

A Thesis Submitted for the Degree of PhD at the University of Warwick

Permanent WRAP URL:

<http://wrap.warwick.ac.uk/181792>

Copyright and reuse:


This thesis is made available online and is protected by original copyright.

Please scroll down to view the document itself.

Please refer to the repository record for this item for information to help you to cite it.

Our policy information is available from the repository home page.

For more information, please contact the WRAP Team at: wrap@warwick.ac.uk



New authentication applications in the protection of caller ID and banknote

by

Shen Wang

Thesis

Submitted to The University of Warwick

in partial fulfilment of the requirements

for admission to the degree of

Doctor of Philosophy in Computer Science

Department of Computer Science

January 2023

Contents

List of Tables	iv
List of Figures	v
Acknowledgments	vii
Declarations	viii
Abstract	ix
Chapter 1 Introduction	1
1.1 Motivation	1
1.2 Contributions	7
1.3 Thesis Outline	8
Chapter 2 Authentication	10
2.1 Overview	10
2.2 Something you know	11
2.3 Something you own	13
2.3.1 One-time password-based	14
2.3.2 Cryptography-based	15
2.4 Something you are	17
2.5 Summary	22
Chapter 3 Spoofing against Spoofing - Caller ID Verification	24
3.1 Introduction	24
3.2 Related work	28
3.2.1 Top-down	29
3.2.2 Bottom-up	30
3.2.3 Comparison between STIR/SHAKEN and CIV	30

3.3	Background	31
3.3.1	Heterogeneous networks	32
3.3.2	Caller ID spoofing	33
3.3.3	Dual-tone multi-frequency	34
3.4	Our proposed system	35
3.4.1	Threat model	36
3.4.2	Protocol description	36
3.4.3	Distinguishing legitimate/illegitimate spoofing	40
3.5	Prototypes	41
3.5.1	Overview	41
3.5.2	CLI-based prototypes	42
3.5.3	DTMF-based prototypes	45
3.6	Evaluations	49
3.6.1	SIP platforms	49
3.6.2	Other platforms	51
3.6.3	Lengths of the challenge and response	52
3.7	Discussion	52
3.7.1	Downgrading attack	52
3.7.2	Attacking the challenge-response protocol	53
3.7.3	Denial-of-service (DoS)	54
3.7.4	Stages of deployment and optimisation	55
3.7.5	Limitations	56
3.8	Conclusion	57

Chapter 4 Counterfeiting No More - Polymer Substrate Fingerprinting 58

4.1	Introduction	58
4.2	Related Work	60
4.3	Production of Polymer Banknote	63
4.4	Proposed Solution	66
4.4.1	Feature Area	66
4.4.2	Experiment Setup	66
4.4.3	Image Processing	69
4.5	Datasets	73
4.5.1	Benchmark	73
4.5.2	Robustness Test	74
4.5.3	Variation Test	75

4.6	Evaluation	76
4.6.1	Framework	76
4.6.2	Results	78
4.6.3	Limitations	84
4.7	Anti-counterfeiting Applications	85
4.7.1	False Acceptance Rate	85
4.7.2	Online Application	85
4.7.3	Offline Application	86
4.8	Security Analysis	87
4.8.1	Threat Model	87
4.8.2	Attack on Fingerprints	88
4.9	Conclusion	90
Chapter 5 Conclusion and Future Work		92
5.1	Summary	92
5.2	Future work	94

List of Tables

2.1	Summary of authentication schemes	10
2.2	Summary of PUFs	23
3.1	Comparison with related works	28
3.2	Comparison of STIR/SHAKEN and CIV	31
4.1	Summary of Datasets	73
4.2	Notations used in PUF metrics. (In the benchmark dataset, S = 100, $L = 2048$ and $T = 10$)	77
4.3	PUF metrics from Benchmark dataset	80
4.4	PUF metrics after robustness tests vs benchmark	82
4.5	PUF metrics after variation tests vs benchmark	82
4.6	False match for one-to-one comparison	86

List of Figures

2.1	Workflow of a basic biometrics system	17
3.1	Overview of DTMF transmission	35
3.2	CIV for heterogeneous telecommunication networks	36
3.3	Intuition behind CIV	37
3.4	Illustration of the CIV user interface	39
3.5	Equipment setup	41
3.6	Overview of CIV prototypes	43
3.7	CIV using CLI/DTMF to send challenge/response	46
3.8	Using DTMF/DTMF to send challenge/response	47
3.9	Delays in CIV between different phones	49
3.10	Cost breakdowns in CIV	50
3.11	DTMF transmission time versus number of digits	52
4.1	Schematic of the Gravure printing process	65
4.2	Feature extraction on different banknotes (the zoomed-in pictures are cropped as a square from the original images for demonstration)	67
4.3	Scanner setup	68
4.4	Camera setup	69
4.5	Physical dimension of the Gabor filter	72
4.6	HD distributions of Benchmark dataset. Decidability $d' \approx 29$	80
4.7	HD histograms after robustness tests	81
4.8	Histograms of variation tests	83
4.9	The Procedure of Fingerprint Registration. The QR code shown in the diagram is generated from a real £10 polymer banknote.	87
4.10	A close-up of translucent patterns	91

4.11 A printed dot ‘.’ in different font sizes using inkjet and laser
printers. The resolution for both printers is 1200 DPI. 91

Acknowledgments

I would like to express my deepest appreciation to my supervisor, Prof. Feng Hao, for his valuable support and encouragement during my PhD study. With his guidance and help, I could focus on pursuing my research interest on the right track and overcome the difficulties through the whole journey. I also thank Dr. Mahshid Delavar and Dr. Ehsan Toreini for their delightful inspiration and help.

This endeavour would not have been possible without the support from my family. I thank my parents, who taught me how to stay optimistic and positive during tough times. Moreover, my wife and my son shared this experience with me and deserved endless gratitude.

I would like to extend my sincere thanks to my friends, my colleagues and the staff at the department of computer science. Without their help, the journey would not be so enjoyable and fruitful.

Declarations

This thesis is submitted to the University of Warwick in support of my application for the degree of Doctor of Philosophy. I hereby declare that the thesis has been composed by myself and has not been submitted in any previous applications for any degree. The work in this thesis has been undertaken by myself under the supervision of Prof. Feng Hao.

- Chapter 3 has been submitted by the author and is currently under review: *Spoofing Against Spoofing: Towards Caller ID Verification In Heterogeneous Telecommunication Systems Without PKI*. My contributions in the paper are as follows: 1) the research on DTMF signalling; 2) the design of CIV prototypes; 3) the implementation of the prototypes on the Android platform in the cellular network.; 4) assisting with the implementation of the prototypes on the other two platforms.; 5) conducting the experiments.; 6) analysing the evaluation results.
- Chapter 4 has been published by the author: *Anti-counterfeiting for polymer banknotes based on polymer substrate fingerprinting* [170]. My contributions in the paper are as follows: 1) research on polymer substrate; 2) design of PSF prototypes.; 3) the implementation of the prototypes with two setups; 4) collection of sample image datasets based on polymer and paper banknotes with different denominations issued in the UK.; 5) conducting experiments.; 6) analysis of evaluation results.

Abstract

In the era of computers and the Internet, where almost everything is interconnected, authentication plays a crucial role in safeguarding online and offline data. As authentication systems face continuous testing from advanced attacking techniques and tools, the need for evolving authentication technology becomes imperative. In this thesis, we study attacks on authentication systems and propose countermeasures. Considering various nominated techniques, the thesis is divided into two parts.

The first part introduces caller ID verification (CIV) protocol to address caller ID spoofing in telecommunication systems. This kind of attack usually follows fraud, which not only inflicts financial losses on victims but also reduces public trust in the telephone system. We propose CIV to authenticate the caller ID based on a challenge-response process. We show that spoofing can be leveraged, in conjunction with dual tone multi-frequency (DTMF), to efficiently implement the challenge-response process, i.e., using spoofing to fight against spoofing. We conduct extensive experiments showing that our solution can work reliably across the legacy and new telephony systems, including landline, cellular and Internet protocol (IP) network, without the cooperation of telecom providers.

