mode. This is mainly because more nearby finding devices pick up the BLE packets while walking or sitting stationary in a public place. Apple's high accuracy level in generating location reports is also because OF is enabled for all *Find My* compatible Apple devices by default when the devices are updated to iOS 13 or later, iPadOS 13.1 or after, and macOS 10.15 or after. This leads to a high-density level of Apple finder devices inside the *Find My* network, positively impacting accuracy.

To ensure the legitimacy of the uploaded location reports, each request is authenticated by Apple before its upload to Apple's servers. However, [47] points out that Apple has to implement better mechanisms to detect whether the uploaded location reports belong to a registered AirTag. With tools like OpenHaystack [79], researchers have been able to create fake AirTags, which act similarly to the real ones. This could potentially be exploited by uploading fake location reports, reducing the accuracy of the combined location reports.

Accuracy can also refer to other personal data stored in the Apple ID, such as the personal E-Mail, Address, or name. To guarantee the correctness and accuracy of this data, Apple lets users access it and manually change it if necessary [93].

(7) Storage Limitation: Apple stores the location reports from finder devices for seven days on their server [44]. Seven days are appropriate and adhere to the *Storage Limitation* principle, as it may take a while until the owner of the lost AirTag starts querying the *Find My* network for it. This behavior can have multiple reasons, such as the owner not noticing that the AirTag is gone. Therefore, *Storage Limitation* is adhered to.

(9) Anonymity and (10) Unlinkability: *Anonymity* is important, as it requires removing personal information when using different communication technologies simultaneously. Apple employs a de-identification process on personal data and considers data to be de-identified if: "*all personal data elements [...] [are] removed, including full IP address and any identifiers linked to personal data*" [92]. AirTags employ MAC Address Randomization, a feature that increases the difficulty of tracking a BLE device over an extensive period, as its Bluetooth address is changed frequently. This is used for the technologies BLE and UWB. The effectiveness of MAC Address Randomization has been criticized, as both the content [96], [97] and timing [98] of the BLE frame can be leveraged. However, these weaknesses have been partially fixed by [99] and for this reason, MAC Address Randomization can be considered an effective measure to increase *Unlinkability.*

Unlinkability states that linking personal information to the data subject should be impossible. Apple employs a technique called *cross-transport key derivation*, which allows a device to use different communication technologies simultaneously [89]. In Apple's case, devices with classic Bluetooth and BLE capabilities, *cross-transport key derivation* allows using different keys. This makes it more challenging for adversaries to link activities and communications from the same device across different channels. This is commonly known as a practice to implement *Unlinkability*.

However, with AirTags, granting *Unlinkability* is rather challenging. AirTags collect geolocation data on the whereabouts of their users. With enough data, an adversary can detect the most visited places from an individual. This method is called *fingerprinting* and [100] proved that four spatiotemporal points are sufficient to identify 95% of all individuals in an anonymized location dataset. Additionally, [101] has implemented an RSSI-based *fingerprinting* mechanism by linking BLE traces emitted by the same device despite MAC address randomization. These examples show that fingerprinting threatens both *Anonymity* and *Unlinkability*.

Apple has implemented an authentication system for uploading location reports using finder devices. As part of the authentication method, the finder device has to reveal a device-specific identifier in the HTTPS request header. This can be utilized to link multiple reports to the same finder device. When downloading a location report, the owner's device includes its Apple ID in the HTTPS request header. This allows Apple to link reports uploaded by a particular finder to the Apple ID of the downloading owners [44]. These are both examples of breaches against the *Unlinkability* requirement. This not only involves linking multiple pieces of information to the same user. By revealing the Apple ID's of both the finder and owner devices, the aspect of *Anonymity* is also not granted and could be exploited. [44] argue that the finder device authentication with an Apple ID is necessary to ensure that no fake reports can be uploaded. Yet, they see no reason the owner device must authenticate to Apple's servers by providing personally identifiable information. Any Apple device can arbitrarily query Apple's servers for location reports [44]. This deems the authentication of owner devices on download unnecessary, and removing it would increase *Anonymity* and *Unlinkability*.

(11) Unobservability: *Unobservability* refers to the fact that the actions of the data subject are hidden from any third party. With Apple's AirTags, this concerns actions like the owner of a lost AirTag querying for location reports and accessing them to check where it is. As stated in the subsection on *Unlinkability*, when downloading a location report, the owner of an AirTag has to authenticate himself by revealing a unique device identifier. Apple is, therefore, aware that the Airtag's owner is accessing the location report. As the data processor, Apple is not a third party in the data lifecycle. However, Apple technically can store the number of data accesses to the data the user has done, which could potentially be leaked if a third party managed to hack into Apple's servers. This is another reason why the owner authentication before downloading location reports can be exploitable.

The scope of *Unobservability* can be broadened to ensure that the actions performed by the AirTag remain hidden from any third parties. However, regarding the emission of BLE packets to nearby listening Bluetooth devices, adhering to this quickly becomes increasingly tricky. Currently, any third-party device equipped with BLE capabilities can sniff for BLE packets and determine whether the packet was sent from an AirTag. This is mainly because inside the BLE advertisement packet, there are company identifiers and information regarding what type of device has emitted it.

For these reasons, *Unobservability* concerning AirTags is challenging to adhere to. Simultaneously, there is no proof that Apple does not abide by it. Therefore, the conclusion that *Unobservability* is adhered to can be drawn.

5.1.3 User-Dependent Privacy Requirements Analysis

Following, the privacy requirements *Awareness*, *Transparency*, *Accountability*, *Lawfulness*, and *Good Faith* will be discussed with the respective use cases. These are user-dependent, meaning their applicability depends on specific usage scenarios.

Normal Use: Use Cases 1 and 2

Use cases 1 and 2 depict everyday usage scenarios of AirTags. Hereby, in use case 1, the AirTag is attached to a personal item of the owner, yet it never leaves the owner, as it carries the AirTag around in a stroll around the city. This leaves the AirTag always in the *connected* state. Use case 2 differs, as the AirTag is attached to a user's suitcase at the airport, and as it checks the suitcase in, the AirTag will lose BLE connection to the owner device and enter *nearby* mode. After 15 minutes in the *nearby* mode, the AirTag's state will change to *lost* mode. The AirTag starts emitting *lost* BLE advertisements, which are picked up by nearby Apple devices. These, in turn, upload encrypted location coordinates to Apple's servers. This allows the user to know the whereabouts of its AirTag as the user moves around the airport. Having landed in a new country, the owner can see that the AirTag has also made it. After a short while, the user picks up their suitcase from the luggage retrieval area, and the AirTag's state returns to the *connected* state.

(1) Awareness: By initially pairing the AirTag, the user is confronted with the screen depicted in figure 5.1. Hereby, the user is informed of the AirTag's purpose of tracking personal belongings and that tracking other people without their consent can lead to prosecutions by law enforcement. On one hand, the AirTag's purpose is evident when a user purchases it. Therefore, on purchase, the user indirectly consents to the AirTags functionality of collecting personal data such as location data. On the other hand, the user must give explicit consent to the collection of its data before collecting and processing it. This is achieved in the pairing process, as the user is asked to allow location data collection by the smartphone, as it improves the location accuracy. Therefore, *Awareness* is granted in use cases 1 and 2 as explicit consent is required.

(2) Transparency: *Transparency* necessitates that the data subject is aware of the exact purpose of the data processing being done. As the user grants consent to collect its data upon setting up the connection with AirTag, it knows the purpose and extent to which its data is being processed. Therefore, the privacy requirement *Transparency* is adhered to in use cases 1 and 2.

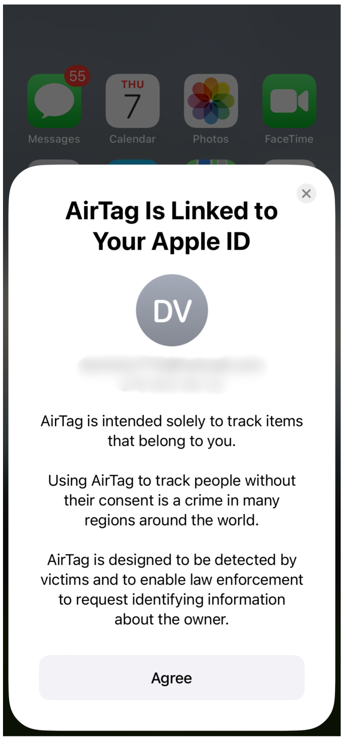

AirTag Is Linked to Your Apple ID

DV

[redacted]

AirTag is intended solely to track items that belong to you.

Using AirTag to track people without their consent is a crime in many regions around the world.

AirTag is designed to be detected by victims and to enable law enforcement to request identifying information about the owner.

Agree

Figure 5.1: AirTag Pairing Screen Notice

(4) Accountability: To comply with the privacy requirement *Accountability*, a user must know who it can hold accountable for collecting and handling its data. In the case of AirTags, Apple takes on the role of the data processor. According to [92], since 2014, Apple has received privacy accountability certifications that adhere to the *APEC Privacy Framework*. In APEC countries, Apple must abide by the *APEC CBPR System* and the *PRP System*. A more in-depth analysis of the APEC CBPR and PRP Systems is discussed in Section 5.2. However, one of APEC's privacy principles, to which Apple abides, is the principle of *Accountability*. Therefore, the user can hold the data controller accountable as it walks around the city or leaves the AirTag in its suitcase, as displayed by use cases 1 and 2.

(8) Lawfulness: As *Lawfulness* varies depending on the jurisdiction where the user is currently situated, use cases 1 and 2 must be viewed separately. Use case 1 states that the owner of the AirTag takes a stroll around Zurich with the AirTag in its backpack. Since Zurich is in Switzerland, the jurisdiction falls under the Swiss data protection laws, the nFADP. Since the user controls their own data and the AirTag is used in an intended manner, this use case can be considered lawful under the nFADP. In use case 2, the AirTag is attached to a user's suitcase and travels to the Frankfurt airport. This scenario also falls under the jurisdiction of the nFADP, as the AirTag is used by a Swiss citizen. The AirTag is used for personal tracking purposes. Therefore, use cases 1 and 2 align with the privacy requirement *Lawfulness*.

(12) Good Faith: Whether *Good Faith* is granted or not can not be conclusively determined. As Apple takes on the role of data collector and processor, it can not be directly assumed that they operate in *Good Faith*. Whether a company acts in *Good Faith* depends on several factors among many stakeholders and individuals. With Apple's statement on unwanted tracking using AirTags [22] in 2022, Apple reminds that they designed the AirTag to help people locate their personal belongings, emphasizing that its intended use is not for tracking other people and that they highly condemn malicious use of it. Nevertheless, it is difficult to determine if Apple handles the data collected and processed from AirTags in *Good Faith*. Research shows that Apple has not yet broken its promise that it can not access and decrypt the data of the location reports from finder iPhones [7]. Yet, this shows that it has not yet been proven otherwise. Considering these reasons, there is no evidence against Apple handling in *Good Faith*; therefore, in use cases 1 and 2, *Good Faith* is adhered to.

Malicious Use: Use Cases 3 and 4

While Use Cases 1 and 2 handle everyday scenarios of the AirTag, Use Cases 3 and 4 were designed to cover niche stalking scenarios. Since the AirTags release, it has been used maliciously by stalkers, who surreptitiously place it with their victims, allowing them to track their victim's whereabouts without having to be near them physically. In both malicious use cases, the victim rides in public transport as the stalker places the AirTag with him. In use case 3, it's put in the victim's jacket; in use case 4, it's put in the victim's backpack. Another aspect where the use cases differ is that in use case 3, the victim has an iPhone; in use case 4, the victim has a non-apple smartphone and, therefore, no automatic tracker detection feature. As Apple has implemented its iOS-based ISA detection method, the victim in use case 3 gets a notification letting him know of the AirTag following him. With the help of a sound played by the AirTag, the victim can locate it and bring it to the police. As of writing this thesis, the rollout of Android's UTA was successful. Yet, as this is a fresh addition to the existing tracker detection mechanisms, it is not considered in this evaluation. Therefore, the only tracker detection mechanism considered in this evaluation is Apple's *tracker detect* app. However, in use case 4, the non-Apple smartphone user has not installed it. This results in the fact that the AirTag moving along with the victim in use case 4 is not detected, and the victim is not made aware of it. It continues its life as the stalker gathers information on the victim's whereabouts. Following, both use cases are analyzed on their adherence to the privacy requirements of *Awareness*, *Transparency*, *Accountability*, *Lawfulness*, and *Good Faith*.

(1) Awareness: *Awareness* assures that the data subject knows the data being processed. As per the definition, it can be extended to the data subject, who has to consent to the collection of its data. In both malicious use cases, the data subject's role, in the beginning, is the stalker. Upon setting up the AirTag, it consents to collecting and processing its location data. However, when the stalker places the AirTag with the victim while riding public transport, the data subject role is transferred to the new stalker's victim. It is unaware of the collection and processing of its location data and never gives consent to it.

There is a slight difference between the two malicious cases as in use case 3, the victim is eventually made aware of the unknown tracker following him with Apple's ISA detection mechanism. He receives a notification informing him of the tracker following him, and then he can locate it with the auditory aid of playing a sound on the AirTag. Apple has significantly improved the detection time of the ISA feature since its release. However, in many real-life stalking scenarios using AirTags, victims deem the feature unreliable, as the time-sensitive alert would be generated promptly. At the point in time when the victim saw it, they would have already returned home and revealed the private address to the stalker. This was the case in [57], and similarly was taken into the design of use case 3. However, importantly, in use case 3, the victim finds the AirTag and can bring it to the police. Research has also criticized the ISA feature, as [12] has demonstrated several ways to disable the triggering of the ISA. This can be achieved by emitting invalid advertisement packets through the wrong bytes set for the battery status or by manually periodically changing the advertisement key of a lost AirTag.

In 2022, Apple announced their collaboration with local law enforcement to combat stalking issues using AirTags [22]. Since every AirTag has a unique serial number, and every AirTag in use is paired with an Apple ID, Apple will provide the account details of the AirTag's owner to law enforcement, given a subpoena or a valid request is in place. Apple emphasizes that there have been successful cases of stalkers being apprehended and charged with the help of Apple's information [22]. [24] and the following analysis of the privacy requirement *Good Faith* show that Apple does not keep this promise. Instead, they tend to complicate law enforcement investigations. In contrast to use case 3, use case 4 differs as the victim is never aware of the tracker collecting its location data, so the victim never becomes aware of the AirTag.