In the second part, we present polymer substrate fingerprinting (PSF) as a method to combat counterfeiting of banknotes in the financial area. Our technique is built on the observation that the opacity coating leaves uneven thickness in the polymer substrate, resulting in random translucent patterns when a polymer banknote is back-lit by a light source. With extensive experiments, we show that our method can reliably authenticate banknotes and is robust against rough daily handling of banknotes. Furthermore, we show that the extracted fingerprints are extremely scalable to identify every polymer note circulated globally. Our method ensures that even when counterfeiters have procured the same printing equipment and ink as used by a legitimate government, counterfeiting banknotes remains infeasible.

Chapter 1

Introduction

1.1 Motivation

The use of authentication can be traced back to the ancient history of mankind. An authentication process was commonly adopted in the military by the sentinel in the fort to verify the returning patrol soldiers based on watchwords during the late BC and early AD periods [125]. Today, with everything digitised, authentication plays an even more important role in people's daily lives. As computers and the Internet are accessible everywhere, most communication occurs between strangers instead of friends and between computers instead of humans. It becomes a necessary mechanism to build the initial trust between the participants during communication.

The security of an authentication method has a close relationship to the environment, implying that the process may be broken when the circumstance changes. Threats are introduced when the new system replaces the old one, while new attacking schemes become possible when the adversaries hold the advanced technologies. Due to its importance for the entire system, authentication techniques need to be evaluated and strengthened continuously.

When a legacy system is replaced by a new system. The Voice over IP (VoIP) protocol has been extensively used as the next generation of the communication system. It has several advantages over the public switched telephone network (PSTN), such as low cost, high call quality, and flexible setup.

The increasing adoption of cloud-based VoIP services help small to medium-sized companies reduce costs while maintaining reliability and flexibility to

support their expanding businesses. In the global telecommunication market, VoIP was estimated to be worth more than 30 billion dollars in 2020 and 85 billion dollars in 2021 [12, 132]. The market is expected to grow continuously at 3.5% each year for the next 5 years.

However, the design of this Internet protocol (IP) based technologies allow caller ID spoofing to be conducted with minimal cost and effort, which was infeasible in the old system.

Caller ID spoofing occurs when the caller intentionally presents a false number to hide their true identity. The technique means to protect the privacy of the caller but is often abused to deceive the receiver to conduct fraud [3, 8]. By deliberately changing its caller ID, the scam call appears as a call from a trustworthy entity, while the spoofer in the call impersonates the staff of the entity to deceive the victim.

Like the identity card, caller ID nowadays is commonly used to represent an individual or organisation. Finance systems or Internet service providers use the incoming caller ID to authenticate or locate their customers, while the emergency services and police use the outgoing caller ID to identify themselves to the receiver. In the legacy system (PSTN), a caller ID is stored and transmitted in plaintext. No protection is implemented to maintain its authenticity or integrity. Similarly, when a call travels through another carrier network, there is no authentication procedure to guarantee that the displayed caller ID belongs to the caller in the previous network. It was not a problem in the debut and early stages of the telephony system. As the telecommunication network was a largely closed system, delivering a caller ID spoofing attack was expensive and there was a high risk of being exposed. The fraudster needed to physically intrude the telephone infrastructure to manipulate its signalling system. However, for financial consideration, more and more add-value services provided by telecommunication carriers force the network to be semi-opened to the public in recent years, making caller ID spoofing possible without access to the core network facility. *Motherboard* from *Vice* reports that a SIM card (White SIMs) circulated underground could spoof any number that the spoofer wants [48]. The operator of White SIMs “likely runs their own Mobile Virtual Network Operator (MVNO)” as a legal add-value service on the established network infrastructure. As buying and using them does not require identity information from the user, White SIMs actually encourage fraud activities with minimal risk.

To satisfy the requirement of backward compatibility, VoIP does not design

the security measure to protect caller ID. The plain text of the caller ID is simply embedded in the message of the text-based session initiation protocol (SIP), providing a cheaper and easier way of spoofing [177]. A VoIP-ready private branch exchange (PBX) or a private SIP server can connect to the network through a leased SIP trunk at the cost of a few dollars per month. The equipment can modify the caller ID and deliver the attack in batches, which is adopted by many fraudsters. An individual attacker would choose the SIP client with a pre-designed spoofing function, requiring no technical background to operate.

In the US, the approach of using caller ID spoofing is legal unless it is “done with the intent to defraud, cause harm, or wrongfully obtain anything of value” [131]. In practice, to avoid caller ID spoofing being abused, the legal spoofing approach needs to follow the restrictions. Different legislation and rules have been passed to protect phone users. For example, in 2010, a law called *Truth in Caller ID Act of 2009* [44] was approved in the US. Due to privacy considerations, the Act allows the caller to block their caller ID information on the remote terminal but forbids one to “transmit misleading or inaccurate caller ID information”. The legislation was expanded and supplemented by *Ray Baum’s Act* [45] adopted in 2019, which gives the Federal Communications Commission (FCC) the right to regulate calls and messages with spoofed caller IDs that originate outside the United States.

Such legislation and rules regulate the behaviour of companies when they advertise their products and services, but scammers are barely influenced. One reason is that the caller ID spoofing scam is difficult to recognise. The caller ID displayed in the attack is an authentic phone number from a famous company or a government emergency service in the local area [10, 11]. From the incoming phone number alone, it is infeasible to tell whether the call originates from the authentic caller or the spoofer. The other reason is that the caller ID spoofing attack is difficult to trace. Such scams are usually initiated from abroad and routed by several proxy agents before entering the domestic area, while the IP address of their equipment can be easily changed when necessary.

In recent years, the number of reported caller ID spoofing has continued to increase worldwide. According to a survey conducted by Which?, more than 8000 cases of impersonation scams were reported in the first half of 2019 in the UK, and about 60 million pounds were stolen by the phone number spoofing scams [15]. Compared to the first six months of 2018, the losses increased by over 50%. In 2021, Ofcom, the UK telecoms regulator, estimated

that more than 65% of UK people “may have received scam calls and text messages during the summer”, and 2% of the scam calls were successful [4]. The fraudulent cases are expected to increase during the post-pandemic period in 2022 [4]. From the statistics in [10], by the end of 2018, more than 26% of all calls in the US were scams, in which more than 65% of these spam calls were using a local phone number (this technique is called *Neighbour Spoofing*). In 2019, the number of scam calls doubled compared to the previous year, in which the call with a spoofed phone number had a significantly higher success rate among these scam calls [11]. By adopting *Enterprise Spoofing*, which uses a legitimate business outbound call number as the spoofed calling number, 33% of the scam calls were answered and more than 30% of these scam calls were successful.

As a global industry problem, caller ID spoofing often acts as a critical enabler for telephone fraud. To address this problem, the Federal Communications Commission (FCC) has mandated telecom providers in the United States to implement STIR/SHAKEN (short for Secure Telephony Identity Revisited/Signature-based Handling of Asserted information using toKENs) [85][84], an industry-driven solution based on digital signatures. However, STIR/SHAKEN has two main limitations. First, it requires a new public key infrastructure (PKI), but scaling up this PKI for the global telecom industry is difficult. Second, it only works with the SIP system (VoIP), leaving the traditional signalling system No.7 (SS7) systems, including landline and cellular networks, unprotected. As STIR/SHAKEN encounters challenges in global adoption, alternatives have not been adequately studied. We will explain more details on STIR/SHAKEN and present a more cost-effective solution in Chapter 3.

When adversaries hold advanced technologies. Despite the rapidly increasing volume of transactions made by debit cards and cryptocurrency, banknotes remain one of the most important payment instruments around the world. Banknotes are convenient, reliable, user-friendly, and function without the extra device, indicating that they are far from being replaced soon. In 2018, McKinsey & Company reported that although the share decreased from 89% to 77% in the last five years, cash remained the main transaction instrument in the world [25]. A survey of Paypers in 2019 drew a similar conclusion [124]. It reported that cash is expected to be the second most frequently used payment method in the UK until at least 2025 [64].

As the value depends on its ability to pay the bearer on demand, banknotes require security measures to maintain their trustworthy property. One solution to protect banknotes is to apply security features in the manufacturing process, which increases the cost and difficulty of counterfeiting. In the UK, raised notes, ultraviolet marks, metallic threads, and hologram images give legitimate banknotes with the unique sight of view and haptic feedback of touch [2].

The alternative solution is to adopt monopolised substrates, such as the cotton paper used to produce the paper banknote. Cotton paper has an exclusive touchable texture, strength, and durability that withstands everyday use. It is the material that has been used to produce banknotes in the past few decades. However, due to the increase in the volume of counterfeit paper banknotes, a demand to replace the paper substrate with another more secure material grows. In 1988, the Reserve Bank of Australia introduced new banknotes with the use of a polymer substrate. The new material is cleaner, more secure and more durable [6]. Alternative security features have been developed to protect polymer banknotes, such as the see-through window and the foil patch.

In the recent two decades, more than 30 countries around the world have issued polymer banknotes as a replacement for paper banknotes. Bank of England issued polymer £5 and £10 notes in 2016 and 2017, respectively. In the past 5 years, few counterfeit polymer banknotes have been reported in circulation in the UK. In 2019, 89% counterfeits are £20 notes and 10% £50 notes as they were still made of paper. In 2020, the total volume of counterfeits discovered was reduced to less than half compared to the previous year due to the issue of £20 polymer notes. In the first half calendar year of 2022, the volume of counterfeit was only one-sixth of that in 2019, given that polymer £50 notes hadn't been issued until the middle of 2021. Less than 0.0031% of the banknotes were counterfeit, even though old £20 and £50 paper notes were still legally used in circulation during data collection [9].

Anti-counterfeiting is the arms race between the authorities and adversaries. Since counterfeits sabotage the economy, authorities try hard to reduce the impact of fake banknotes. In addition to protecting banknotes with security features and substrate, another method is to periodically review the design of legitimate banknotes. In the UK, the authorities have issued at least five generations of banknotes in different denominations in the last 30 years [16]. Banknotes with the old design are withdrawn; advanced security features are added to the new design; the new mass production process is implemented; the

corresponding hardware is upgraded. As an example of the activity of £20 notes issuing, the total volume of £20 paper notes in circulation is approximately 2100 million in 2018 [9]. These paper banknotes need to be withdrawn after the new £20 polymer notes are issued, while about 70000 ATMs and countless detectors wait to be upgraded to adopt the new £20 note design [64]. Not only does it increase total economic operating costs for our society, but it also requires a lot of effort from the individual user. Millions of users are asked to learn to recognise the new design. The whole process consumes time and effort, making issuing new banknotes expensive and complicated.

However, despite the great effort put into securing banknotes, there are two main reasons why counterfeits, especially high-quality ones, are difficult to remove from circulation. One reason is the rapid iteration of printing technologies. These new technologies not only help counterfeits mimic the appearance of legitimate banknotes, but also enable the types of material used as the substrate. Counterfeits made by alternative printing techniques do not require to be as durable as their genuine equivalence but only survive at least one transaction. For example, the metallic thread woven through the paper of the legitimate banknote could be replaced by a high-definition printed pattern with a similar appearance. With the maturity of ink-jet printing technologies, printing on plastic is possible at a low cost. The off-the-shelf equipment with these techniques is accessible to all people, including professional forgers. This results in an increasing volume of polymer counterfeits appearing in circulation. In 2019, the Reserve Bank of Australia reports that the total value of fake plastic money has been gradually increasing in the past five years, from less than 1 ppm (parts per million - the number of counterfeits per million genuine banknotes in circulation) to 15 ppm, in which 40% counterfeits have been considered high quality in the report [24]. The new £20 and £50 notes issued in the UK recently are expected to face more challenges in the arms race of counterfeits, as forgers always gain a high payoff from the high-denomination counterfeit.

The other reason is that current anti-counterfeiting principles are based on differential detection. Significant differences should be found by the naked eye or sophisticated detectors when the banknote is considered a counterfeit. This method is efficient when counterfeits are of low quality, but it easily fails with high-quality counterfeits. Even with the help of detectors, it is difficult for an untrained user to spot high-quality counterfeits accurately. The process requires the examiner to memorise several locations and appearances of security

features. Therefore, the accuracy of the counterfeits detection is a subjective result that depends on the sophistication of the equipment and the experience of the examiner [59]. For example, since there is no complete and consistent information on how to detect the well-known “Superdollar”, these supernotes had existed in circulation outside the banking system for a long time [118].

Several researchers from different communities have attempted to solve this counterfeiting problem. Some of them use a spectrometer to check the spectra of inks printed on genuine banknotes issued from different countries [46, 52, 92, 134, 144]. Some analyse the background texture printed on legitimate banknotes using supervised learning[27, 39]. Such studies focus on precision in reproduction. They believe in the existence of such features that legitimate banknotes bear, but counterfeits do not. To obtain a veracity result, several fake banknotes are collected to evaluate the proposed methods in the experiments. However, the results of these studies are highly dependent on the quality of the specimens collected. Since the ability to detect low-quality counterfeits has no value in practice, the subjective problem of how to obtain appropriate counterfeits cannot be solved in such studies. We will explain more details on polymer-banknote counterfeiting and will present a new anti-counterfeiting solution in Chapter 4.

1.2 Contributions

In this section, we summarised our contributions in this dissertation as follows.

1. In solving the problem of caller ID spoofing, we contribute the following.
 - We propose Caller ID Verification (CIV), a new bottom-up solution to authenticate the caller ID based on a challenge-response one-time-password process without depending on the digital signature.
 - We present concrete prototypes of CIV for landline, cellular, and VoIP phones, showing how our concept works across heterogeneous telecom networks (PSTN/SIP).
 - We systematically evaluate the performance of CIV in landline, cellular, and SIP networks to show feasibility. This is the first demonstration of a caller ID authentication solution that works on all three types of phone systems across heterogeneous networks.

2. In solving the problem of banknote counterfeiting, we contribute the following.

- We propose Polymer Substrate Fingerprinting (PSF), a novel anti-counterfeiting technique for polymer banknotes based on the paper-based physically unclonable function, which analyses the naturally occurring, unique, and unrepeatable imperfections in the opacity coating layer of a polymer substrate.
- We present a proof-of-concept implementation that uses a commodity negative-film scanner to capture those imperfections by photographing the random translucent patterns of a polymer substrate when it is back-lit and transforming them into a compact fingerprint for authentication.
- We collect an extensive dataset using polymer banknotes from the UK and conduct experiments to show that our technique can reliably authenticate banknotes with high accuracy, is robust against rough daily handling, and is highly scalable to identify every polymer banknote circulated in the world.

1.3 Thesis Outline

Chapter 2 discusses the state-of-the-art research and practice in authentication solutions as the background of our research. The content of the chapter is classified into three sections based on the adoption of different factors: password-based authentication, token-based authentication, and biometrics-based authentication. In each section, we review existing studies and discuss challenges.

In Chapter 3, we introduce our proposed solution to address the caller-ID spoofing attack with token-based authentication techniques. We present the literature review of the previous solutions, explaining their pros and cons. Then we demonstrate that caller ID spoofing can be leveraged, in conjunction with DTMF, to efficiently fight against spoofing. We conduct the experiment on and across platforms and heterogeneous telecommunication networks to support our proposed method.

In Chapter 4, we present a solution to combat the counterfeiting of banknotes with biometrics-based authentication techniques. We review previous research and then introduce our solution to authenticate the banknote with its

stochastic translucent pattern. We then evaluate our proposed method with various security characteristics through extensive experiments.

In Chapter 5, we conclude this dissertation and suggest future research.

Chapter 2

Authentication

2.1 Overview

Authentication is a process to validate that an object is which this object claims to be. One authentication process can be referred to authenticate the object's integrity, or the object's identity. In this thesis, we concern the latter rather than the former. Authentication can take place locally or remotely between two or more principals, which include people, equipment, and services. It is one of the most fundamental building blocks of security measures to protect the assets of the system [20].

Table 2.1: Summary of authentication schemes

Factor	Property	Application
Something you know	Primary	<ul style="list-style-type: none">• system-generated passwords• user-generated passwords
	Secondary	<ul style="list-style-type: none">• secret questions• associative passwords
Something you own	OTPs-based	<ul style="list-style-type: none">• HID ActivID[®]Token [5]• CAPTCHAs¹
	Cryptography-based	<ul style="list-style-type: none">• smart cards• JSON² web token
Something you are	Behavioural	<ul style="list-style-type: none">• handwriting• keystroke dynamics
	Physiological/Physical	<ul style="list-style-type: none">• fingerprint• iris pattern

¹Completely Automated Public Turing test to tell Computers and Humans Apart

²JavaScript Object Notation

Authentication typically involves three parties: the user, the system, and the verifier. The user, who could be a person or a computer, seeking access to the services provided by the system. The verifier is responsible for determining whether the user is authenticated to access these services by conducting the authentication process [22].