To conclude the privacy requirement *Awareness*, after transferring the data subject's role from the AirTag owner to the stalker victim, no previous consent is acquired to collect and process its location data. This applies to both malicious use cases. Out of this reasoning, *Awareness* is not adhered to. As in use case 3, the victim is made aware of the processing of its data with the ISA after it reaches its home, *Awareness* per definition[2] is granted. However, this must be put into perspective. On one hand, *Awareness* is only achieved *during* the data collection and processing process. On the other hand, no consent was given by the data subject. These aspects will be simultaneously viewed while analyzing the element of *Awareness*. For these reasons, a partial adherence to *Awareness* can be concluded for use case 3. Use case 4 does not adhere to it at all.

(2) Transparency: *Transparency* requires the data subject to be informed of the precise purpose of the data processing. In use cases 3 and 4, the owner of the AirTag, the stalker, is aware of the purpose, as it initially sets up the connection to the AirTag. During this time, *Transparency* is granted. However, as soon as the stalker slips the AirTag with the victim, the data subject role transfers to the victim, who is unaware that it is being tracked. There is a dependency of *Transparency* on the previously discussed privacy requirement *Awareness*. If *Awareness* is not granted, i.e., the data subject is unaware of its data being processed, *Transparency* is not adhered to either. As use case 4 does not adhere to the *Awareness* requirement, *Transparency* is not granted either.

The adherence to *Transparency* in use case 3 can not be determined trivially. As the user finds the AirTag and brings it to the police, one could argue that if the victim knows the product, it also knows for what purposes the data is being collected and processed. If

[2] The definition of *Awareness* solely requires that the data subject must be aware of its data being processed. Therefore, consent from the data subject is not required but strongly encouraged.

the product were unknown to the stalker's victim, it would get the ISA notification and think there would be no harm to it. In use case 3, the victim tracks down the AirTag and consequently brings it to the police, indicating some knowledge of the product and its capabilities. The basis of this analysis lies in implicit versus explicit data subject informing on the purposes of the data collection and processing. Even though explicit informing of the data subject is not clearly defined in the definition of *Transparency* in Section 4, a more careful approach suggests that it should be considered. For this reason, *Transparency* is not granted by either malicious use case, as no explicit informing of the data subject on the purposes of the data collection and processing takes place.

(4) Accountability: *Accountability* is essential, as it guarantees that the data subject and other third parties can hold the data controllers accountable for their actions. It is necessary if the data subject has experienced any misuse or breach of their personal information. Through *Accountability*, Apple, as the data controller, can be held responsible. Misuse occurs in the use cases 3 and 4. Apple's item finder, the AirTag, tracks strangers without their consent. Therefore, it is essential that Apple can be held accountable for the misuse of the AirTag.

There is an ongoing lawsuit [24] by several plaintiffs against Apple. The plaintiffs claim that the AirTag has enabled their stalkers, calling it "*the weapon of choice of stalkers and abusers*" [23]. Extensively and descriptively, the lawsuit describes the functionality of the AirTag, the different detection measures available, each with a set of deficiencies, and the plaintiffs' personal experiences. As this lawsuit is integral to understanding to what extent Apple abides by protecting user privacy from a victim's perspective, the lawsuit is discussed and thoroughly analyzed in Section 5.3. To summarize, Apple is being held accountable for its actions regarding the AirTag, therefore abiding by the privacy requirement *Accountability.* This applies to both use cases 3 and 4.

(8) Lawfulness: In the malicious use cases, a victim is unknowingly and unwillingly tracked using an AirTag. Not all privacy requirements are adhered to, for example, *Awareness*, and *Transparency* can not be granted. Whether Apple, as the data controller and processor, abides by the law depends on the jurisdiction, as privacy laws differ worldwide. As *Awareness* and *Transparency* are both privacy requirements in the analyzed legal frameworks, the GDPR and nFADP, *Lawfulness* is consequently not granted in their respective jurisdictions.

(12) Good Faith: It is generally challenging to determine whether a data controller handles the data subject's information in *Good Faith.* As an outside observer, one can not know what happens behind Apple's closed doors. An analysis of *Good Faith* requires a combination of (1) transparency in Apple's privacy policies and its adherence to the data protection regulations and (2) a track record of responsible data handling practices. With several publishments [22], [89], [90], [91], [92], [93] Apple asserts that they maintain transparent, user-focused privacy policies. Yet, this is the case with many companies. While they may claim to prioritize user privacy, the actual implementation and safeguards must be evaluated to determine the user's trust in handling its data. As an outside observer, one must rely on reported user experiences to assess whether *Good Faith* is adhered to. The lawsuit, *Hughes versus Apple*, analyzed later in Section 5.3, is taken for reference.

With this lawsuit, it becomes evident that Apple does not keep all its promises. Concerning Apple's claim that they cooperate with local law enforcement when identifying the owners of AirTags used for stalking-related purposes, there are several negative experiences made by plaintiffs where no full cooperation was shown. In [22], Apple promises to "*provide the paired account details in response to a subpoena or valid request from local law enforcement*". Yet, experiences made by plaintiffs in [24] show quite the contrary. Apple is restrictive when handing out identifying information on the owners of AirTags. As further analyzed in Section 5.3, Apple can only hand out such relevant information in cases where the pairing was fewer than 25 days ago [102]. Additionally, in another example, plaintiff Kacz's case demonstrates that Apple prefers protecting the identity of stalkers over aiding their victims. Both cases demonstrate Apple's limited cooperation with local law enforcement in identifying individuals who misuse AirTags for malicious practices like stalking. This advocates firmly against handling in *Good Faith*.

Before the AirTag's release in April 2021, there were several concerns about the AirTag's tracking capabilities being used maliciously and the respective mitigation efforts in place being insufficient. In an interview with NPR, Eva Galperin, the Director of Cybersecurity at the *Electronic Frontier Foundation*, an international non-profit digital rights group, expressed her worries as follows: "*I was concerned ahead of their release as soon as I figured out how they worked. [...] The fact that they chose to bring the product out to market in the state it was in last year is shameful*" [103]. Other voices expressed similar concerns, yet Apple responded to allegations like these with a press campaign dismissing and minimizing the concerns and going as far as calling AirTags *Stalker-Proof* [24].

This marketing campaign suggests that Apple knowingly misled the press and public opinion on the potential risks of the AirTag. They released an unfinished product to the market, intending to address known safety and privacy issues only after the product had already been launched. This negligence in fulfilling their responsibility resulted in the endangerment of numerous lives, as the AirTag was exploited for malicious purposes, including theft, stalking, and even murder. There is little *Good Faith* to be seen in this, as Apple prioritizes profits over safety and privacy.

The privacy requirement *Good Faith* can also be viewed from a manufacturer's perspective. Apple could assume that its consumers act in *Good Faith* when buying their products. After all, Figure 5.1 reminds the owner of an AirTag that its purpose is to act as an item finder, not a people tracker. Yet, many items can be misused for malicious purposes. Objects, like knives or axes, can diverge from their official uses when used to stab other people. Similar to AirTags, these can be bought anonymously. In contrast, other weapons, such as firearms, usually require official licensing from the holder registered with the government. This traceability of firearms through licensing facilitates a more efficient process of identifying shooters from gunshot scenes and conducting thorough investigations. These are clear examples of how items can be misused for malicious reasons. It is careless behavior by the manufacturer to not implement enough safeguards that limit the misuse potential of certain products. Regarding *Good Faith*, Apple shows apparent negligence in two different aspects:

- Safeguards: As previously analyzed and Section 5.3 shows, the safeguards implemented to protect victims from stalkers have been insufficient since the AirTags' release and still are to the point in time of writing this thesis. Especially regarding Android users, Apple has been negligent in its protective duty and has not taken appropriate measures. Even for Apple users, the detection system of stalker AirTags at the point of the AirTags' release was faulty and only slowly improved over time.

- Cooperation with stalking victims and law enforcement: Disturbingly, examples have shown that Apple prefers to protect their consumers, who misused the AirTag for stalking purposes, over aiding stalking investigations initiated by the victim. With a pairing information access restriction after 25 days, Apple appears to have intentionally complicated the identification of stalkers without a clear justification. Arguments like *Storage Limitation*[3] make little sense, as aiding stalking victims to identify their stalkers outweighs the other.

Out of these reasons, concerning malicious use cases 3 and 4, there is enough evidence to conclude that Apple does not act in *Good Faith.* With the *Find My* ecosystem, any device is turned into a tracking device. Yet, the Apple AirTag has established itself as the primary stalking device due to its small size, long-lasting battery life, and cheap price. Daily encounters with Apple devices are almost impossible to evade, making the AirTag inside the *Find My* network a gift to stalkers. Apple ignored several warnings before Airtag's release. Instead, they initiated a campaign referring to the AirTag as *stalker-proof* to successfully launch a potentially hazardous product. There is little *Good Faith* to be seen in this, and clearly, Apple does not adhere to it.

5.2 Analysis of the CBPR and PRP Systems by APEC

As previously mentioned, since 2014, Apple has received yearly privacy accountability certifications on its adherence to the *APEC Privacy Framework.* In APEC Countries, Apple abides by the *APEC CBPR* and *APEC PRP* systems [92]. This section provides an in-depth analysis of the APEC with the different frameworks Apple abides by.

The APEC is an economic forum with the primary goal of supporting sustainable economic growth and prosperity in the Asia-Pacific region. It ensures that goods, services, investments, and people move more easily across borders inside participating countries. Currently, there are 21 countries in the APEC, with the largest economies belonging to the United States, Russia, Japan, South Korea, and Canada [104]. The following sections review and analyze the APEC systems CBPR and PRP.

[3] Apple could refer to *Storage Limitation* as an argument, stating that 25 days would be storing the data for as long as necessary. Yet, from an outside perspective, we do not know whether Apple decides to (1) delete the pairing information or (2) elect not to share it under any given circumstances. In either case, the argument could be treated more as an excuse rather than having any validity.

5.2.1 CBPR Analysis

The CBPR was developed to grant "*continued free flow of personal information across borders while establishing meaningful protection for the privacy and security of personal information*" [105]. It was endorsed in 2004, updated in 2015, and comprises APEC's nine guiding privacy principles from the APEC privacy framework [106]. There are four purposes of the framework [105]:

1. Development of appropriate privacy protections for personal information.
2. Enable global organizations that collect, access, use, or process data in APEC economies to develop and implement uniform approaches within their organizations that can be used across APEC borders.
3. Assist enforcement agencies in fulfilling their mandate to protect information privacy.
4. Advance international mechanisms to promote and enforce information privacy and maintain the continuity of information flows among APEC economies and their trading partners.

Therefore, while the APEC CBPR's primary goal is to protect the data subject, it also aims to facilitate the collection and processing of data for corporations in APEC countries by creating standardized regulations on data privacy. It is not intended to replace or change an APEC member's domestic laws and regulations. Instead, it is meant to act as a minimum level of protection for countries with little to no domestic privacy protection requirements. On the other hand, if a country's privacy protection regulations exceed what the CBPR expects, the full extent of the domestic law and regulations apply.

The APEC CBPR will be read consistently with the APEC Privacy Framework [106]. The initial APEC Privacy Framework was modeled on the guidelines of the *Protection of Privacy and Trans-Border Flows of Personal Data* from the *Organisation for Economic Co-operation and Development* (OECD). Since its release in 2005, it was updated in 2015, drawing upon the concepts introduced into the OECD guidelines [107] in 2013. Following, the nine APEC information privacy principles are demonstrated and explained [106]:

- Preventing Harm: Protecting personal information should prevent misuse, with obligations and remedies proportionate to the associated risks and potential harm.
- Notice: Personal information controllers should provide clear and easily accessible statements about their practices and policies with respect to personal information that should include:
  - the fact that personal information is being collected.
  - the purpose for which personal information is collected.
  - the types of persons or organizations to whom personal information might be disclosed.

  - the identity, location, and contact information of the personal information controller.
  - the choices and means the personal information controller offers individuals for limiting the use and disclosure of, and for accessing and correcting, their personal information.

  All reasonable practicable steps shall be taken to ensure that such notice is provided either before or at the time of collection of personal information. Otherwise, such notice should be provided as soon after as possible.

- Collection Limitation: The collection of personal information should be limited to information that is relevant to the purposes of collection and any such information should be obtained by lawful and fair means, and where appropriate, with notice to, or consent of, the individual concerned.
- Uses of Personal Information: Personal information should be used only to fulfill the purposes of collection and other compatible or related purposes except:
  - with the consent of the individual whose personal data is collected.
  - when necessary to provide a service or product requested by the individual.
  - by the authority of law and other legal instruments, proclamations, and pronouncements of legal effect.
- Choice: Where appropriate, individuals should be provided with clear, prominent, easily understandable, accessible, and affordable mechanisms to exercise choice to the collection, use, and disclosure of their personal information.
- Integrity of Personal Information: Personal information should be accurate, complete, and up-to-date to the extent necessary.
- Security Safeguards: Personal information controllers should protect personal information with appropriate safeguards against risks, such as loss or unauthorized access to personal information or unauthorized destruction, use, modification, or disclosure of information or other misuses.
- Access and Correction: Individuals should be able to;
  - obtained from the personal information controller confirmation of whether or not the personal information controller holds personal information about them;
  - have communicated to them after having provided sufficient proof of their identity and personal information about them;
    1. within a reasonable time;
    2. at a charge, if any, that is not excessive;
    3. in a reasonable manner;
    4. in a form that is generally understandable;
    5. Challenge the accuracy of the personal information relating to them and, if possible and appropriate, have the information rectified, completed, amended, or deleted.

- Accountability: A personal information controller should be accountable for complying with measures that give effect to the previously stated principles.

5.2.2 Mapping APEC Privacy Requirements to Privacy Requirements

To understand the depth of the APEC privacy requirements, a comparison with the privacy requirements mapping from Section 4.1 is provided in the following section. Each requirement is analyzed for any potential commonalities and differences.

Preventing Harm: Although *Preventing Harm's* definition is rather widely applicable, ranging from privacy protection mechanisms such as education and awareness campaigns, laws, regulations, and law enforcement mechanisms to organizational controls, it focuses mainly on preventing harm to the data subject as a result of wrongful collection or misuse of personal information. From an organization's perspective, this can be achieved through appropriate security measures to protect the data subject's data. This overlaps with the definition of *Confidentiality*, which requires precisely this.

Notice: *Notice* ensures that individuals know what information is collected and its purpose. This overlaps with the privacy requirement *Awareness* from the privacy requirement mapping, yet with APEC, no explicit consent is required to collect the data. Additionally, with APEC's *Notice* requirement, the data subject should be informed before or at the time of the data collection. There are exemptions where notice can be provided after data collection. This differs from the previously analyzed frameworks, which all required consent from the data subject before collecting its data. However, in general, there is a big overlapping between APEC's *Notice* definition and the privacy requirement *Awareness* from the privacy requirements mapping.