In authentication, three types of factors are typically used to confirm a user's identity: something you know, something you own, and something you are. The first factor, "something you know," refers to something that the user knows, such as a password or a personal identification number (PIN). The second factor, "something you own," refers to something that the user possesses, such as a security token or a mobile phone. And the third factor, "something you are," refers to something that is inherent to the user, such as a fingerprint or the facial features. If authentication utilises only one factor, it is called single-factor authentication (SFA), such as entering the password during the system log-in. When an application requires two or more factors to provide higher levels of security against unauthenticated access, it is called multi-factor authentication (MFA) [119], such as the withdrawal of cash from an ATM with a debit card.

Based on different factors, applications with different properties have been deployed to handle various scenarios (see Table 2.1). Details of these applications and relevant scenarios are given and explained in the rest of this chapter.

2.2 Something you know

Verifying the secret to ensure one's identity is simple but efficient. It has been adopted in the authentication as early as late BC in the military: If a soldier wanted to verify himself towards the sentry, presenting the codeword was necessary [125].

In an authentication process, a user's secret, such as a password, social security number, or other information, is compared to the secret stored in the system. If the two match, the authentication is successful, and the user is granted access to the system. Textual passwords are the most commonly used form of secret, as they are relatively easy to change when needed.

A password-enabled system allows easier design and maintenance without the need for auxiliary devices [80]. Based on the different levels of resource access, the primary password or the secondary password is used for the authen-

tication. The primary password is verified when a user tries to log in to the system, while the secondary passwords are preferred by the software designers to assist the primary passwords for further control and protection [179].

Despite its ubiquitous usage, password-based authentication is considered to have relatively low-security assurance [32, 97]. As a result, the authentication is often used in combination with other factors such as the token or biometrics based authentication to increase the security of the system. When used alone, password-based authentication is vulnerable to various attacks, including those that exploit weaknesses in users' password choices, the storage and handling of passwords within the system, and the communication between users and the authentication system [98].

Vulnerability of the password choices. The main disadvantage of knowledge-based authentication is that the success of the process highly relies on the limited memory of the user. Two fundamental principles are suggested in choosing a password [89, 148]. The first principle is to use a unique password for each account. However, many people reuse passwords, despite the security risks associated with this practice [116]. The second principle is that the password should be a long enough combination of numbers, symbols, and letters. Nevertheless, strong passwords are difficult for humans to remember [89].

Vulnerability of the passwords storage and handling. To prolong the time of a password-guessing attack success on a leaked database, the passwords stored in the database are usually encrypted during enrolment. The effectiveness of this approach depends on the choice of the encryption algorithm and its implementation. Thus, the security of a password-based system is greatly impacted by the knowledge and skills of its designer and developer.

Vulnerability in the authentication communication. To secure a remote user authentication process, the applied password-based authentication protocol should at least resist offline password guessing attacks, stolen verifier attacks, and denial of service (DoS) attacks [80]. Several studies have put effort into the subject, trying to design a protocol that could resist all these attacks using only a one-way hash function without the need for complex computation or public key certificates [82, 83, 90, 126, 136, 158, 159]. However, B.T Hsieh et al. point out that "based on our observation, designing a password authentication protocol using only the hash function is almost impossible to defeat all well-known

attacks” [80].

2.3 Something you own

When a token is authenticated, the process verifies the factor of ownership. If the token passes the validation, the owner is granted access to the asset. An example of token authentication is using a key to open the door.

To enhance the security of the process, tokens can be used in combination with the password. For instance, when withdrawing cash from an ATM, a user needs to present a smart card and enter a numeric password. In this way, even if the card is lost, the password serves as an additional layer of protection for the user’s assets. The use of a token and a password is considered a more secure authentication mechanism according to the user authentication level system (UALS) [97].

In terms of the storage medium, tokens can be divided into two types: hardware-based and software-based.

Hardware tokens include memory tokens and smart tokens. Memory tokens simply store information without any processing, an example of which is the magnetic stripe card commonly used for human-computer authentication. Smart tokens, on the other hand, use embedded microprocessors to compute and output cryptographic results based on the stored information [179]. Smart tokens offer higher security than memory tokens but are more costly to deploy due to the use of microprocessors. A common example of a smart token is a credit card with an embedded chip.

Software tokens are often considered a cost-effective and convenient solution for authentication conducted over an insecure channel. For example, Transport Layer Security (TLS) certificates are used to verify that a client is communicating with the correct server that owns a specific domain. Another example is the use of JSON web token (JWT), which securely transports authentication data between the parties over the Internet.

In terms of the methods used for secret generation and verification, token authentication can be categorised into two types: one-time password-based and cryptography-based.

2.3.1 One-time password-based

One-time password (OTP) is a password that is valid for only one authentication session. Compared to static passwords, OTPs are not vulnerable to replay attacks [21].

In OTP applications, the applied password is not fixed and therefore generating a valid password requires synchronisation between the token and the verifier. This synchronisation process can be achieved through the use of a counter, timestamp, or challenge-response process.

Counter synchronised OTP. The idea of OTP was first suggested by Leslie Lamport [98] in the early 1980s. He proposed using the mathematical hash chain to generate passwords so that each password is used and available for only one authentication instance. Lamport's scheme has three components: a counter value n , a secret s , and a one-way hash function h . During enrolment, the original password $p_0 = h^n(s)$ (The equation represents that s is hashed n times) is calculated. The p_0 is then sent to the verifier via an authenticated channel as the initial secret for the authentication. At the i th time of the authentication, $p_i = h^{n-i}(s)$ is computed, and then is sent to the verifier. The verifier hashes the received value one more time and compares it with the saved p_{i+1} , which should be equal to $h^{n-i+1}(s)$. If two match, the validation succeeds. The p_{i+1} will be replaced by the p_i in the database and the counter value is reduced by one after the validation. It is obvious that n should be large enough, as the token will be invalid when n reaches zero. Most counter-synchronised OTP solutions are designed with the same principle [75].

Time synchronised OTP. The high computation cost of counter synchronised OTP limits its usage on devices with low computational power. Time-synchronised OTP, which uses the internal clock to synchronise the password, avoids the high computation requirement in Lamport's scheme [81]. New passwords are generated from the current timestamp, rather than a previous password.

To address the issue of timing skew, an acceptance window or valid period threshold is used by the verifier [81, 173]. This reduces the number of false alarms when the clocks on the token and the verifier are slightly out-of-sync.

Challenge-response-based OTP. Challenge-response-based OTP is a method of token authentication that uses a series of challenges and correspond-

ing responses to verify the identity of a user. For example, CAPTCHAs is designed to protect against automated malicious activities that can exploit online systems. It presents users with challenges or puzzles and the user needs to send back the desired answer to pass the authentication. Such challenges are easy for humans to solve but difficult for computers or bots.

In a challenge-response-based OTP system, the authentication server generates a unique challenge for the user during the login or transaction process. This challenge is typically a random number or a combination of characters. The server then sends the challenge to the user's registered device, usually a mobile phone or a dedicated OTP device, via a secure channel. The user's device receives the challenge and bounces the response back. The authentication server compares the received response with the expected response it calculates. If the two responses match, the user is granted access. Otherwise, the authentication fails.

The key feature of challenge-response-based OTP is that each challenge is unique and used only once. Once a challenge-response pair is used, it becomes invalid for future authentication attempts, providing an additional layer of security. This makes OTP particularly effective against various attacks like replay attacks, where intercepted authentication data are used to deceive the verifier.

We proposed a solution based on the challenge-response OTP to address the caller ID spoofing attack on telecommunication systems. The details of this solution are discussed in Chapter 3.

2.3.2 Cryptography-based

OTP-based authentication relies on the generation of one-time passwords for each authentication attempt, while cryptography-based token authentication generates and validates tokens using cryptographic algorithms and shared secret keys.