Collection Limitation: This principle overlaps with three separate privacy requirements from the mapping. The most obvious one is a match with the *Data Minimization* principle, which is the primary goal of *Collection Limitation*, the "*collection of personal information should be limited to information that is relevant to the purposes of collection*" [106]. *Data Minimization* is further divided into the (a) strict minimal collection or (b) proportionate collection of personal data. By analyzing APEC's definition of *Collection Limitation*, it becomes evident that it can be categorized into a (b) proportionate collection of personal information. Additionally, by stating that obtaining the information should be done lawfully and fairly, there is a match with the privacy requirement *Lawfulness*, which is included in the nFADP [76]. Lastly, the definition of *Collection Limitation* extends the definition of *Notice*, requiring, if appropriate, the individual to consent to collecting its data. This shows that if applicable under given circumstances, consent from the data subject is required. A case is described where consent is unnecessary, as the data collection would serve a more significant cause. Yet, Apple's AirTag data collection process does not fall into this category and, therefore, requires the data subject to give explicit consent.

Uses of Personal Information: This requirement handles using personal data only to fulfill the specified collection purposes. This overlaps with the definition of the privacy requirement *Transparency*, which necessitates that the personal data may only be collected for the specific purpose stated and that the data subject is informed of that purpose.

Uses of Personal Information requires considering the individual's expectations which overlap with the purpose that the data subject is aware of. For this reason, *Transparency* from the privacy requirements mapping overlaps with the *Uses of Personal Information* requirement.

Choice: *Choice* refers to the principle that the data subject has control over collecting, using, and disclosing its personal information. Although there are some minor overlaps with the privacy requirements, *Awareness*[4] and *Transparency*[5], APEC's definition of *Choice* can be seen as a proprietary privacy requirement, which is not contained in the privacy requirements mapping, as it adds an additional level of control to the user, giving it control over the collection of its data.

Integrity of Personal Information: With *Integrity*, the personal information collected, processed, and stored should be accurate, complete, and kept up-to-date. This overlaps with the principle of *Accuracy* from the privacy requirement mapping.

Security Safeguards: To adhere to the *Security Safeguards* requirement, the data controller should implement appropriate security safeguards against risks. There is some overlapping with APEC's *Preventing Harm* principle, yet *Security Safeguards* focuses more on the security safeguards aspect. In contrast, *Preventing Harm* mainly aims to protect the data subject on a broader scope. Nevertheless, there exists an overlap with the privacy requirement of *Confidentiality*, as both address the deployment of safeguards to mitigate potential risks.

Access and Correction: This privacy requirement refers to the capability of the data subject to access and correct its personal information. Access should be provided in a reasonable manner and form. This privacy requirement is not included in the privacy requirements mapping from section 4.1. Therefore, it can be seen as a proprietary privacy requirement.

Accountability: Similarly to the definition of *Accountability* in the privacy requirements mapping, to adhere to APEC's requirement of *Accountability*, the data controller can be held accountable for the data processing and collection of a data subjects' data. Upon transferring personal data to another person or organization, the personal information controller should verify that the recipient protects the data consistently with the APEC privacy principles. Like in NIST's Privacy Framework [68], it is not explicitly mentioned that the data subject should be able to hold the data controller accountable. However, it can be implied that any stakeholder in the data collection and processing process can hold the data controller accountable for its actions.

To conclude, the APEC privacy framework adds two new requirements to the mapping: *Choice* and *Access and Correction*. Otherwise, APEC matches many of the GDPR's privacy requirements.

[4] *Awareness* requires the data subject to be aware of the processing of its data. By giving the data subject a choice on to what extent its data should be processed, it is made aware of the processing.

[5] *Transparency* can be matched, as the organization should provide clear notice about the choices available to the data subject

Privacy Requirement	**NIST**	**COPri**	**GDPR**	**nFADP**	**APEC**
Awareness	✓	✓	✓	✓	✓
Transparency	✓	✓	✓	✓	✓
Confidentiality	✓	✓	✓		✓
Accountability	(✓)	✓	✓		✓
Data Minimization **(a) minimum** **(b) proportionate**	(✓) (b)	✓(a)	✓(b)		✓(b)
Accuracy	(✓)		✓	✓	✓
Storage Limitation	(✓)		✓	✓	
Lawfulness	✓		✓	✓	✓
Anonymity	✓	✓			
Unlinkability	✓	✓			
Unobservability	✓	✓			
Good Faith				✓	
Choice					✓
Access and Correction					✓

Table 5.2: Adjusted Privacy Requirements Mapping with APEC's Privacy Requirements

5.2.3 PRP Analysis

The APEC *Privacy Recognition for Processors* is designed to help data processors demonstrate their ability to effectively implement a data controller's privacy obligations related to processing personal information. It also allows data controllers to identify qualified and accountable data processors. The APEC CBPR [105] and the APEC Privacy Framework [106] only apply to data controllers. In a data controller-processor relationship, according to the accountability principle, the data controller is responsible for the activities the data processor performs on their behalf. Conclusively, the PRP is a certification for data processors [108].

5.2.4 Conclusions on APEC Systems

Since 2014, Apple has adhered to the APEC Privacy Framework. For this, they have received privacy accountability certifications. In APEC countries, Apple voluntarily abides by the APEC CBPR and the APEC PRP [92]. With the CBPR and PRP, APEC aims to facilitate data transfer inside APEC economies. The CBPR aims to develop appropriate privacy protections for personal information. With a uniform approach to data collection and handling across different countries, a simplified approach to data transfer for global organizations is achieved, promoting and enforcing information privacy and maintaining the continuity of information flows. The PRP acts as a certification guided towards data processors. On the one hand, it helps data controllers to identify qualified and accountable processors. On the other, it promotes more minor to medium-sized data processors to become a part of a global data processing network [108]. The APEC privacy framework contains many similarities with other privacy frameworks. Table 5.2 shows the addition of the APEC privacy framework to the existing privacy requirements mapping from Section 4.1. While there are many similarities between the APEC privacy framework and the mapping, two requirements from the APEC framework are not contained: *Choice* and *Access and Correction.* As these are important standalone requirements adding to the core definition of privacy, they are added as separate requirements.

5.3 Lawsuit Analysis: Hughes et al. versus Apple

Following, a thorough analysis of the ongoing lawsuit *Hughes et al. versus Apple* is presented. It is integral to this thesis to give insights into the victims' perspective on the AirTag stalking cases. So far, only Apple's perspective and research work have been presented while evaluating the adherence to the different privacy requirements. References in this section are from the official 140-page lawsuit [24]. Other references are cited accordingly. As this thesis is written with no prior legal knowledge and in the context of a bachelor thesis with the Communication Systems Group at the Department of Informatics, the following lawsuit analysis will focus on its technical rather than its legal perspective.

5.3.1 Class Action Divison

The lawsuit was filed in October 2023 in the state of California. Along with the lead plaintiff, 53 plaintiffs form part of a class action lawsuit against Apple. In total, four main class allegations exist, with additional sub-classes. These consist of the following:

- The iOS Stalked Class (31 plaintiffs): All persons residing in the United States who own iOS devices and were tracked, without consent, by Apple's AirTag.
- The Android Stalked Class (7 plaintiffs): All persons residing in the United States who own Android devices (and who do not own iOS devices) were tracked, without consent, by Apple's AirTag.
- The iOS At-Risk-Of-Stalking Class (31 plaintiffs): All persons residing in the United States who own iOS devices.
- The Android At-Risk-Of-Stalking Class (7 plaintiffs): All persons living in the United States who own Android devices

With this, it is important to note that the respective plaintiffs of the *iOS Stalked Class* and *iOS At-Risk-Of-Stalking Class* are the same members. Equally, this can also be applied to the Android-related classes.

5.3.2 Main Allegations

The lawsuit begins with an extensive introduction to various topics such as stalking with AirTags, its technologies and functionality, the plaintiffs involved with some of their experiences, and the tracking methods employed by Apple. As it is written extensively and descriptively, the key takeaways of this introduction are presented below:

- Modern technology has increased the tools available for stalkers. Significantly, the use of real-time location information is one of the most dangerous and frightening instruments available to stalkers.

- Apple's AirTag stands out from any competitor product due to its unparalleled accuracy, ease of use, and affordability.

- Before the release of the AirTag, several advocates and technologists urged Apple to rethink the product and to consider its inevitable use in stalking. Apple responded by dismissing concerns and pointing to mitigation features, claiming the AirTag is *stalker proof*.

- A survey by *Motherboard*, an online tech news publication by *VICE Media*, shows that within a year of the release of the AirTag, there were at least 150 police reports filed, where AirTags were used by stalkers to track their victims. It is believed that the exact number of AirTag stalking cases is significantly higher, as the 150 reports only contain cases that (1) were reported and (2) could be obtained [109].

- Consequently, due to the release of the AirTag, multiple murders have occurred in which the murderer used an AirTag to track the victim. Similarly, individuals have been murdered when using AirTags to track down stolen property and confront the thieves.

- International stalking cases with AirTags have spiked with an increased amount of reports of malicious AirTags in use, e.g., in the United Kingdom [110] and India [111].

- Several deficiencies in detecting AirTags for iOS and Android users exist. An overview is displayed in Table 5.3. The severity and details of these deficiencies are discussed later on.

- Victims of stalking with AirTags have little meaningful recourse in the criminal justice system. Apple claims to cooperate with law enforcement on AirTag-related requests [22]. There are, however, several critical aspects as to why this statement can be deemed misleading:

  - Stalking behavior is not a crime in many jurisdictions. As there is no criminalization, no charges can be brought forward. For this reason, police are often disinterested in pursuing stalking cases, leading to no fulsome investigation at all. Several plaintiffs and news reports have described this [110].

  - Apple has not kept its full promise in aiding law enforcement agencies combat AirTag stalking. In an example with plaintiff Hopkins, the police sent a valid law enforcement request to find the owner of an AirTag used in a stalking incident. Apple only responded with the information that the AirTags were bought in a four-pack. With another plaintiff, Apple was sent a subpoena, to which they responded with an Excel spreadsheet containing no identifying information about the owner of the AirTag, i.e., the stalker.

    A possible explanation for this reluctance to hand out identifying information of an AirTags' owner can be found inside Apple's *Legal Process Guidelines* [102]. Here, Apple mentions "*With a serial number, Apple may be able to provide the paired account details in response to a subpoena or greater legal process. AirTag pairing history is available for a period of up to 25 days*" [102]. This concludes that to receive useful identifying information on the owner of an AirTag, (1) the AirTag must have been paired to the owner's device within the last 25 days, and (2) the subpoena or valid request must also be sent within the same period. This shows that Apple does not fully cooperate with law enforcement agencies, deeming Apple's statement in [22] unreliable and misleading.

  - In a related matter, plaintiff Kacz's experience shows Apple preferring to protect the owners of AirTags that are using them maliciously, as opposed to aiding the stalker's victims.

- Children are particularly vulnerable to stalking with AirTags. Various plaintiffs claim that a stranger or estranged spouse has placed an AirTag with their child to (1) track the child's or (2) parents' location. As younger children are not usually equipped with smartphones, no automatic detection features are in place, making them susceptible to unnoticed tracking.

Table 5.3 shows different safety features employed in Apple AirTags and lists their deficiencies. The *Unknown AirTag Screen Alerts* refers to the previously discussed ISA feature. It works on devices running iOS 14.5 or later. Although Apple has significantly reduced the alert time of a nearby unknown AirTag to around 30 minutes [12], some plaintiffs claim it may still take up to a day to receive the alert after being tracked. The ISA can be disabled inadvertently if the iPhone owner turns off *Location Services* in their smartphone settings. This is usually done for privacy-related reasons. Yet, with this, the owner unknowingly disables the ISA feature. Other deficiencies regarding the ISA feature are that it cannot be triggered independently and can not be relied upon. There is no feature for an owner of an iPhone to trigger a scan manually if they are suspicious of being stalked. Apple's algorithms trigger the alerts. If the alert is shown once, there is no guarantee that it will reappear. This issue happened with plaintiff Araujo, whose daughter received the ISA while Araujo was driving. When they pulled over, the alert disappeared, and they could not locate the nearby AirTag.

Another detection feature under heavy critique is AirTag's sound alerts. If an Apple AirTag has been disconnected from its owner extensively, a sound alert on the AirTag gets triggered. Apple has not officially disclosed the alert sound's volume, but according to [24], it is estimated to be approximately 60 dB. This level compares to a normal conversation or ambient background music. Additionally, the sound is not very distinct,

Safety Feature	**Operating System**	**Deficiency**
Unknown AirTag Screen Alerts	iOS	Alert is not immediate
		Can be disabled inadvertently
		Cannot be triggered independently
		Reliability
Sound Alerts	iOS and Android	Alert is not immediate
		Cannot be triggered independently
		Volume is insufficient
		Sound is not distinct
		Duration
		Can be disabled with ease
Disabling AirTags	iOS and Android	Physical possession of the AirTag is required
AirTag Identifier Reset	iOS (and potentially Android)	Resetting identifiers also resets Apple's unknown tracker search logic
AirTag Firmware Updates	iOS (and potentially Android)	Relies on AirTag owners to implement
Tracker Detect App	Android	Low Awareness
		Does not run in the background
		Triggering sound alerts

Table 5.3: Summary of deficiencies of various AirTag safety features [24]

meaning it could be mistaken for noise from other devices. This combination leads to a problematic conclusion. If the victim is hearing impaired, in a loud environment, or the sound played by the AirTag is muffled, there is no way for the victim to pick up on it. Additionally, a victim will not hear the sound alert if an AirTag is placed behind a car's license plate. While the sound alert is flawed in its design, it can also be disabled quite easily by manually removing the speaker [112]. This goes further, as so-called *silent* AirTags started being sold on mainstream e-commerce sites like eBay and Etsy [113].

Dealing with a discovered AirTag from a stalker is not as trivial as it may seem. AirTags can be deactivated by removing the battery. However, law enforcement agencies have advised against that, as removing the battery could contaminate the evidence. Jennifer Landhuis, the director of the Stalking Prevention Awareness and Resource Center, additionally points out that bringing the AirTag to the local police might also bear risks: "*If the offender is monitoring the victim's actions and sees the AirTag has now gone to [...] [a] police station, that can escalate the situation and put a victim more in danger*" [24]. Apple has released a protocol [114] for victims to follow if they find an unknown AirTag. They recommend disabling the AirTag, and if the victim feels like its safety is at risk, to contact local law enforcement [114]. These measures might increase the danger level that the victim is in for a short period. Yet, local law enforcement must react adequately to any potential danger and protect the stalking victim. Not doing anything with a found AirTag just to appease the stalker is not a viable solution either.