Despite the goal of replacing static password-based systems, the security of OTPs still heavily relies on the security of the one-way hash function and the communication channel. If the hashed passwords are obtained by an attacker, they reveal a lot of information about the original secrets. Thus, to enhance the security of the challenge-response process, symmetric and asymmetric cryptography are proposed as additional layers of protection [28, 164].

Symmetric cryptography Symmetric cryptography, also known as secret key cryptography, uses the same shared secret key for both encryption and decryption. The general workflow of symmetric-based authentication is as follows: During enrolment, a self-certified key is shared through a secret channel. During authentication, a challenge, such as a large integer or array of integers, is sent to the end user. The user encrypts the challenge with the shared key and sends the result back to the verifier. If the response matches the one generated locally, the authentication is successful.

Due to its balance of power consumption and security level, symmetric challenge-response-based applications are widely used in the authentication process of the Internet of Things (IoT) [93, 108, 143].

Asymmetric cryptography Asymmetric cryptography, also known as public key cryptography, uses key pairs for encryption and decryption. Each participant in the authentication holds a pair of keys. The public key is used for encryption and digital signature verification, while the private key is used for decryption and digital signature generation [110, 146]. Since the self-certified key is not sufficient to prove one's identity, the public key must be provisioned by the verifier or a trusted third party (TTP).

The simplified process of asymmetric-based authentication is as follows: The verifier sends a random number as a challenge to the end user. The user uses their private key to sign the challenge and produces a digital signature. If the signature sent back to the verifier can be verified using the user's public key, the authentication is successful. To ensure that the public key used in the authentication belongs to the end user, the digital certificate of the user's public key, vouched by a trusted authority, is checked beforehand.

The main challenge of asymmetric cryptography is protecting the user's private key. Two methods are commonly used in asymmetric cryptography authentication systems. One method is to store the private key in a single piece, encrypted with a user-selected password. When the password is entered, the private key is retrieved in its entirety and can be used immediately. The other method is to use passwords as part of the private key [35, 66, 67]. This approach involves dividing the private key into two parts, where the user key is derived from the password, which can be memorised by the user, while the server key is stored and used within secure storage.

Authentication methods that use asymmetric cryptography have high complexity of implementation and high computational power consumption.

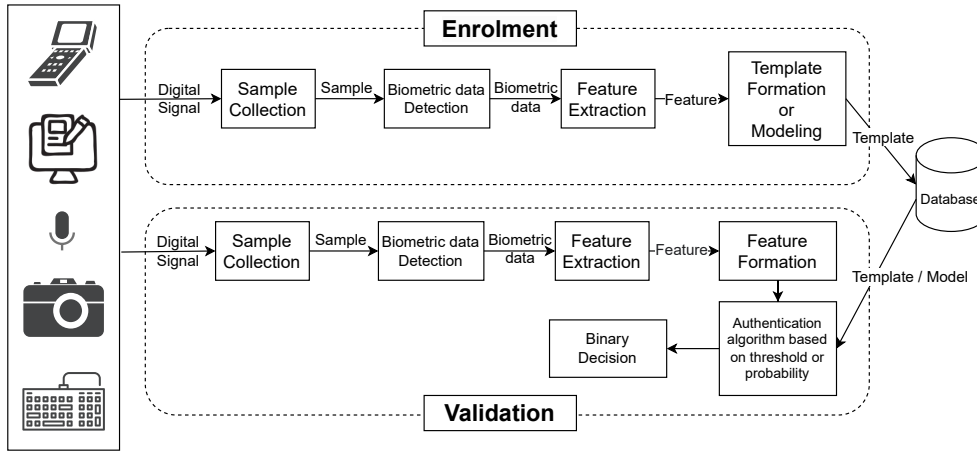


Figure 2.1: Workflow of a basic biometrics system

Additionally, maintaining the availability and security of the Public Key Infrastructure (PKI) can be very expensive. This is why it is not as widely adopted as symmetric cryptography, especially in resource-constrained environments such as IoT devices, where power consumption and cost are a concern.

2.4 Something you are

Compared to passwords and tokens, which can be easily shared, manipulated, or stolen, biometrics are considered a natural and reliable solution for identity authentication in various scenarios [91]. This is because biometric information, such as fingerprints, facial features, or iris patterns, are unique to each individual and difficult to replicate or steal. In 2019, the market for biometric systems was valued at more than \$30 billion, showing significant growth over the past two decades. The market has grown over 600 times in the last 20 years, with the technology becoming more sophisticated, accurate and cost-effective [20].

In a biometric system, the enrolment process begins with the use of capture equipment to collect samples of an individual’s biometric data, such as fingerprints, facial features, or iris patterns (shown in Figure 2.1). The samples are then processed by an image algorithm, which detects the biometric data and extracts the unique feature patterns. These feature patterns are then used to create a digital representation of the individual’s biometric data, which is stored in a database for later use during the authentication process.

To generate the representation, two techniques, template-based and model-based, are utilised in biometric systems for the enrolment process. In a

template-based enrolment, the extracted biometric data is formatted and protected using a cryptographic algorithm before being stored as a reference template in a database for the corresponding user. This template is then used as a reference point for comparison during the authentication process. In a model-based enrolment, the extracted biometric data is collected and used to train a machine-learning algorithm. The output of this algorithm, known as the user's model, is then stored in the database. The user's model is used to generate predictions during the authentication process by comparing the new data against the model. Template-based enrolment is considered to be more straightforward, as it simply stores the biometric data, while model-based enrolment is considered to be more flexible, as it allows for more sophisticated processing of the biometric data using machine learning algorithms.

The validation stage in a biometric system follows similar steps as the enrolment process; these include the capture of the biometric sample, processing the sample with an image algorithm, and extracting the feature patterns (shown in Figure 2.1). The extracted feature is then compared with the stored template or model of the user; if a match is found the user is granted access to the system.

In authentication methods that rely on factors of knowledge and possession, the secret is directly gathered from digital signals, whereas in biometrics authentication, data is captured by measuring the physical characteristics of an individual, which generates analogue signals that are then converted to binary sample data. However, the physical measurement process can introduce noise that leads to fuzzy output [163]. Since the comparison of two fuzzy outputs does not always lead to the binary results, two techniques are used in the validation stage depending on the enrolment approach: The similarity-based technique compares the current sample with a stored template to grant access if the similarity score is higher than a predefined threshold. The probability-based technique, on the other hand, uses a pre-trained model to determine access when using a model-based enrolment [17].

As biometric data are unique to an individual and sensitive to the illegal disclosure, security measures such as tamper-resistant hardware and cryptography are necessary for data protection. Additionally, passwords or other supervised methods may be used to control access to the data capture equipment.

Biometric systems can generally be divided into two categories: behavioural and physiological.

Behavioural biometrics. Behavioural biometrics is a method of identifying or verifying the identity of a user based on their unique behaviour patterns, such as typing rhythm, mouse movement, and navigational patterns. It is typically used as an additional layer of security in conjunction with traditional authentication methods, such as passwords or fingerprints, to provide more robust and secure authentication [17]. It can be used as point-to-entry authentication, but it also has the potential to be used for continuous authentication. A typical example of this is keystroke dynamics [18, 72, 95, 111], gait [71, 121, 156], and handwriting [23, 31, 61, 150]. Researchers often prefer to use machine learning algorithms to improve the accuracy of behavioural biometrics systems. This approach is effective in analysing and identifying unique patterns and traits in user behaviour to improve the accuracy of the system [40, 49, 133].

Keystroke dynamics is considered a natural way to verify a user’s identity because it relies on the user’s unique typing pattern. This method is non-intrusive, cost-efficient and transparent to the user [18, 33]. The feature extraction algorithm in keystroke dynamics is mainly based on the timing of the key pressing and releasing [26, 72, 111].

Gait recognition is a biometric technology that uses computer vision or sensory data to identify an individual based on their unique walking pattern [56]. It has potential applications in forensics, security, immigration, and surveillance. Cross-view gait recognition in the controlled environments is a promising strategy that enhances the accuracy of the system [149]. Sensor-based gait recognition, which uses sensors on mobile phones or wearable equipment to measure a user’s gait, has also gained increasing interest in recent years due to the widespread adoption of these devices.