Next to Apple's ISA, their *Tracker Detect* app, designed to protect Android users from AirTag stalking, is also heavily criticized in the lawsuit. It was released only in late 2021, and the app seems deeply flawed in its design. It is minimalistic in design and functionality, only allowing users to scan for nearby trackers manually. There is no background scanning functionality, and it can not issue push notifications. Android users have to selectively conduct scans on suspicion of being tracked. It was rendered useless in the case of plaintiff Jane Doe 1. She lives in a densely populated area, and after manually scanning for nearby trackers, the *Tracker Detect* app would tell her that AirTags are nearby. Its functionality is limited to detecting nearby *Find My* devices. It cannot tell whether a specific AirTag has been following a user for extensive time. It is evident that Apple has not tried its hardest to design an Android app with the same capabilities as Apple's ISA feature. Yet, what makes this discussion interesting is that in March 2022, soon after the release of the *Tracker Detect* app, Google noted that it would already be possible for Apple's *Tracker Detect* app to run in the background [115]. Apple responded to this by claiming that "*Continuous background scanning with Tracker Detect on Android would negatively impact battery life and other features of Bluetooth*" [115]. Apple insists the only power-efficient solution for background scanning would be on an Android Operating System (OS) level. No data is available on how high Apple's *Tracker Detect* app's battery consumption would be with a background detection feature. However, Apple knowingly dismissed background scanning as a privacy mechanism to prioritize battery life, shifting the burden onto its competitor, Google.

The *AirTag Firmware Updates* safety feature in Table 5.3 refers to the fact that safety measure improvements to the AirTag are usually employed through firmware updates. The lawsuit claims these updates do not happen automatically with a background process. Instead, the owner would have to implement the updates manually. While it is true that the firmware updates are installed through a connection to the owner's iPhone, this statement contradicts [10] findings. [10] claims that the "*AirTag silently updates in the background without indicating this to the owner, and the owner cannot stop this process*". In the lawsuit, no reference is given to the aforementioned statement, so it is difficult to determine where such information was gathered and evaluate its integrity. Because of this, [10] finding outweighs the lawsuit's claim that the owner of an AirTag has control over firmware updates.

The *AirTag Identifier Reset* Safety Feature refers to the pseudo-random public key, which appears in lost messages. The lawsuit claims that it would change regularly, hindering the ISA from detecting an AirTag as the iPhone would assume, that each different key would belong to a different AirTag. As a reference, an article [116] by the tech media outlet, *macworld*, is cited. Again, research suggests quite the contrary. According to [17], the pseudo-random public key of a lost message is only exchanged daily. As the ISA's detection time is around 30 minutes as of June 2021, there should be no interference [12]. However, [12] has proved that manually changing the key every few minutes can prevent triggering the ISA feature. Generally, Apple has not published any information on the functionality and flagging mechanisms of the ISA. There has been lots of research on it [12], [17], [47], yet no findings on the exact mechanisms have been published yet.

All deficiencies in Table 5.3 are considered stalking threats as they enable the successful stalking of victims. Stalking is possible when the victim is not aware that it has been tracked for an extensive amount of time. Therefore, these threats can be seen as breaches of the privacy requirement *Awareness*.

To conclude, these allegations indicate a compelling argument suggesting that Apple may have been negligent in its duty of taking responsibility when releasing a small but powerful tracker. The AirTag, with its cutting-edge technology, has unfortunately been misused on numerous occasions, with instances ranging from stalking and theft to even more severe cases like murder. The plaintiffs accuse Apple on multiple accounts of *Negligence* in its duty of care in its design, marketing, and introduction into the market of its AirTags. The previously mentioned design defects played a substantial factor in causing harm to the plaintiffs. Simultaneously, Apple is accused of *Unjust Enrichment*, profiting from selling a potentially dangerous and defective product. On a count of *Intrusion upon Seclusion*, the plaintiffs argue their right to privacy was breached by Apple's unique position to monitor individuals through AirTags, resulting in an intentional intrusion into the plaintiff's privacy affairs.

5.4 Privacy Risk Assessment

With several privacy requirement breaches presented by Sections 5.1 and 5.3, these must be classified on their severity concerning user privacy. To systematically classify the risks associated with these threats, Microsoft's DREAD risk assessment model [117] is used. Why DREAD was chosen over other existing risk assessment frameworks is explained in Section 5.4.1. Furthermore, it is explained and analyzed in 5.4.2 and applied to AirTag threats in Section 5.4.3.

5.4.1 Risk Assessment Frameworks

Risk assessment tools are used to identify, estimate, and prioritize risks to the assets and operations of an organization. Similar to threat modeling, multiple risk assessment tools are widely adopted. Some examples include the NIST Risk Management Framework [118], OWASP Risk Rating Methodology [119], Factor Analysis of Information Risk (FAIR) [120], and DREAD by Microsoft [117]. Out of these, DREAD by Microsoft stands out as it has been adopted by recent IoT-related work to rank threats based on severities [121]. Additionally, it adds a high level of clarity through its simplicity. Therefore, DREAD is used in this thesis and further described in Section 5.4.2.

5.4.2 DREAD Risk Assessment Framework

DREAD is an acronym that stands for the following five criteria when assessing potential threats. The five criteria are evaluated for each given threat, and a threat score between one and ten is given. [117]:

- **D**amage: The damage resulting from an attack. 1 = Low Damage; 10 = High Damage.
- **R**eproducibility: How often a specified type of attack will succeed. 1 = Low Reproducibility; 10 = High Reproducibility.
- **E**xploitability: The required expertise and effort necessary to mount an attack. 1 = Low Exploitability; 10 = High Exploitability.
- **A**ffected Users: The number of users that could be affected by an attack. 1 = Low Amount of Affected Users; 10 = High Amount of Affected Users.
- **D**iscoverability: The likelihood a threat will be exploited. 1 = Low Likelihood; 10 = High Likelihood.

An average of the five scores is a threat's risk indicator. A higher number indicates a more serious threat and should, therefore, be given a higher priority. A threat with a low risk indicator can be treated with a low priority.

There are some drawbacks to the DREAD framework, as it has been criticized for not having sufficient objectivity in the risk assessment process and that the numerical values of risk attributes assigned to threats are not constant throughout the life span of the system applying the model. There has been work improving this by creating systems based on DREAD with a higher stability and resolution rate concerning subjective choices [122], yet for the following risk assessment, Microsoft's model will be applied due to its high simplicity and clarity. DREAD's subjectivity can be overlooked, as it importantly establishes relative comparisons on the severity of different threats. Generally, Risk assessment models must analyze different factors independent of each other. DREAD achieves this, as none of the five criteria correlate to each other [123].

5.4.3 AirTag DREAD Model

Figure 5.2 shows a diagram of the several AirTag-related threats identified in this thesis. A majority of the threats can be considered stalking-related, as they mainly focus on feature vulnerabilities like the ISA, *tracker detect* app, or *sound alert* features.

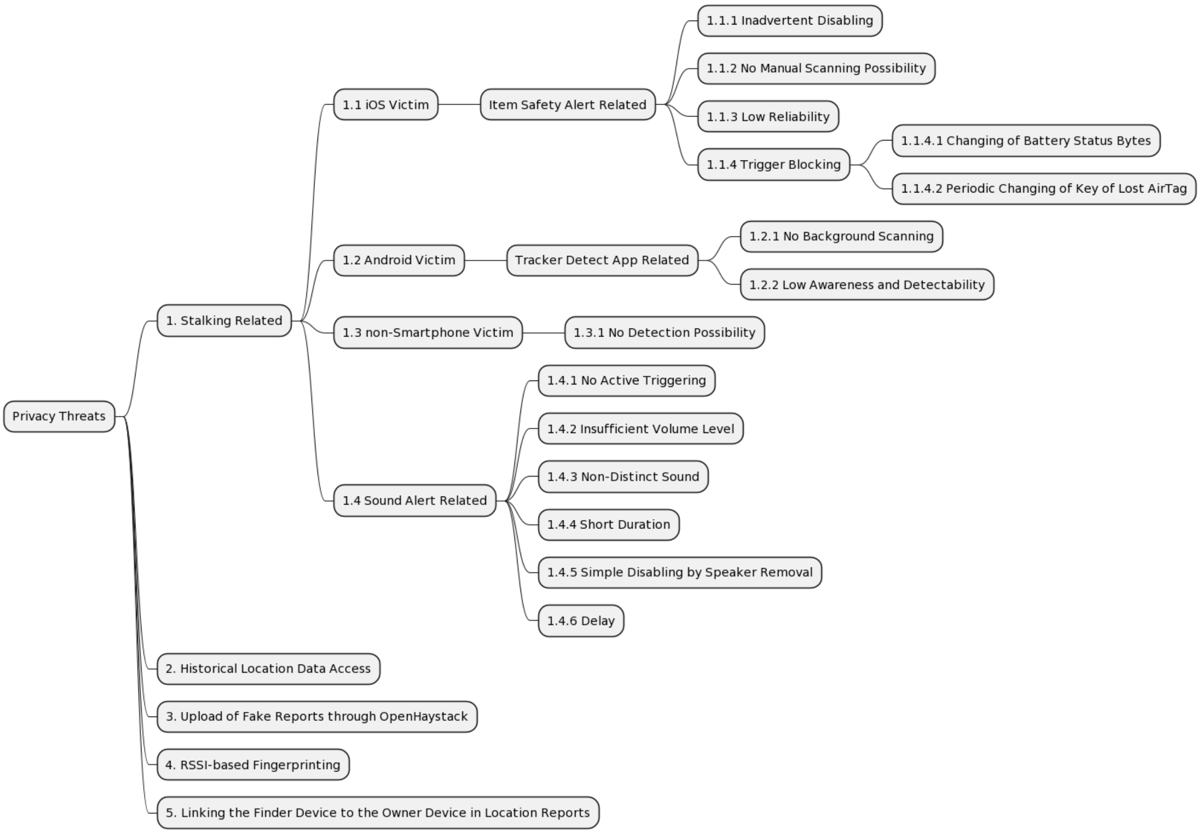
Privacy Threats
1. Stalking Related
1.1 iOS Victim
Item Safety Alert Related
1.1.1 Inadvertent Disabling
1.1.2 No Manual Scanning Possibility
1.1.3 Low Reliability
1.1.4 Trigger Blocking
1.1.4.1 Changing of Battery Status Bytes
1.1.4.2 Periodic Changing of Key of Lost AirTag
1.2 Android Victim
Tracker Detect App Related
1.2.1 No Background Scanning
1.2.2 Low Awareness and Detectability
1.3 non-Smartphone Victim
1.3.1 No Detection Possibility
1.4 Sound Alert Related
1.4.1 No Active Triggering
1.4.2 Insufficient Volume Level
1.4.3 Non-Distinct Sound
1.4.4 Short Duration
1.4.5 Simple Disabling by Speaker Removal
1.4.6 Delay
2. Historical Location Data Access
3. Upload of Fake Reports through OpenHaystack
4. RSSI-based Fingerprinting
5. Linking the Finder Device to the Owner Device in Location Reports

Figure 5.2: AirTag-related Threats

To conduct the DREAD risk assessment, some assumptions on the different factors and threat scores have to be made:

- **Threat Scores**: While DREAD assumes a scale of one through ten, there have been debates [123] on the significance of little score differences (e.g., the difference between a discoverability score of six versus seven). [123] recommends simplifying risk levels 1-3, signifying a low, medium, and high rating. Recent research has also applied this adaptation [121]. Consequently, to avoid comparisons with small margins, this will also be applied in this risk assessment. Additionally, the sums will be calculated instead of the average threat scores to show more distinct results.
- **Affected Users**: Each threat primarily affects individual victims rather than a widespread user database. Therefore, in this assessment, *Affected Users* refers to the number of users that could be affected by a threat based on the OS of the devices in use[6]. This assumption applies to the stalking-related threats (1.). For non-stalking-related threats (2. - 5.), a separate DREAD risk assessment will be held, where a high (3) *Affected Users* value refers to a high percentage or even an entire population of AirTag users. A low (1) score on the other hand signifies a low percentage of AirTag users affected by a threat.
- **Damage**: The damage caused by each threat varies a lot. Damage can be inflicted by breaching user privacy which can cause reputational harm, discrimination, physical violence, and emotional distress among others [24]. It is difficult to compare these with each other, determining which one outweighs the other. For this reason, any threat relating to the revelation of sensitive data is automatically assigned the threat value 3[7].
- **Reproducibility**: As all threats can be reproduced easily, all are automatically assigned a high (3) threat score.
- **Discoverability**: To score the likelihood of a threat being exploited, it is compared with the reported cases and how many malicious users have abused the threat.

By conducting the DREAD risk assessment, each threat has been analyzed and given a threat score. Tables 5.4 and 5.5 show the output of DREAD with a threat score assigned to each threat. The range of threat scores ranges from a minimum of 9 and a maximum of 15[8].

Table 5.4 shows that threats 1.2.1, 1.2.2, 1.2.3, and 1.4.5 have the highest threat scores. Threat 1.2.1 has a score of 15, and the others have a value of 14. Apple should give these threats the highest attention and work on providing sufficient mitigation tactics. On the

[6] The global OS smartphone marketshare [124] is taken for reference. This is dominated by Android and Apple with respective marketshares of close to 70% and 29%.

[7] Any of the analyzed threats can eventually lead to the unwanted sharing of personal location data. This is categorized as sensitive data, therefore all threats are assigned a high (3) threat score.

[8] With the assumptions that Reproducibility and Damage are both assigned a high (3) value:
Min: 3 (D) + 3 (R) + 1 (E) + 1 (A) + 1 (A) = 9.
Max: 3 (D) + 3 (R) + 3 (E) + 3 (A) + 3 (A) = 15.