Handwriting can be used not only to identify an individual but also to detect certain health issues [38, 53, 57]. Depending on the equipment used, handwriting recognition can be categorised as stylus-based or in-air movement-based. The former uses smartphones and pen-based devices to analyse dynamic features, such as handwriting velocity and duration [87]. Similar to gait recognition, the latter uses multiple cameras or motion capture sensors to detect movement, acceleration, orientation, and angular velocity for user identification and verification [61].

Physiological biometrics. Physiological biometrics measures the unique characteristics of a person’s body and is known as static authentication (also called point-of-entry authentication [17]), as any changes to the user’s body

after the authentication process will not be detected by the system. In this section, we introduce iris scanning, fingerprint recognition, and face recognition, which are three biometric systems that are commonly encountered in people's daily lives.

Iris is the connective tissue that surrounds the pupil of the eye. The structure of the iris forms during the first year of life, and the same characteristics of the iris remain genetically unchanged throughout the life of that person [51, 60]. Iris recognition is generally accepted to be the most accurate method to identify a person [102]. Millions of iris pairing tests report the false match rate of zero [50]. As an internal organ, the iris is difficult to trace and counterfeit compared to the fingerprint and face [51]. One limitation of iris recognition is that it requires controlled environments in order to accurately scan the iris pattern. As light stimulates the iris dilator muscle to constrict, the measurement result may vary under different illumination conditions. Another limitation of iris recognition is that the scanning process can cause discomfort for the user. This is because the subtle infrared illumination used during the process shines directly into the iris to acquire images of its detailed structures, which can be uncomfortable for some users.

Fingerprints are considered one of the oldest serving, most popular biometric technologies for person identification and verification [60]. The biometric system that uses fingerprints for identification and forensics is based on the fact that every individual has unique fingerprints on their fingertips. The pattern on the surface of the human fingertip is made up of ridges and valleys, and the uniqueness of each fingerprint is determined by the characteristics of its local ridges and their arrangement. Four types of fingerprint representation schemes are used for identification or forensics: greyscale image, phase image, skeleton image, and minutiae. Among these, minutiae-based representation is the most preferred scheme as it is highly distinctive and compact [178]. To prevent a fake fingerprint, which can be created with or without the knowledge of the person, from passing authentication, various methods have been proposed to increase the security of the system. These include liveness detection based on perspiration [123], skin deformation [145], quality-related features [65], and Local Phase Quantisation (LPQ) [69]. These methods help to ensure that the fingerprint being presented is from a live finger and not a fake or replicated one.

The facial recognition system is widely used because of its non-intrusiveness, universality and uniqueness [94, 120]. Compared to other biometrics, facial

recognition, which recognises the individual user based on his unique facial characteristics, has a high score in terms of availability, universality, and acceptability [94]. According to the report of the National Institute of Standards and Technology (NIST) in 2018, the accuracy of commercialised software to identify a person from the captured image has a false match rate of less than 0.2% [34]. Apple claims that face ID has less than one in one million probability of being unlocked by a random person’s face, which is much more secure than a 4-digit password (one in ten thousand) [1]. However, the accuracy of the side face recognition is only about 50% as reported in [141], since the performance of the system relies on the quality of the captured image, the light conditions, and the angles of the face [19].

Physically Unclonable Function While biometrics studies human beings, Physically Unclonable Function (PUF) measures the intrinsic random physical features of the object. A PUF leverages inherent variations in the manufacturing process or physical properties of a device to create a unique response or “fingerprint” for that specific device. These variations can include manufacturing defects, random variations in material properties, or environmental factors, which are used to generate the unique, irreproducible output based on the measured object’s physical characteristics.

PUFs are the biometrics of inanimate objects in many ways. For example, an object’s shape, size, colour, or texture can be used as a biometric for identification and authentication purposes [105]. Like biometrics for human beings, PUFs are accepted as promising solutions for digital device identification, authentication, and tamper detection, especially in the area of IoT [128, 137, 138].

The output of a PUF is typically based on a challenge-response mechanism. A challenge is provided to the PUF, and it generates a response that is unique to that specific device. The response is typically difficult to predict or replicate, making it useful for applications requiring device authentication, key generation, or secure identification. However, it is important to note that PUFs are not without challenges. Variations in environmental conditions or ageing effects can impact the reliability and stability of PUF responses. Additionally, PUF implementations must carefully consider the security and integrity of the challenge-response process to mitigate potential attacks or vulnerabilities.

The idea of PUFs was first proposed by Pappu et al. [122] in 2001. Several studies expand the idea, and different PUFs are proposed according to their

usage scenarios. Generally, PUFs can be categorised into two groups (as seen in Table 2.2): **non-electric** and **electric**.

In non-electric PUFs, the uniqueness of measured objects source from their random structure during the manufacture. For example, the speckle pattern of optical micro-structure can identify individual glass, while the fingerprint of a paper document can be extracted from its random fibre structure. Individual compact disc, radio frequency, magnetic media and acoustical equipment is distinguishable when they are measured by their unique features.

Electric PUFs, in which an electric quantity in the systems renders their PUF behaviour [105], can be further categorised into three subgroups: analogue electric, delay-based intrinsic and memory-based intrinsic. Analogue electric PUFs are based on the object's behaviour when it is measured by electric. Different transistors have various threshold voltages. The voltage drops and resistances help to spot the different power distribution systems. A random thickness of the protective coating on the read-proof hardware is used to generate a secure key when needed, while different LC circuits absorb various amount of power when set in a Radio frequency (RF) field. Delay-based intrinsic and memory-based intrinsic PUFs are measured by digital signals, including arbiter PUF, ring oscillator PUF, static random-access memory PUF, butterfly PUF, latch PUF and flip-flop PUF (as seen in Table 2.2).

We proposed a non-electric, paper-based PUF solution to address the issue of banknote counterfeiting. The details are presented in Chapter 4.

2.5 Summary

In this chapter, the current approaches in the field of authentication are reviewed. The design of an authentication system is based on three factors: something you know (knowledge-based authentication), something you own (token-based authentication), and something you are (biometrics-based authentication). For each of these factors, the current applications, challenges, and ongoing research are discussed and summarised.

As mentioned earlier, the choice of authentication factors depends on the problem being addressed and the environment in which it is used. In different circumstances, one factor may need to be replaced by another to ensure the security of the system. In the following chapters, we will examine various security threats and propose solutions to enhance the security of affected systems.

Table 2.2: Summary of PUFs

Source of randomness	Category	Measurement	Object	Nature of component	Research
Non-electric	Non-electric PUFs	Analogue	Optical	speckle pattern of optical microstructure on optical token	[86, 122]
			Paper	random fibre structure on paper document	[37, 47]
Electric	Analogue Electric PUFs	Analogue	Compact disc	deviation of lengths of lands and pits on compact disc	[76]
			Radio Frequency	near-field scattering of electromagnetic radiation waves	[54]
	Magnetic		particle patterns in magnetic media	[88]	
	Acoustical		unique frequency spectrum of acoustical delay line	[167]	
	Threshold voltage		variations on the threshold voltages of transistors	[104]	
	Power distribution		voltage drops and resistances in power distribution system	[78]	
	Coating		variability on passive dielectric sprayed coating	[162]	
	Resonant circuit		power absorbed by LC circuit ³ when it is in RF ⁴ field	[74]	
	Arbiter		random offset delay between two paths of arbiter circuit	[99, 103]	
	Ring oscillator		device dependent frequency of ring oscillator	[68]	
Memory-based intrinsic PUFs	Digital	SRAM ⁵	random power-up behaviour of SRAM cells	[79]	
		Butterfly	random power-up behaviour of FPGA ⁶ latch	[73]	
		Latch	random power-up behaviour of IC ⁷ latch	[147]	
		Flip-flop	random power-up behaviour of regular flip-flops on FPGA	[106]	

³ An electric circuit consisting of an inductor (L) and a capacitor (C)

⁴ Radio frequency

⁵ Static random-access memory

⁶ Field-programmable gate array

⁷ Integrated circuit

Chapter 3

Spooing against Spooing - Caller ID Verification

In this chapter, we propose Caller ID Verification (CIV), a novel solution to fight against caller-ID spoofing without the need for public key infrastructure (PKI). The solution is based on token authentication via challenge-response process. We first review the existing solutions, as well as their limitations. Then we introduce the design of our protocol and its prototypes, followed by the evaluation of the experiment results on various platforms. The evaluation shows that as a bottom-up solution, CIV can practically work across telephony networks, and can efficiently address the caller-ID spoofing with little to no cost by leveraging the techniques used in the spoofing attack.

3.1 Introduction

Telephone scams have been increasing at an alarming rate, especially during the Covid pandemic [117]. According to a survey in 2022, 1 in 3 Americans (33%) report having been targeted by phone scams, causing a total loss of nearly \$40 billion in the past 12 months [157]. In these scams, fraudsters frequently use *caller ID spoofing* to modify the caller's phone number in order to hide identity or to pretend to call from trusted sources (e.g., banks, tax revenue offices) [177].

In fact, caller ID spoofing is not anything new [20]. It has been known possible since the calling line identification (CLI) was first introduced in 1987 as a telephone service to display the incoming call number [100]. Sometimes, there are legitimate reasons for the caller to modify the number, e.g., showing a

single outgoing number for an organisation or a toll-free number for the callee to dial back. In traditional phone systems, modifying a caller ID requires special hardware or access to the telecom infrastructure. However, with VoIP [112], it has become effortless to modify the caller ID using the software. Once a modified number is permitted by the originating network, it will be trusted by the subsequent networks without validation. The ease of modifying caller ID has unfortunately enabled many frauds [63].

Besides inflicting financial losses on victims, caller ID spoofing attacks also reduce public trust in the telephone infrastructure. Ofcom, the telecom regulator in the UK, has been warning the public not to trust the caller ID display, and instead, users should “hang up and call the phone number to check whether the call was genuine” [117]. The Federal Communications Commission (FCC) has similar advice: “Don’t answer calls from unknown numbers. If you answer such a call, hang up immediately.” [63]. While these warnings serve to raise public awareness of the untrustworthiness of the caller ID display, they also have a significant side effect: according to YouGov, nearly 90% of people simply stop answering phone calls with unknown numbers [176]. Legitimate personal and business calls are blocked as well.

To restore the public confidence in caller ID display and stop the spoofing attacks, the FCC has recently mandated telecom providers in the US to implement STIR/SHAKEN by 30 June 2022 (gateway providers by 30 June 2023) [62]. STIR/SHAKEN represents a suite of protocols developed by an Internet Engineering Task Force (IETF) working group to combat caller ID spoofing. It works by attaching a digital signature as part of a Session Initiation Protocol (SIP) header together with a VoIP call. In practice, the digital signature needs to be accompanied by a certificate chain in the transmission to allow verification by the receiving party.

The FCC describes STIR/SHAKEN as “an industry-standard caller ID authentication technology” [62]. But actually, STIR/SHAKEN does not authenticate any caller ID. Instead, it authenticates the *originating carrier* where the call is made (or the *gateway carrier* for international calls that arrive inbound at the gateway), based on the carrier’s exclusive possession of a private signing key. Arguably this solves a different problem (which may not exist as any major issue in the telecom industry). The problem of authenticating caller IDs remains unsolved and is left to carriers: carriers must attest, as part of the digital signature, whether the caller is authorised to use the phone number. The difficulty here is to distinguish between legitimate and illegitimate

modifications of the number. Unfortunately, carriers do not always have the knowledge to tell them apart (if they do, caller ID spoofing would have been a much easier problem to tackle). For example, when a user modifies the VoIP caller ID to a mobile phone number, the carrier may not know if the user is authorised to use a number that belongs to a different carrier. To cater for this, STIR/SHAKEN introduces A, B, and C levels of attestations (for full, partial, and gateway attestations respectively) to indicate the carrier’s knowledge with different levels of confidence under different conditions. Level A attests that the caller is authenticated and is authorised to use the number; level B attests that the caller is authenticated but it is unknown if they are authorised to use the number; level C attests that the call is signed at a gateway without knowing if the caller is authenticated or is allowed to use the number. Interpreting validation results for signed calls of different levels has proved hard for users [140].

Apart from the ambiguity in the definition of “authentication”, STIR/SHAKEN suffers from two inherent limitations in the system design. First, it critically relies on trusted certificate authorities (CAs) to certify signing keys in a public-key infrastructure (PKI). (VoIP networks normally involve a PKI when using SSL/TLS to protect the data transmission in certain paths, but the PKI we discuss here in STIR/SHAKEN is a *new* infrastructure.) To spur the deployment of STIR/SHAKEN, the FCC has appointed several telecom companies in the US as the CAs. To comply with STIR/SHAKEN under the FCC rule, telecom providers in the US shall pay these CAs subscription fees for the issuance of certificates, normally based on the company’s annual revenue [152]. Although the FCC has been urging a global adoption of STIR/SHAKEN, it is extremely unlikely that the FCC-appointed CAs will be trusted by all other countries. (Similarly, if China appoints its own CAs, it is equally unlikely that they will be trusted by the FCC.)

Second, STIR/SHAKEN involves the transmission of digital signatures and a chain of certificates (several kilobytes) as part of the signalling data. The original design is to support only the SIP (VoIP) system, which leaves the traditional SS7 (landline and cellular) systems unprotected [140]. Although there are retrospective proposals to support STIR/SHAKEN in SS7 systems, e.g., by transmitting signature data out-of-band through a trusted third party, such a trusted third party is difficult to find in reality. The FCC acknowledges that “the STIR/SHAKEN framework is only operational on IP networks”, but requires that “providers using older forms of network technology to either

upgrade their networks to IP or actively work to develop a caller ID authentication solution that is operational on non-IP networks” [62]. However, the “caller ID authentication solution” for non-IP networks has not been specified, which leaves a gap in the regulatory rules [177].

Public data show that since STIR/SHAKEN was mandated in June 2022, this solution has not been as effective as expected. First of all, after the mandate of STIR/SHAKEN, the number of robocalls actually went up and reached a record of 5.5 billion calls in October 2022 [154]. Many of the robocalls are now signed with STIR/SHAKEN to look more legitimate. Among all the signed calls with the B-attestation, about a quarter are robocalls; for calls signed with the C-attestation, about a third are robocalls [155]. Many of these signed robocalls present a different caller ID from where they are calling. The statistics [153] also show that although nearly 70% of the outbound VoIP calls are signed with STIR/SHAKEN, only 15-24% of the calls received by the terminating networks have valid signatures; for many calls, the signatures are removed as they traverse intermediate non-IP networks. It is also reported that many calls are signed with wrong attestation levels, which should not be surprising given that STIR/SHAKEN only authenticates the “carrier” and the attestation of the caller ID is entirely based on the carrier’s “word of mouth” [177].

STIR/SHAKEN represents an industry-driven approach: representatives from several telecom companies form the core of an IETF working group to specify a signature-based framework called STIR, followed by implementation details called SHAKEN [85]. Some of these companies were later appointed by the FCC as the CAs. Besides serving a trusted role, there is also an economic incentive to be a CA, since other companies are obliged to pay them subscription fees in compliance with the FCC regulation. So far, this industry-driven solution has received limited scrutiny from the security research community. In particular, alternatives to STIR/SHAKEN have not been sufficiently studied. Given the fundamental limitations of STIR/SHAKEN and the prevalence of caller ID spoofing attacks, we believe that it has become more urgent than ever to explore more secure and effective solutions to stop spoofing attacks. Since many spoofing calls originate from overseas [140], it is crucial that we address it as a *global* problem rather than any regional or country-specific problem. Furthermore, instead of treating caller ID authentication for SIP and SS7 (non-IP) networks as separate problems, we propose to tackle them together.

Table 3.1: Comparison with related works

Scheme	Auth subject	Delay	Properties										
			Requires TTP	Requires registration	Heuristic authentication	Deterministic authentication	Dif. legitimate spoofing	Works with SIP	Works with SS7	SIP phone prototype	Mobile phone prototype	Landline phone prototype	Hetero. networks experiment
Top-down													
STIR/SHAKEN [85]	Carrier	N/A	●	●	●	●	●	●	●	●	●	●	●
AuthLoop [129]	Call center	4.8-8.9 s	●	●	●	●	●	●	●	●	●	●	●
AuthentiCall [130]	Caller ID	1-1.41 s	●	●	●	●	●	●	●	●	●	●	●
Bottom-up													
CallerDec [114]	Caller ID	8.4-20 s			●	●	●	●	●	●	●	●	●
CEIVE [55]	Caller ID	10-23 s			●	●	●	●	●	●	●	●	●
CIV (this paper)	Caller ID	4.7-29 s	●		●	●	●	●	●	●	●	●	●

●/○: have a full/partial *undesirable* property. ●/○: have a full/partial *desirable* property.