Threat	Damage	Repro-ducibility	Exploit-ability	Affected Users	Discover-ability	Total
1.1 Stalking of iOS Victims						
1.1.1 ISA: Inadvertent Disabling	3	3	2	2	2	**12**
1.1.2 ISA: No Manual Scanning Possibility	3	3	2	2	3	**13**
1.1.3 ISA: Low Reliability	3	3	2	2	2	**12**
1.1.4.1 ISA Trigger Blocking: Battery Status Byte Change	3	3	1	2	1	**10**
1.1.4.2 ISA Trigger Blocking: Periodic Key Change	3	3	1	2	1	**10**
1.2 Stalking of Android Victims						
1.2.1 Tracker Detect: No Background Scanning	3	3	3	3	3	**15**
1.2.2 Tracker Detect: Low Awareness/Detectability	3	3	2	3	3	**14**
1.3 Stalking of non-smartphone victims						
1.3.1 Non-Smartphone: No Detection Possibility	3	3	3	2	3	**14**
1.4 Sound Alert related Threats						
1.4.1 No Active Triggering	3	3	2	3	2	**13**
1.4.2 Insufficient Volume Level	3	3	2	3	2	**13**
1.4.3 Non-Distinct Sound	3	3	2	3	2	**13**
1.4.4 Short Duration	3	3	2	3	2	**13**
1.4.5 Simple Disabling by Speaker Removal	3	3	3	3	2	**14**
1.4.6 Delay	3	3	2	3	2	**13**

Table 5.4: DREAD risk assessment on AirTag-related stalking threats

Threat	Damage	Repro-ducibility	Exploit-ability	Affected Users	Discover-ability	Total
Other Threats						
2. Historical Location Data Access	3	3	1	2	1	**10**
3. Upload of Fake Reports through OpenHaystack	3	3	1	2	1	**10**
4. RSSI Fingerprinting	3	3	1	3	1	**11**
5. Linking the Finder Device to the Owner Device	3	3	1	2	1	**10**

Table 5.5: DREAD risk assessment on other AirTag-related threats

other hand, Table 5.5 shows the more specific threats and their respective rankings. They are in a different table, as their assessment regarding the category *Affected Users* follows a different assumption. Nonetheless, by analyzing their total threat scores, it becomes evident that these are comparatively low-priority threats, close to the minimum value of 9. In the following sections, the following threats 1.2.1, 1.3.1, and 4 are analyzed, and their respective threat scores concerning the criteria *Exploitability*, *Affected Users* and *Discoverability* are explained.

Threat Risk Analysis: 1.2.1 Tracker Detect - No Background Scanning

The *No Background Scanning* threat refers to Apple's detection app *Tracker Detect* for Android users. The *Tracker Detect* app manually scans for nearby Apple *Find My* devices. Critically, it does not have a background scanning feature, as Apple claims, it would lead to too high battery usage [115]. The risk assessment gave it an *Exploitability* score of 3. This is because, in a stalking case, no technical knowledge is required. A stalker must be aware of the *Tracker Detect* app and its limited capabilities to successfully exploit the fact that no background scanning is conducted. After all, a victim with no stalking suspicions would not independently trigger a manual scanning for no reason. Concerning the *Affected Users* category, a high (3) score is also assigned. According to [124], close to 70% of worldwide Smartphones run the Android OS. This leads to a high likelihood that a victim is an Android victim. A high (3) *Discoverability* is scored. This is accredited to the high number of reports made by Android stalking victims.

Threat Risk Analysis: 1.3.1 Non-Smartphone Victim - No Detection Possibility

Victims without a smartphone are more susceptible to stalking, as they do not have the opportunity to rely on any detection system. This has been the case in stalking reports, where stalkers placed AirTags inside children's backpacks or shoes [24]. Additionally, the elderly could be targeted, as many have not adopted the use of smartphones. A high (3) *Exploitability* score is assigned, as stalking with AirTags is trivial. No technical expertise is required, especially when dealing with more susceptible types of people like children or the elderly. A medium (2) score in the *Affected Users* implies that the likelihood of stalking on non-smartphone users is possible, yet not that frequent. This is shown by [125], as in 2022, 68% of the worldwide population owns a smartphone. [24] shows that there have been an extensive amount of stalking reports where especially children have become the victims of stalking. In many cases, the stalkers aimed to locate the parent of the child by collecting the child's location data and whereabouts. For this reason, a high (3) *Discoverability* score has been assigned.

Threat Risk Analysis - 4 RSSI Fingerprinting

RSSI-based Fingerprinting is a recently discovered threat, where different BLE traces emitted by the same device can be linked together despite MAC address randomization as a countermeasure. [101] proved this and obtained a 99% re-identification accuracy of a device inside a pool of 30 motionless devices. To put these findings in perspective, [101] states that the identification success rate decreases against mobile targets. Scenarios involving AirTags are moving scenarios, as the user constantly moves around with its AirTag or smartphone. Nonetheless, RSSI Fingerprinting can still be considered a threat to AirTags, as there still is the possibility of linking different BLE traces to the same emitting device. A low (1) *Exploitability* score is assigned, as a high level of expertise on BLE and RSSI values is required to conduct a successful fingerprinting attack. The potentially *Affected Users* are any users with an AirTag, therefore referring to the entire AirTag population, which allows for a high (3) *Affected Users* score. While there has been research on *(RSSI-) Fingerprinting*, there have been no reported attacks using this exact method. Therefore, there is a low likelihood of it being exploited currently, which is the reason for its low (1) *Discoverability* score.

Conclusively, using the DREAD risk assessment model, the different threats identified in this thesis are assessed and given a total threat score. Some assumptions and adaptations from Microsoft's initial DREAD risk assessment proposal were made to ensure that the model applies to the AirTag use case. The output can be viewed on a relative scale, assigning a higher priority to threats with higher DREAD scores and a respective lower priority to threats with lower DREAD scores.

5.5 Threat Privacy Classification and Evaluation

With the completed DREAD risk assessment, evaluating the severity of the AirTag's non-adherence to certain privacy requirements becomes integral. All corresponding threats are listed in Table 5.6 for each requirement. Hereby, it becomes evident that there is a heavy bias towards *Awareness* directed threats. With a total of 20 threats identified, 14 of them breach the *Awareness* requirement.

There are several aspects to consider when evaluating how severely each privacy requirement is affected. There are two main values of importance: The DREAD score of each threat and the number of threats concerning a privacy requirement. Considering each without the other leads to a non-conclusive result. For example, in a quantitative approach, one could take an average DREAD score for each privacy requirement and then rank the severity with a high DREAD score leading to a high value. This could lead to the fact that a privacy requirement with only one high DREAD score threat to its name could outrank a privacy requirement with ten low DREAD score threats. On the other hand, if a privacy requirement has ten out of 20 threats assigned to it, yet these would all be threats with a very low DREAD score, they could outrank another privacy requirement with eight high DREAD scores.

Privacy Requirement	Threat	Avg. DREAD	Amount	Severity
Awareness	1.1.1, 1.1.2, 1.1.3, 1.1.4.1, 1.1.4.1, 1.2.1, 1.2.2, 1.3.1, 1.4.1, 1.4.2, 1.4.3, 1.4.4, 1.4.5, 1.4.6	12.8	14	3
Transparency				0
Confidentiality	2	10	1	1
Accountability				0
Data Minimization				0
Accuracy	3	10	1	1
Storage Limitation				0
Lawfulness				0
Anonymity	4, 5	10.5	2	2
Unlinkability	4, 5	10.5	2	2
Unobservability				0
Good Faith				0
Choice				0
Access and Correction				0

Table 5.6: Threat Privacy Requirement Classification

Therefore, a combination of the two values must be considered when evaluating the severity. Finally, it is assessed on a scale of 0 to 3. A severity score of 0 shows no impact, while a high score of 3 indicates the highest severity level. Following the results are displayed:

- (3) High Severity: *Awareness*
- (2) Medium Severity: *Anonymity*, *Unlinkability*
- (1) Low Severity: *Confidentiatlity*, *Accuracy*
- (0) No Impact: *Transparency*, *Accountability*, *Data Minimization*, *Storage Limitation*, *Lawfulness*, *Unobservability*, *Good Faith*, *Choice*, *Access and Correction*

5.6 Threat Classification Tree

Taking Figure 5.2 as an inspiration and design, a simple classification tree in Python was created[9]. The idea is to explore the flaws in the AirTag system, such as BLE, iOS's ISA feature, or the sound alert feature. Each threat is assigned an information part, which describes the flaw. Additionally, a concrete example is given, together with the privacy requirements, that the threat breaches. As this thesis heavily focuses on literary review and evaluation, implementing the threat classification tree serves a purely explorative purpose. By viewing different threats, a history is saved and presented after the threat classification tree analysis is done.

[9] https://github.com/dominic1712/AirTagThreatClassification

Someone unfamiliar with the problems of the AirTag could use it to determine if its privacy-related aspects are met. Similarly, the threat classification tree could demonstrate the issues a new tracker manufacturer must be aware of when placing a novel tracker inside Apple's *Find My* network. In conclusion, the developed threat classification tree provides a tool for exploring the vulnerabilities discovered by the AirTag system. It offers insights into privacy breaches and serves as a guide for assessing existing trackers and informing the development of new devices within Apple's *Find My* network.

Chapter 6

Discussion

The following sections discuss the evaluation and results of this thesis. The discussion will range from the set of privacy requirements, the use cases, and their application to the privacy requirements to analyzing the lawsuit with the set of privacy threats. It ends with the DREAD Risk Assessment, the output of which will be thoroughly analyzed regarding its broader relevance.

6.1 Privacy Requirements Mapping

The four frameworks considered for this thesis can be divided into legal (GDPR and nFADP) and technical (COPri V.2, NIST) frameworks. They differ quite a bit, as the legal frameworks focus on protecting a data subject and providing a legal basis for it. The technical frameworks focus more on a data handler's point of view to conceptualize its data subject protection mechanisms to simplify the implementation and prioritization of adequate privacy-preserving measures. However, they have a main commonality, the privacy requirements, which allow for comparison between them.

While many privacy frameworks exist, these four were chosen because of their temporal relevance and importance to the topic. The GDPR is one of the most critical data regulations, and it came into effect in 2016 as one of the strictest privacy frameworks, which laid the basis for many legal and technical frameworks published afterward. As this thesis is written at the University of Zurich, the nFADP, Switzerland's legal privacy framework, is particularly important. Reworked and republished in 2023, it contains the most recent privacy-related developments and shows minor differences from the GDPR.

Apple is a US-based company. Therefore, including a US-based framework with the NIST privacy framework makes sense. The US has many state-wide legal frameworks, such as [88], which all differ in strictness and applicability. For this reason, a nationwide accepted framework was elected, which instead focuses on providing a complete set of privacy guidelines (NIST's subcategories). Additionally, with Gharib's COPri V.2, a set of privacy requirements are evaluated, which should be addressed during the design phase of a system. This thesis revolves around the Apple AirTag and its respective deficiencies, many of which could have been avoided in the design phase. Therefore, adding COPri V.2 as the fourth privacy framework analyzed in this thesis makes sense.

The initial privacy requirements mapping results show twelve privacy requirements (Table 4.1), displaying which framework contains which privacy requirement. Especially regarding NIST's privacy framework, mapping the requirements on top of NIST's subcategories was difficult. This is mainly due to NIST's ambiguity, where specific privacy requirements are mentioned as goals of categories, yet no respective subcategory covers the requirement. These implications are evaluated as partial overlappings with the privacy requirement.

After evaluating the set of privacy requirements, it stands out that the only requirement missing is enabling the data subject to access their data. The data subject should be able to look into the collected data and adjust preferences accordingly. This is added to the set of privacy requirements after the mapping with APEC's privacy framework to the privacy requirements set. With the inclusion of APEC's privacy framework, there is a geographical full circle, as the significant economic driving continents, North America (NIST), Europe (GDPR/nFADP), and Asia (APEC), are all included in the mapping. As a result, in an *Own Requirements Framework*, a complete set of privacy requirements is contained as depicted by Table 6.1.

Privacy Requirement	Own Framework
(1) Awareness	✓
(2) Transparency	✓
(3) Confidentiality	✓
(4) Accountability	✓
(5) Data Minimization	✓
(6) Accuracy	✓
(7) Storage Limitation	✓
(8) Lawfulness	✓
(9) Anonymity	✓
(10) Unlinkability	✓
(11) Unobservability	✓
(12) Good Faith	✓
(13) Choice	✓
(14) Access and Correction	✓

Table 6.1: Own Requirements Framework

There has been research focusing on comparing different legal and technical privacy-focused frameworks. Yet, this thesis is the first to explore the commonalities and differences between selected frameworks from different geographical, cultural, legal, and technical backgrounds with an increased focus on the defined privacy requirements. For example, many legal aspects of the GDPR and nFADP were not analyzed. Similarly, the technicalities of implementing COPri V.2 and NIST's Privacy Framework were briefly summarized. This was elected to focus on the overlapping definitions between the different frameworks. The privacy requirements are a small yet defining aspect of how a particular framework sees privacy and where the focus should be. This shift between the technical and legal frameworks becomes evident in Table 5.2 where there are primarily overlappings between COPri and NIST (technical) and the GDPR, nFADP, and APEC (legal) frameworks. Conclusively, the *Own Framework* resulting from this thesis contains both legal and technical privacy requirements from various frameworks. It is a complete set consisting of 14 different requirements, of which all are equally important.

6.2 Use Case Analysis with Privacy Requirements

The four use cases only represent a small portion of what is possible using the AirTag. Yet, they provide insight into the most popular (mis-)uses of the AirTag. Recent research has focused on exploring the vulnerabilities of AirTags and then placing the vulnerabilities into certain niche use cases (e.g., linking and identifying protestors through OF advertisements [44]). The approach taken by this thesis is to explore the different threats discovered by research and place them inside different use cases. With this, the threats are categorized according to different privacy requirements. Table 6.2 shows the mapping of different research paper findings and the privacy requirements.

Paper Title, Year, Reference	**Threat**	**Privacy Requirement**
Who Tracks The Trackers? 2021, [12]	ISA Trigger Blocking: Changing of Battery Status Bytes	**(1) Awareness**
Who Tracks The Trackers? 2021, [12]	ISA Trigger Blocking: Periodic Changing of Key of a Lost AirTag	**(1) Awareness**
Who Can Find My Devices? 2021, [44]	Historical Location Data Access	**(3) Confidentiality**
Track You: A Deep Dive into Safety Alerts for Apple AirTags 2023, [47]	Upload of Fake Reports through OpenHaystack	**(6) Accuracy**
Who Can Find My Devices? 2021, [44]	Linking the Finder Device to the Owner Device in Location Reports	**(9) Anonymity, (10) Unlinkability**
RSSI-based Fingerprinting of Bluetooth Low Energy Devices 2023, [101]	RSSI-based Fingerprinting	**(9) Anonymity, (10) Unlinkability**

Table 6.2: AirTag-related threats categorized by Privacy Requirement

The use case analysis contained information gathered from three separate sources: (1) Official publications and documentation by Apple, (2) Research papers, and (3) the lawsuit *Hughes et al versus Apple*. This presents insights into the manufacturer's perspective, the actual users of AirTags or victims of AirTag stalking perspective, and the scientific perspective of researchers. While the personal experiences made by the plaintiffs in the lawsuit convey a clear image, the integrity of its scientific references to technological

AirTag-related aspects is questionable. This is mainly due to the poor choice of references made in the lawsuit, which only consists of new articles. Therefore, all technological statements were cross-referenced with actual scientific literature to provide a coherent analysis of the assertions made by the lawsuit. With this, it turns out that issues like the *AirTag Firmware Updates* or *AirTag Identifier Reset* are wrongly criticized, as it was possible to find contradicting information in scientific literature. However, the inclusion of the lawsuit itself is critical, as it provides a deep dive into the experiences made by AirTag stalking victims, which is the biggest threat associated with Apple's AirTag.

Generally, most issues discovered in the use case analysis for each privacy requirement can be seen as threats to the AirTag system. This applies to threats like the disabling of iOS ISAs or the issues related to the AirTag's sound alerts. However, the threat risk assessment does not represent the findings in *Good Faith*. This is because the issues in *Good Faith* are not direct threats to the AirTag system. Apple is unwilling to aid stalker victims and prefers to protect its customers. These are not direct threats but simply examples of generic non-adherence of *Good Faith*. Therefore, no *Good Faith* related threats are considered in Section 5.4.3.

Nonetheless, Apple's non-adherence to *Good Faith* is an alarming finding within this thesis. With several marketing campaigns during the product roll-out phase of the AirTag, Apple successfully influenced the general public's opinion. By spreading fake information and going as far as calling the AirTag *Stalker-Proof*, Apple neglected its role as a data collector and processor of sensitive data. This, along with inadequate automatic detection methods and the protection of stalkers by restricting access to stalker-identifying data, presents a convincing argument that Apple violates the privacy requirement of Good Faith.

6.3 DREAD Risk Assessment and Severity Scale

To conduct the DREAD risk assessment, several assumptions had to be made about the model to increase its applicability to the use case at hand. Similar assumptions have been produced by recent IoT-related work [121], yet by changing the *Affected Users* to the potentially affected users and considering the global smartphone OS market, this thesis adopts a new version of DREAD. The threats are further divided into separate risk assessment categories. This is because of the difference in the *Affected Users*. While stalking-related threats target non-AirTag owners and, therefore, an uncertain number of users, the other threats are aimed at the owners of AirTags. The users affected by the different types of threats are consequently categorized into separate groups.

Tables 5.4 and 5.5 display the results of the risk assessments. Due to the assumptions (Damage and Reproducibility automatically assigned the high (3) value), the scores range from 9 to 15. The differences in scores, therefore, come from varying DREAD scores in the categories: *Exploitability*, *Affected Users*, and *Discoverability*. Out of these, *Discoverability* is challenging to evaluate. This is due to many stalking-related incidents never being reported to the police, possibly because they are not found, or the victim does not know what to do with it. Another reason why *Discoverability* is challenging to quantify is because only a small number of the reported AirTag-related stalking cases are published.

This was discovered by *Motherboard* in an investigation in April 2022, where they found that in a year since the AirTag's release, over 150 reports have amassed [109].

The risk assessment gives the threats a quantitative value, facilitating the comparison between single threats. This aids in identifying the worst threats from Apple's point of view, which should be given the highest priority in solving. However, the DREAD scores should not be taken out of the context of this risk assessment. They serve a comparative manner, and for example, a low threat score of 10 does not directly imply a low-risk threat in general. Instead, in the set of threats provided, it can be viewed as a lower-priority threat, given that most threats achieved a higher threat score. Therefore, the reach and depth of this DREAD risk analysis must be put into perspective. This analysis serves its purpose, as it coherently analyzes the list of threats found and quantifies them by giving them a risk score. If compared with other threats that are not included in this risk assessment, a new DREAD risk assessment with possible adjustments to the assumptions must be conducted.

Each threat was further categorized into one of the 14 privacy requirements. The outcome, Table 5.6, shows that 70% of the identified threats breach the *Awareness* privacy requirement. On the one hand, it makes sense, as stalking is the biggest AirTag-related threat existing to the point of writing this thesis, and most of the threats that breach *Awareness* are closely related to stalking purposes. However, some stalking-related threats are very similar to each other. Significantly, the sound-alert threats could technically be viewed as one threat, as they all base themselves on some fundamental flaw in the sound-alert feature. Furthermore, a severity scale is created to quantify how much a privacy requirement is affected by threats. It takes both the amount of threats and an average DREAD score as input and displays the severity on a scale of no (0), low (1), medium (2), and high (3) severity. This evaluation guides toward identifying the breached privacy requirements and categorizing how heavily each one is affected.

This summarizes the findings in the entire thesis, taking the privacy requirements explored in Section 4.1 for reference and categorizing the threats discovered in literature and the lawsuit using the DREAD risk assessment model. To this day, no similar work has been done. The privacy requirements in different frameworks define ground rules that must be followed. Work has been done on analyzing privacy requirements in IoT devices, yet most of these add focus on security requirements. This thesis is the first paper to explore several technical and legal frameworks on their defined privacy requirements comparatively. Several use cases were tested by evaluating the AirTag system, and their results show some concerning aspects. Despite several concerns raised by respectable voices inside the technical community, Apple successfully launched a flawed and dangerous product that can seriously endanger the lives of any human being if used for malicious reasons. Apple boasts of certifications like the APEC accountability certification, which shows its adherence to the APEC privacy framework. However, the experiences made by several plaintiffs in [24] show where Apple's priorities lie: *Profit maximization with minimal concern for the experiences made by victims of their products.* This has led to the AirTags' exploitation for malicious purposes, including theft, stalking, and even murder.

6.4 AirTag Privacy Classification Tree

The privacy classification tree visually represents a walk-through of the existing AirTag-related privacy threats. While the mindmap in Figure 5.2 was taken as an inspiration, the tree follows a different classification, allowing all threats to be categorized. This is achieved by an initial classification into the following six categories:

1. iOS-related AirTag issues
2. Android-related issues
3. Sound-related issues
4. BLE-related issues
5. macOS-related issues
6. FindMy Authentication-related issues

The first three categories are taken from the mindmap, yet the others initialize a different categorization. With BLE-, macOS-, and FindMy Authentication-related issues, the categorization process is more granular, allowing for a more detailed classification of specific threats. Generally, the classification tree framework aims to serve as a tool for future research and development within the AirTag ecosystem. It presents threats in an explorative manner, facilitating the analysis of threats on the AirTag system from a third party's perspective. To the point of writing this thesis, no classification tree exists on the existing AirTag threats. The prototype proposed in this thesis is the first to gather AirTag threats and display them exploratively.

6.5 Omittance of RSSI-Values Experiment

In Section 4, a design for gathering RSSI values in a controlled environment was proposed. This was omitted from the thesis for the following reasons. A decision was made to focus on a profound threat analysis by incorporating additional steps, such as analyzing the lawsuit and delving into the APEC privacy framework. These were initially not planned during the design stage; the entire lawsuit was published after completing the background section. Awareness of APEC's privacy framework arose while evaluating the use cases.

The RSSI values experiment would have moved outside the scope of the AirTag threats and privacy requirements, which is another reason why it was omitted. Instead, the threat classification tree in Section 5.6 was designed. With this, a full circle was achieved by placing the discovered threats within the set of privacy requirements discovered.

Chapter 7

Conclusions and Future Work

The following chapter summarizes the topics, findings, and conclusive remarks discovered in this thesis and answers the initially proposed research questions. Further, it outlines future work to be conducted in this research area.

7.1 Conclusions

This paper presents a cohesive analysis of the AirTag system and places it in a broader context within the privacy landscape. By analyzing threats and mapping them to a corresponding privacy requirement, an initial framework for AirTag threat and risk analysis within the vast privacy landscape is proposed.

7.1.1 Commonalities and Differences of Privacy Requirements in Privacy Frameworks

Table 5.2 shows the differences and commonalities regarding the defined privacy requirements by analyzing and comparing various privacy frameworks. This summarizes the answer to RQ1: Although the results are not that evident due to the NIST privacy framework having an almost complete set of privacy requirements, there is a distinction that can be made between the technical and legal frameworks. While all set a primary focus on *Awareness* and *Transparency*, the legal frameworks have an increased focal point on requirements like *Accuracy*, *Storage Limitation*, and *Lawfulness*. It makes sense that these are contained within the legal frameworks as they all promote responsible data handling practices by ensuring data integrity and focusing on individuals' rights.

The technical frameworks instead lay an increased focus on the privacy requirements *Anonymity*, *Unlinkability*, and *Unobservability*. These can all be implemented with the technical know-how of the technologies used within the system. Adding to these apparent differences between the analyzed frameworks, some outlying privacy requirements are only contained within a single privacy requirement. These include aspects like *Good Faith* (nFADP) or *Choice* (APEC). These distinctions show how various frameworks define the scope of privacy. They all have common ground, yet certain differences stand out when delving deeper into more specific definitions of privacy.

7.1.2 Apple's adherence to the Privacy Requirements

With an extensive analysis of the privacy requirements and to what extent Apple abides by them, concerns regarding specific privacy requirements are raised. The results of the use case analysis and risk assessment show that Apple's AirTag system breaches the following privacy requirements to a certain extent: (1) *Awareness*, (2) *Transparency*, (3) *Confidentiality*, (6) *Accuracy*, (9) *Anonymity*, (10) *Unlinkability*, and (12) *Good Faith*. Especially, *Awareness* and *Good Faith* are worrying. *Awareness* achieved the highest (3) severity score possible in the severity assessment due to its high average DREAD score and the large number of threats targeting it. *Good Faith* did not undergo the DREAD risk assessment, as no specific threats were breaching it. However, with the analysis of the lawsuit [24], the experiences made by several plaintiffs display how Apple does not abide by *Good Faith*. Restrictive data policies complicating stalking investigations and marketing campaigns conveying a false sense of security in the AirTag product show evident defiance in *Good Faith*.

These results answer RQ2, showing where Apple has to focus on improving its privacy mechanisms. Apple implemented common practice safeguards like MAC randomization, which protect the AirTag's system to a great extent. Yet, as the AirTag is a novel product with highly advanced technologies, there has been a lot of research on it, pointing out flaws and strengths within its system. Simultaneously, Apple is regularly improving its privacy mechanisms. Therefore, the list of threats on the AirTag system is constantly changing, with new threats being identified and older ones being neutralized regularly.

7.1.3 BLE Tracker Detection Improvements

Many threats were identified within this thesis, which displays the wide range of attacking possibilities on the AirTag system. However, with an analysis of Apple's tracker detection methods and several research papers focusing on improving it [7], [16], [81], a numerous amount of improvements to the current ISA feature and tracker detect app are proposed. These are the initial answers to RQ3. With Google's implementation of the UTA in July 2023, some improvements to the ISA, such as active scanning, were introduced. The collaboration between Google and Apple is still an ongoing project, with Apple reportedly handing in the first version of the industry-wide specification to the IETF in December 2023 [126]. An overarching industry-wide standard would greatly improve the detection mechanisms of BLE trackers.

Detection of BLE trackers should not rely solely on automatic detection methods of BLE devices. Research into real-life scenarios has shown that auditory cues are necessary as a complementary detection method. Especially considering that the most susceptible people prone to stalking are children and the elderly. Children are usually unaware of the malevolence of other people. Many of the elderly are unaware of the capabilities of novel technologies and, therefore, would suspect less of an unknown AirTag moving along with them. Most importantly, many children and a large amount of the elderly do not have smartphones. Their primary detection methods of unknown BLE trackers are auditory hints. Other hints may include visual cues like a blinking flashlight implemented within the tracker. This feature is used by the Pebblebee clip. Therefore, in addition to improving OS-based tracker detection methods, a second focus should be laid on improving the auditory and visual features of modern trackers. Conclusively, these findings answer RQ3 and highlight two aspects of tracker detection mechanisms that can improve user privacy significantly.

7.2 Future Work

With the threats directed towards AirTag constantly changing, the list of threats discovered in this thesis derives from a snapshot of the current AirTag-related threats. This list will change with new threats being discovered by researchers and Apple users and simultaneously, old threats being neutralized with improvements to the system, like the industry-wide tracker detection specification. Therefore, future work entails delving into the AirTag system and aggregating future threats.

Similarly, privacy frameworks are also constantly evolving. They are being adapted to changes in technology, and considering that the IoT area is one of the most promising with a huge impact on our daily lives, there is a lot of ongoing work in adapting privacy frameworks to the developments in IoT technology. Future work should capture these changes and add them to the framework developed within this thesis. Especially if there are changes directed toward the inclusion of novel privacy requirements.

Furthermore, the threat classification tree designed in this thesis is only the first prototype. Several improvements can be made. The tree could be expanded to include additional classification layers in a future implementation, providing a higher granularity and a more comprehensive approach to threat categorization. However, this necessitates a higher number of threats to be categorized. Currently, the tree is written as a terminal-based input/output Python code. The functionality of the tree could be enhanced. An example would be the implementation of a machine learning algorithm that could accurately define the threat based on a set of input factors.

Lastly, the experiments designed in Section 4.3 could be carried out in future work. Its results are significant in the context of improved tracker detection methods. By gathering many RSSI values at various distances, it could be possible to accurately determine the distance between BLE-advertisement emitting and receiving devices. This could be achieved through a machine learning algorithm. Ultimately, this could be applied in real-life scenarios, as malicious trackers moving along with a stalker's victim (e.g., in its backpack) could be detected earlier based on the RSSI values.

Bibliography

[1] H. Smith, “Ed Tuck, investor who made GPS usable for boaters and Bubbas, dies at 85,” *Washington Post*, Apr. 2023, last visit 16 of March, 2024. [Online]. Available: https://www.washingtonpost.com/local/obituaries/ed-tuck-investor-who-made-gps-usable-for-boaters-and-bubbas-dies-at-85/2017/07/05/69e96744-618e-11e7-a4f7-af34fc1d9d39_story.html

[2] S. Wolpin, “Commerical GPS Turns 25: How the Unwanted Military Tech Found Its True Calling,” May 2014, last visit 16 of March, 2024. [Online]. Available: https://mashable.com/archive/commercial-gps-25-anniversary

[3] B. Xiao, K. Zhang, R. Grenfell, and T. Norton, “Handheld gps–today and tomorrow,” in *FIG XXII International Congress Washington, DC USA*, 2002.

[4] “Find Your Everything,” last visit 16 of March, 2024. [Online]. Available: https://chipolo.net/en/

[5] “Tile Trackers | Bluetooth Trackers for Keys, Wallets, Pets, and More,” last visit 16 of March, 2024. [Online]. Available: https://www.tile.com

[6] G. Celosia and M. Cunche, “Saving private addresses: An analysis of privacy issues in the bluetooth-low-energy advertising mechanism,” in *Proceedings of the 16th EAI International Conference on Mobile and Ubiquitous Systems: Computing, Networking and Services*, 2019, pp. 444–453.

[7] A. Heinrich, N. Bittner, and M. Hollick, “Airguard-protecting android users from stalking attacks by apple find my devices,” in *Proceedings of the 15th ACM Conference on Security and Privacy in Wireless and Mobile Networks*, 2022, pp. 26–38.

[8] C. Garg, A. Machiry, A. Continella, C. Kruegel, and G. Vigna, “Toward a secure crowdsourced location tracking system,” in *Proceedings of the 14th ACM Conference on Security and Privacy in Wireless and Mobile Networks*, 2021, pp. 311–322.

[9] “AirTag,” last visit 16 of March, 2024. [Online]. Available: https://www.apple.com/airtag/

[10] T. Roth, F. Freyer, M. Hollick, and J. Classen, “Airtag of the clones: Shenanigans with liberated item finders,” in *2022 IEEE Security and Privacy Workshops (SPW)*. IEEE, 2022, pp. 301–311.

[11] J. O'Malley, "The imminent ultra-wideband revolution: After decades in the lab, big tech has finally found a use for a unique wireless technology," *Engineering & Technology*, vol. 16, no. 9, pp. 1–4, 2021.

[12] T. Mayberry, E. Fenske, D. Brown, J. Martin, C. Fossaceca, E. C. Rye, S. Teplov, and L. Foppe, "Who tracks the trackers? circumventing apple's anti-tracking alerts in the find my network," in *Proceedings of the 20th Workshop on Workshop on Privacy in the Electronic Society*, 2021, pp. 181–186.

[13] G. A. Fowler, "Review | Apple's AirTag trackers made it frighteningly easy to 'stalk' me in a test," *Washington Post*, May 2021, last visit 16 of March, 2024. [Online]. Available: https://www.washingtonpost.com/technology/2021/05/05/apple-airtags-stalking/

[14] J. Koetsier, "How To Track People With Apple AirTags," last visit 16 of March, 2024. [Online]. Available: https://www.forbes.com/sites/johnkoetsier/2021/04/22/how-to-track-people-with-apple-airtags/

[15] I. C. Campbell, "Apple's new Find My feature could let you know if you're the one being tracked," Mar. 2021, last visit 16 of March, 2024. [Online]. Available: https://www.theverge.com/2021/3/4/22313659/apple-find-my-app-item-safety-alerts-prevent-stalking

[16] K. O. Müller, L. Bienz, B. Rodrigues, C. Feng, and B. Stiller, "Homescout: Anti-stalking mobile app for bluetooth low energy devices," in *2023 IEEE 48th Conference on Local Computer Networks (LCN)*. IEEE, 2023, pp. 1–9.

[17] T. Mayberry, E.-O. Blass, and E. Fenske, "Blind my-an improved cryptographic protocol to prevent stalking in apple's find my network," *Proceedings on Privacy Enhancing Technologies*, 2023.

[18] "Apple AirTags Are the Best Luggage Trackers," Jul. 2023, last visit 16 of March, 2024. [Online]. Available: https://www.nytimes.com/wirecutter/blog/best-luggage-tracker-apple-airtag/

[19] S. Tabahriti, "Eric Adams is handing out hundreds of free Apple AirTags in a bid to cut the number of auto thefts in New York," last visit 16 of March, 2024. [Online]. Available: https://www.businessinsider.com/eric-adams-new-york-giving-away-apple-airtags-auto-thefts-2023-5

[20] "Apple AirTags and Bluetooth Trackers Are Officially a Billion-Dollar Industry - Here's What To Know, Trends, and the Best Ways To Invest," Dec. 2022, last visit 16 of March, 2024. [Online]. Available: https://finance.yahoo.com/news/apple-airtags-bluetooth-trackers-officially-175911565.html

[21] M. Gault, "Woman Allegedly Used Apple AirTag to Track and Kill Her Boyfriend," Jun. 2022, last visit 16 of March, 2024. [Online]. Available: https://www.vice.com/en/article/xgy8qz/woman-allegedly-used-apple-airtag-to-track-and-kill-her-boyfriend

[22] "An update on AirTag and unwanted tracking," last visit 16 of March, 2024. [Online]. Available: https://www.apple.com/newsroom/2022/02/an-update-on-airtag-and-unwanted-tracking/

[23] J. Stempel, "Apple is sued by women who say AirTag lets stalkers track victims," *Reuters*, Dec. 2022, last visit 16 of March, 2024. [Online]. Available: https://www.reuters.com/legal/apple-is-sued-by-women-who-say-airtag-lets-stalkers-track-victims-2022-12-06/

[24] "Hughes v Apple Amended Complaint 10-12-2023," Oct. 2023. [Online]. Available: https://cdn.arstechnica.net/wp-content/uploads/2023/10/Hughes-v-Apple-Amended-Complaint-10-12-2023.pdf

[25] "Apple, Google partner on an industry specification to address unwanted tracking," last visit 16 of March, 2024. [Online]. Available: https://www.apple.com/newsroom/2023/05/apple-google-partner-on-an-industry-specification-to-address-unwanted-tracking/

[26] "Tracker im Praxistest - Bei Hindernissen versagen viele Bluetooth-Tracker schnell," Sep. 2023, last visit 16 of March, 2024. [Online]. Available: https://www.srf.ch/sendungen/kassensturz-espresso/tests/gadgets-elektronik/tracker-im-praxistest-bei-hindernissen-versagen-viele-bluetooth-tracker-schnell

[27] "The Bluetooth Low Energy Primer," May 2022, last visit 16 of March, 2024. [Online]. Available: https://www.bluetooth.com/bluetooth-resources/the-bluetooth-low-energy-primer/

[28] "Topologie-Optionen | Bluetooth Technologie Website," last visit 16 of March, 2024. [Online]. Available: https://www.bluetooth.com/de/learn-about-bluetooth/topology-options/

[29] "Core Specification," Jul. 2021, last visit 16 of March, 2024. [Online]. Available: https://www.bluetooth.com/specifications/specs/core-specification-5-3/

[30] "rssi [Bluetooth LE Wiki]," last visit 16 of March, 2024. [Online]. Available: https://bluetoothle.wiki/rssi

[31] "Proximity and RSSI," Sep. 2015, last visit 16 of March, 2024. [Online]. Available: https://www.bluetooth.com/blog/proximity-and-rssi/

[32] D. Coppens, A. Shahid, S. Lemey, B. Van Herbruggen, C. Marshall, and E. De Poorter, "An overview of uwb standards and organizations (ieee 802.15. 4, fira, apple): Interoperability aspects and future research directions," *IEEE Access*, vol. 10, pp. 70 219–70 241, 2022.

[33] H.-J. Pirch and F. Leong, "Introduction to impulse radio uwb seamless access systems," *Proceedings of the Fraunhofer SIT ID: SMART Worksho, Darmstadt, Germany*, pp. 19–20, 2020.

[34] S. Aditya, A. F. Molisch, and H. M. Behairy, "A survey on the impact of multipath on wideband time-of-arrival based localization," *Proceedings of the IEEE*, vol. 106, no. 7, pp. 1183–1203, 2018.

[35] A. Zignani and S. Tomsett, "Ultra-wideband (uwb) for the iot–a fine ranging revolution," *ABI Research*, 2021.

[36] "FiRa Consortium," last visit 16 of March, 2024. [Online]. Available: https://www.firaconsortium.org/

[37] "How UWB Works | FiRa Consortium," last visit 16 of March, 2024. [Online]. Available: https://www.firaconsortium.org/discover/how-uwb-works

[38] T. Li, J. Liang, Y. Ding, K. Zheng, X. Zhang, and K. Xu, "On design and performance of offline finding network," in *IEEE INFOCOM 2023-IEEE Conference on Computer Communications.* IEEE, 2023, pp. 1–10.

[39] "iCloud+ - Find My," last visit 16 of March, 2024. [Online]. Available: https://www.apple.com/icloud/find-my/

[40] "SmartThings Find," last visit 16 of March, 2024. [Online]. Available: https://smartthingsfind.samsung.com/login

[41] J. Kastrenakes, "Apple says there are now over 1 billion active iPhones," Jan. 2021, last visit 16 of March, 2024. [Online]. Available: https://www.theverge.com/2021/1/27/22253162/iphone-users-total-number-billion-apple-tim-cook-q1-2021

[42] T. Haselton, "Here's how Apple's AirTag trackers compare to Tile, and why the company is so upset with Apple," Apr. 2021, last visit 16 of March, 2024. [Online]. Available: https://www.cnbc.com/2021/04/27/apple-airtags-versus-tile-tracker-how-they-compare.html

[43] "Apple's Find My network now offers new third-party finding experiences," last visit 16 of March, 2024. [Online]. Available: https://www.apple.com/newsroom/2021/04/apples-find-my-network-now-offers-new-third-party-finding-experiences/

[44] A. Heinrich, M. Stute, T. Kornhuber, and M. Hollick, "Who can find my devices? security and privacy of apple's crowd-sourced bluetooth location tracking system," *arXiv preprint arXiv:2103.02282*, 2021.

[45] H. Ibrahim, R. Asim, M. Varvello, and Y. Zaki, "I tag, you tag, everybody tags!" in *Proceedings of the 2023 ACM on Internet Measurement Conference*, 2023, pp. 561–568.

[46] "How to reset your AirTag," last visit 16 of March, 2024. [Online]. Available: https://support.apple.com/en-us/102577

[47] N. Shafqat, N. Gerzon, M. Van Nortwick, V. Sun, A. Mislove, and A. Ranganathan, "Track you: A deep dive into safety alerts for apple airtags," *Proceedings on Privacy Enhancing Technologies*, 2023.

[48] angorb, "[Bluetooth Company Identifiers] Company identifiers are unique numbers assigned by the Bluetooth SIG to member companies requesting one." last visit 16 of March, 2024. [Online]. Available: https://gist.github.com/angorb/f92f76108b98bb0d81c74f60671e9c67

[49] "Mark an AirTag or other item as lost in Find My on iPhone," last visit 16 of March, 2024. [Online]. Available: https://support.apple.com/guide/iphone/mark-an-item-as-lost-iph1b451b75f/ios

[50] Samsung, "Samsung Galaxy Unpacked January 2021: Official Replay," Jan. 2021, last visit 16 of March, 2024. [Online]. Available: https://www.youtube.com/watch?v=TD_BZN0bn_U

[51] "[Update] Introducing the New Galaxy SmartTag+: The Smart Way to Find Lost Items," Apr. 2021, last visit 16 of March, 2024. [Online]. Available: https://news.samsung.com/us/introducing-the-new-galaxy-smarttag-plus/

[52] "Samsung SmartThings Find Rapidly Expands With Over 300 Million Nodes Helping To Locate Devices," last visit 16 of March, 2024. [Online]. Available: https://bit.ly/494I0v2

[53] "Pebblebee Clip," last visit 16 of March, 2024. [Online]. Available: https://pebblebee.com/products/pebblebee-clip

[54] "Jiobit Smart Tag | Buy Jiobit GPS Tracker | Jiobit," last visit 16 of March, 2024. [Online]. Available: https://www.jiobit.com/product

[55] "100 things we announced at I/O 2023," May 2023, last visit 16 of March, 2024. [Online]. Available: https://blog.google/technology/developers/google-io-2023-100-announcements/

[56] "3 ways unknown tracker alerts on Android help keep you safe," Jul. 2023, last visit 16 of March, 2024. [Online]. Available: https://blog.google/products/android/unknown-tracker-alert-google-android/

[57] P. Somasundaram, "Two women sue Apple, saying stalkers used AirTags to track them," *Washington Post*, Dec. 2022, last visit 16 of March, 2024. [Online]. Available: https://www.washingtonpost.com/nation/2022/12/07/apple-airtag-lawsuit-stalking/

[58] J. Martin, D. Alpuche, K. Bodeman, L. Brown, E. Fenske, L. Foppe, T. Mayberry, E. C. Rye, B. Sipes, and S. Teplov, "Handoff all your privacy: A review of apple's bluetooth low energy continuity protocol," *arXiv preprint arXiv:1904.10600*, 2019.

[59] K. Fawaz, K.-H. Kim, and K. G. Shin, "Protecting privacy of {BLE} device users," in *25th USENIX Security Symposium (USENIX Security 16)*, 2016, pp. 1205–1221.

[60] D. Banisar, *Privacy and Human Rights...: An International Survey of Privacy Laws and Developments.* Electronic Privacy Information Center, 1999.

[61] J. Q. Whitman, "The Two Western Cultures of Privacy: Dignity versus Liberty," *The Yale Law Journal*, vol. 113, no. 6, pp. 1151–1221, 2004, publisher: The Yale Law Journal Company, Inc. [Online]. Available: https://www.jstor.org/stable/4135723

[62] A. Lukács, "What is privacy? the history and definition of privacy," 2016.

[63] A. F. Westin, “Privacy and freedom,” *Washington and Lee Law Review*, vol. 25, no. 1, p. 166, 1968.

[64] H. F. Atlam and G. B. Wills, “IoT Security, Privacy, Safety and Ethics,” in *Digital Twin Technologies and Smart Cities*, M. Farsi, A. Daneshkhah, A. Hosseinian-Far, and H. Jahankhani, Eds. Cham: Springer International Publishing, 2020, pp. 123–149, series Title: Internet of Things. [Online]. Available: http://link.springer.com/10.1007/978-3-030-18732-3_8

[65] “What is GDPR, the EU’s new data protection law?” Nov. 2018, last visit 16 of March, 2024. [Online]. Available: https://gdpr.eu/what-is-gdpr/

[66] O. Radley-Gardner, H. Beale, and R. Zimmermann, Eds., *Fundamental Texts On European Private Law.* Hart Publishing, 2016. [Online]. Available: http://www.bloomsburycollections.com/book/fundamental-texts-on-european-private-law-1

[67] “About NIST,” Jul. 2009, last visit 16 of March, 2024. [Online]. Available: https://www.nist.gov/about-nist

[68] National Institute of Standards and Technology, “NIST PRIVACY FRAMEWORK:: A TOOL FOR IMPROVING PRIVACY THROUGH ENTERPRISE RISK MANAGEMENT, VERSION 1.0,” National Institute of Standards and Technology, Gaithersburg, MD, Tech. Rep. NIST CSWP 01162020, Jan. 2020. [Online]. Available: https://nvlpubs.nist.gov/nistpubs/CSWP/NIST.CSWP.01162020.pdf

[69] N. Lefkovitz and K. Boeckl, “Nist privacy framework: An overview,” 2020.

[70] NIST, “Getting Started,” *NIST*, Jan. 2020, last visit 16 of March, 2024. [Online]. Available: https://www.nist.gov/privacy-framework/getting-started-0

[71] W. Labda, N. Mehandjiev, and P. Sampaio, “Modeling of privacy-aware business processes in bpmn to protect personal data,” in *Proceedings of the 29th Annual ACM Symposium on Applied Computing*, ser. SAC ’14. New York, NY, USA: Association for Computing Machinery, 2014, pp. 1399–1405. [Online]. Available: https://doi.org/10.1145/2554850.2555014

[72] C. Kalloniatis, E. Kavakli, and S. Gritzalis, “Addressing privacy requirements in system design: the PriS method,” *Requirements Engineering*, vol. 13, no. 3, pp. 241–255, Sep. 2008. [Online]. Available: https://doi.org/10.1007/s00766-008-0067-3

[73] M. Gharib, P. Giorgini, and J. Mylopoulos, “Towards an Ontology for Privacy Requirements via a Systematic Literature Review,” in *Conceptual Modeling*, H. C. Mayr, G. Guizzardi, H. Ma, and O. Pastor, Eds. Cham: Springer International Publishing, 2017, vol. 10650, pp. 193–208, series Title: Lecture Notes in Computer Science. [Online]. Available: http://link.springer.com/10.1007/978-3-319-69904-2_16

[74] M. Gharib, J. Mylopoulos, and P. Giorgini, “Copri v. 2 - a core ontology for privacy requirements,” *Data & Knowledge Engineering*, vol. 133, p. 101888, 2021.

[75] ——, “COPri - A Core Ontology for Privacy Requirements Engineering,” in *Research Challenges in Information Science*, ser. Lecture Notes in Business Information Processing, F. Dalpiaz, J. Zdravkovic, and P. Loucopoulos, Eds. Cham: Springer International Publishing, 2020, pp. 472–489.

[76] “SR 235.1 - Federal Act of 25 September 2020 on Data Protection (Data Protection Act, FADP),” last visit 16 of March, 2024. [Online]. Available: https://www.fedlex.admin.ch/eli/cc/2022/491/en

[77] “7 Major Differences Between the New FADP and GDPR,” last visit 16 of March, 2024. [Online]. Available: https://www.adnovum.com/blog/swiss-data-protection-law-how-the-new-fadp-differs-from-the-gdpr

[78] S. M. E. Portal, “New Federal Act on Data Protection (nFADP),” last visit 16 of March, 2024. [Online]. Available: https://www.kmu.admin.ch/kmu/en/home/fakten-und-trends/digitalisierung/datenschutz/neues-datenschutzgesetz-revdsg.html

[79] A. Heinrich, M. Stute, and M. Hollick, “OpenHaystack: a framework for tracking personal bluetooth devices via Apple’s massive find my network,” in *Proceedings of the 14th ACM Conference on Security and Privacy in Wireless and Mobile Networks.* Abu Dhabi United Arab Emirates: ACM, Jun. 2021, pp. 374–376. [Online]. Available: https://dl.acm.org/doi/10.1145/3448300.3468251

[80] “Find You: Building a stealth AirTag clone | Positive Security,” last visit 16 of March, 2024. [Online]. Available: https://positive.security/blog/find-you

[81] J. Briggs and C. Geeng, “BLE-Doubt: Smartphone-Based Detection of Malicious Bluetooth Trackers,” in *2022 IEEE Security and Privacy Workshops (SPW).* San Francisco, CA, USA: IEEE, May 2022, pp. 208–214. [Online]. Available: https://ieeexplore.ieee.org/document/9833870/

[82] G. Celosia and M. Cunche, “Discontinued Privacy: Personal Data Leaks in Apple Bluetooth-Low-Energy Continuity Protocols,” *Proceedings on Privacy Enhancing Technologies*, vol. 2020, no. 1, pp. 26–46, Jan. 2020. [Online]. Available: https://petsymposium.org/popets/2020/popets-2020-0003.php

[83] M. Stute, A. Heinrich, J. Lorenz, and M. Hollick, “Disrupting continuity of apple’s wireless ecosystem security: New tracking, DoS, and MitM attacks on iOS and macOS through bluetooth low energy,AWDL, and Wi-Fi,” in *30th USENIX security symposium (USENIX Security 21)*, 2021, pp. 3917–3934.

[84] P. Leu, G. Camurati, A. Heinrich, M. Roeschlin, C. Anliker, M. Hollick, S. Capkun, and J. Classen, “Ghost peak: Practical distance reduction attacks against {HRP}{UWB} ranging,” in *31st USENIX Security Symposium (USENIX Security 22)*, 2022, pp. 1343–1359.

[85] L. Tawalbeh, F. Muheidat, M. Tawalbeh, and M. Quwaider, “Iot privacy and security: Challenges and solutions,” *Applied Sciences*, vol. 10, no. 12, p. 4102, 2020.

[86] G. Yang, "An Overview of Current Solutions for Privacy in the Internet of Things," *Frontiers in Artificial Intelligence*, vol. 5, p. 812732, Mar. 2022. [Online]. Available: https://www.frontiersin.org/articles/10.3389/frai.2022.812732/full

[87] N. Okui, V. Bracamonte, S. Kiyomoto, and A. Duke, "Iot data privacy," *The Internet of Things: From Data to Insight*, pp. 121–139, 2020.

[88] "California Consumer Privacy Act (CCPA)," Oct. 2018, last visit 16 of March, 2024. [Online]. Available: https://oag.ca.gov/privacy/ccpa

[89] Apple, "Apple Platform Security," may 2022.

[90] ——, "Location Services Privacy Overview," nov 2019.

[91] ——, "A Day in the Life of Your Data," apr 2021.

[92] "Legal - Privacy Governance - Apple," last visit 16 of March, 2024. [Online]. Available: https://www.apple.com/legal/privacy/en-ww/governance/

[93] "Privacy - Features," last visit 16 of March, 2024. [Online]. Available: https://www.apple.com/privacy/features/

[94] L. Chen, D. Moody, A. Regenscheid, and K. Randall, "Recommendations for discrete logarithm-based cryptography: Elliptic curve domain parameters," National Institute of Standards and Technology, Tech. Rep., 2019.

[95] "SafeCurves: Introduction," last visit 16 of March, 2024. [Online]. Available: https://safecurves.cr.yp.to/

[96] J. Martin, T. Mayberry, C. Donahue, L. Foppe, L. Brown, C. Riggins, E. C. Rye, and D. Brown, "A study of mac address randomization in mobile devices and when it fails," *arXiv preprint arXiv:1703.02874*, 2017.

[97] J. K. Becker, D. Li, and D. Starobinski, "Tracking anonymized bluetooth devices," *Proceedings on Privacy Enhancing Technologies*, 2019.

[98] C. Matte, M. Cunche, F. Rousseau, and M. Vanhoef, "Defeating mac address randomization through timing attacks," in *Proceedings of the 9th ACM Conference on Security & Privacy in Wireless and Mobile Networks*, 2016, pp. 15–20.

[99] E. Fenske, D. Brown, J. Martin, T. Mayberry, P. Ryan, and E. Rye, "Three years later: A study of mac address randomization in mobile devices and when it succeeds," *Proceedings on Privacy Enhancing Technologies*, 2021.

[100] Y.-A. De Montjoye, C. A. Hidalgo, M. Verleysen, and V. D. Blondel, "Unique in the crowd: The privacy bounds of human mobility," *Scientific reports*, vol. 3, no. 1, pp. 1–5, 2013.

[101] G. Gagnon, S. Gambs, and M. Cunche, "Rssi-based fingerprinting of bluetooth low energy devices," in *International Conference on Security and Cryptography (SECRYPT 2023)*, 2023.

[102] "Legal Process Guidelines," last visit 16 of March, 2024. [Online]. Available: https://www.apple.com/legal/privacy/law-enforcement-guidelines-us.pdf

[103] M. Levitt, "AirTags are being used to track people and cars. Here's what is being done about it," *NPR*, Feb. 2022, last visit 16 of March, 2024. [Online]. Available: https://www.npr.org/2022/02/18/1080944193/apple-airtags-theft-stalking-privacy-tech

[104] "About APEC," last visit 16 of March, 2024. [Online]. Available: https://www.apec.org/about-us/about-apec

[105] "Apec CROSS-BORDER PRIVACY RULES SYSTEM," 2019, last visit 16 of March, 2024. [Online]. Available: https://cbprs.org/documents/

[106] "Apec PRIVACY FRAMEWORK," 2015, last visit 16 of March, 2024. [Online]. Available: https://www.apec.org/publications/2017/08/apec-privacy-framework-(2015)

[107] "The OECD PRIVACY FRAMEWORK," 2013, last visit 16 of March, 2024. [Online]. Available: https://www.oecd.org/sti/ieconomy/oecd_privacy_framework.pdf

[108] "Apec Privacy Recognition for Processors Purpose and Background," 2015, last visit 16 of March, 2024. [Online]. Available: https://cbprs.blob.core.windows.net/files/PRP%20-%20Purpose%20and%20Background.pdf

[109] S. Cole, "Police Records Show Women Are Being Stalked With Apple AirTags Across the Country," Apr. 2022, last visit 16 of March, 2024. [Online]. Available: https://www.vice.com/en/article/y3vj3y/apple-airtags-police-reports-stalking-harassment

[110] E. Fry, "Chilling rise of AirTag stalking as Apple device is branded 'gift to abusers'," Apr. 2023, last visit 16 of March, 2024. [Online]. Available: https://www.mirror.co.uk/tech/inside-chilling-rise-airtag-stalking-29746435

[111] "Ex-partner uses Apple AirTag to stalk Ahmedabad woman, device found hidden under driver's seat," last visit 16 of March, 2024. [Online]. Available: https://bit.ly/4amwV9r

[112] AirtagAlex, "Howto remove the speaker coil from the Airtag (make your Airtag very silent!)," Feb. 2021, last visit 16 of March, 2024. [Online]. Available: https://www.youtube.com/watch?v=sgGNShP9H8A

[113] "'Silent AirTags' With Speakers Removed Pop Up on Etsy, eBay," Feb. 2022, last visit 16 of March, 2024. [Online]. Available: https://uk.pcmag.com/mobile-phone-accessories/138509/silent-airtags-with-speakers-removed-pop-up-on-etsy-ebay

[114] "What to do if you get an alert that an AirTag, Find My network accessory, or set of AirPods is with you," Apr. 2023, last visit 16 of March, 2024. [Online]. Available: https://support.apple.com/en-us/HT212227

[115] "Review | Am I being tracked? Anti-stalking tech from Apple, Tile falls short." Mar. 2022, last visit 16 of March, 2024. [Online]. Available: https://www.washingtonpost.com/technology/2022/03/31/airtags-stalking/

[116] "How to find, block, and disable an AirTag that's tracking you | Macworld," last visit 16 of March, 2024. [Online]. Available: https://www.macworld.com/article/345863/how-to-find-block-disable-airtag-moving-with-you.html

[117] DOMARS, "Threat modeling for drivers - Windows drivers," Aug. 2023, last visit 16 of March, 2024. [Online]. Available: https://learn.microsoft.com/en-us/windows-hardware/drivers/driversecurity/threat-modeling-for-drivers

[118] I. T. L. Computer Security Division, "About the RMF - NIST Risk Management Framework | CSRC | CSRC," Nov. 2016, last visit 16 of March, 2024. [Online]. Available: https://csrc.nist.gov/projects/risk-management/about-rmf

[119] "OWASP Risk Rating Methodology | OWASP Foundation," last visit 16 of March, 2024. [Online]. Available: https://owasp.org/www-community/OWASP_Risk_Rating_Methodology

[120] "The Importance and Effectiveness of Cyber Risk Quantification," last visit 16 of March, 2024. [Online]. Available: https://www.fairinstitute.org/what-is-fair

[121] A. R. Mahlous, "Threat model and risk management for a smart home iot system," *Informatica*, vol. 47, no. 1, 2023.

[122] L. Zhang, A. Taal, R. Cushing, C. de Laat, and P. Grosso, "A risk-level assessment system based on the stride/dread model for digital data marketplaces," *International Journal of Information Security*, pp. 1–17, 2022.

[123] kexugit, "DREADful," Aug. 2007, last visit 16 of March, 2024. [Online]. Available: https://learn.microsoft.com/en-us/archive/blogs/david_leblanc/dreadful

[124] "Mobile OS market share worldwide 2009-2023," last visit 16 of March, 2024. [Online]. Available: https://www.statista.com/statistics/272698/global-market-share-held-by-mobile-operating-systems-since-2009/

[125] "Global smartphone penetration 2016-2022," last visit 16 of March, 2024. [Online]. Available: https://www.statista.com/statistics/203734/global-smartphone-penetration-per-capita-since-2005/

[126] C. Silva, "Google and Apple are closer to making AirTags stalker free," Dec. 2023, last visit 16 of March, 2024. [Online]. Available: https://mashable.com/article/protection-privacy-airtag-google-apple

Abbreviations

AD	Advertising Data
AdvA	Advertising Address
APEC	Asia-Pacific Economic Cooperation
BLE	Bluetooth Low-Energy
BR/EDR	Basic Rate/Enhanced Data Rate
CBPR	Cross-Border Privacy Rules
COFN	Crowdsourced Offline Finding Network
COPri	Core Ontology for Privacy requirements engineering
CRC	Cyclic Redundancy Check
FAIR	Factor Analysis of Information Risk
FiRa	Fine Ranging
GDPR	General Data Protection Regulation
GFSK	Gaussian Frequency Shift Keying
GHz	Gigahertz
GPS	Global Positioning System
IoT	Internet of Things
IR-UWB	Impulse Radio Ultra Wideband
ISA	Item Safety Alert
LoRaWAN	Long Range Wide Area Network
LSB	Least Significant Bit
MAC	Media Access Control
MSB	Most Significant Bit
nFADP	New Federal Act on Data Protection
NFC	Near Field Communication
NIST	National Institute of Standards and Technology
OECD	Organisation for Economic Co-operation and Development
OS	Operating System
PDU	Protocol Data Unit
PPM	Privacy Preference Management
RQ	Research Question
PRP	Privacy Recognition for Processors
RSSI	Received Signal Strength Indication
ToF	Time of Flight
UTA	Unwanted Tracker Alerts
UWB	Ultra Wideband
WiFi	Wireless Fidelity

List of Figures

List of Tables