In this chapter, we propose Caller ID Verification (CIV). CIV authenticates caller ID based on a challenge-response protocol. As we will explain, it does not require a PKI and works with existing heterogeneous networks (SS7/SIP), hence addressing the two major limitations of STIR/SHAKEN. Our solution does not require any trusted third party and can be deployed by updating the software on the user’s phone. This follows a bottom-up approach as opposed to STIR/SHAKEN’s top-down approach. There are previous bottom-up proposals [55, 113], which *probabilistically* infer the authenticity of the caller ID based on heuristics. By contrast, CIV *deterministically* authenticates the caller ID based on a challenge-response protocol.

3.2 Related work

We broadly divide previous solutions into two types: top-down and bottom-up. A top-down solution involves introducing a trusted third party (TTP), while a bottom-up solution does not need one. Table 3.1 presents an overview of the representative schemes for each type. There are several *undesirable* properties that we wish to avoid. In particular, we want to avoid relying on a TTP¹ and an extra registration process if possible. We consider two types of authentication: heuristic and deterministic. The former performs authentication *probabilistically* based on heuristics (e.g., network characteristics and prior training data), but the authentication result may vary under different

¹Recall a TTP is “a third party that can break your security policy” [20].

conditions and troubleshooting failures can prove difficult. The latter performs authentication *deterministically* based on explicit rules (e.g., possession of a secret key or a secret number), which is considered more desirable. In addition, it is necessary that a caller ID authentication solution should distinguish legitimate and illegitimate modifications of a caller ID. Finally, the solution should work with heterogeneous networks (SS7/SIP) and be tested in such conditions.

3.2.1 Top-down

STIR/SHAKEN was jointly developed by several telecom companies. To our knowledge, there is no peer-reviewed academic paper on the original proposal of STIR/SHAKEN, but the scheme is described in a series of IETF RFCs [85]. STIR/SHAKEN critically relies on trusted certificate authorities (CAs) in a PKI to issue digital certificates for the subscribed carriers. It authenticates the carrier (not the caller) based on their possession of a unique private signing key. It is designed to work with an IP-based SIP network and has been commercially deployed on VoIP and (IP-based) mobile phones. STIR/SHAKEN does not authenticate any caller ID, or distinguish legitimate/illegitimate spoofing; it relies on the carrier’s ‘word of mouth’ for attesting to the authenticity of the caller ID.

AuthLoop was proposed by Reaves et al. at USENIX Security’16 [129] based on adapting TLS 1.2 to a telephony network. Same as STIR/SHAKEN, it requires a PKI (called a “Telephony PKI” in their paper). This scheme is designed for a client-server setting, where the server (a call centre) is authenticated based on digital signatures, but the client caller is not authenticated. The system does not work with existing SIP or SS7 signalling. A prototype of AuthLoop was implemented between PCs (as clients and servers) but not on telephones or connected to telephone networks. The verification delay was reported to be 4.8–8.9 seconds.

AuthentiCall was proposed by Reaves et al. at USENIX Security’17 [130]. To address the lack of caller authentication in AuthLoop, the authors propose to introduce a trusted central server. Telephone users need to register their numbers (including legitimately modified numbers) with this server in an enrolment process to obtain certificates. When the user makes a call, the phone first contacts the central server, and then the server mediates the authentication process between the caller and the callee. A prototype was implemented using

an Android app connected to a server through the Internet, but the app was not connected to telephone networks. The authors report 1–1.41 seconds delay for the verification process (and 22–25 seconds for the enrolment process, which is not included in Table 3.1).

3.2.2 Bottom-up

CallerDec was proposed by Mustafa et al. at DSN’14 [113] (with a journal version in 2016 [114]). It is designed for a circuit-switched telephone network. It authenticates caller ID by calling back the number. This scheme is the closest to ours. It has the advantage of not needing any PKI or extra registration. However, CallerDec authenticates the caller ID *probabilistically* based on applying heuristics to infer the caller’s state. The authentication outcome critically depends on physical network characteristics, such as timing in the call setup, as well as the choice of classifiers and the prior training data. CallerDec, on its own, does not distinguish legitimate and illegitimate spoofing; instead, it requires the user to press keys on the keypad to indicate if a spoofed call is legitimate or not. A prototype based on an Android mobile phone was implemented and tested in a circuit-switched cellular (3G) network. The verification delay was reported to be 8.4 (call-based) and 20 (SMS-based) seconds.

CEIVE was proposed by Deng et al. at MobiCom’18 to authenticate caller ID in a 4G wireless network [55]. It uses a similar call-back idea as CallerDec to infer the caller’s state. The main difference is that it is a *callee-only* solution without needing any cooperation from the caller. Same as CallerDec, it authenticates the caller ID *probabilistically* based on heuristics. The authors acknowledge that this is reliable for a single carrier, but the performance varies for a different carrier or across carriers. CEIVE does not distinguish legitimate and illegitimate spoofing. It was implemented on an Android phone and tested on a 4G network. A verification delay of 10–23 seconds was reported.

3.2.3 Comparison between STIR/SHAKEN and CIV

CIV authenticates the caller ID of individual call through a challenge-response protocol. It can be implemented by using a combination of CLI and DTMF, which are nearly universally supported in existing telephone networks. For telephone systems that use CNAM databases for sharing caller names, we require registering the support for CIV by appending a flag in the caller name

in CNAM; in other systems (in particular VoIP), the caller name and the flag can be directly transmitted along with the caller ID. CIV distinguishes legitimate and illegitimate spoofing; it supports both SS7 and SIP; and it has been implemented on all three types of phone systems and tested across heterogeneous networks. As a comparison, STIR/SHAKEN only authenticates the carrier, while the task of identifying whether an outgoing call is legitimate or not relies on the carrier. It requires the setup of PKI and trusted third-party, and it supports SIP only (seen in Table 3.2).

We do not claim that our solution is perfect (see limitations in Section 3.7.5), however, CIV is the first solution that has been shown to work on all three phone systems across heterogeneous networks. While many people believe that STIR/SHAKEN is the ultimate solution, our work shows that relying on a PKI (or TTP) is not a necessity. We hope this will encourage more research into tackling this important real-world problem from the bottom up.

Table 3.2: Comparison of STIR/SHAKEN and CIV

	STIR/SHAKEN (Top-down)	CIV (Bottom-up)
Method	Security primitive (non-interactive)	Security protocol (interactive)
Authenticated party	Carrier	Caller
Distinguish legitimate/ illegitimate spoofing	No	Yes
PKI & trusted third parties	Yes	No
Date transmission	Signature + certificate chains	Decimal digits
Telephony networks	SIP-only	All networks
Overhead	N/A	4.7 s

3.3 Background

Over time, telecommunication systems have evolved to be exceedingly complex, spanning heterogeneous networks. In this section, we explain the relevant background below.

3.3.1 Heterogeneous networks

Public Switched Telephone Network (PSTN). PSTN represents circuit-switched telephone networks, which are traditionally connected by copper wires to transmit analogue voice signals. Later, Integrated Services Digital Network (ISDN) was developed to allow the digital transmission of data over copper lines, which also introduced new telephony features, such as voice mail, call forwarding, and caller ID/name display. The Signalling System 7 (SS7) is the dominant protocol to control calls in PSTN, including call setup, connection, tear-down, and billing [58]. Maintaining dedicated wires for PSTN has become increasingly expensive. BT has announced a plan to phase out PSTN by 2025 [7]. However, for the foreseeable future, PSTN will still be used in many parts of the world [161].

Cellular Network. Cellular technology allows transmitting voice data wirelessly rather than over wired connections. The first generation of mobile phones (1G) used analogue signals, and the handset simply sent the serial number in the air, which was vulnerable to cloning attacks [20]. The second generation (2G) adopted digital technology with encryption. GSM was introduced in 1992. Each GSM handset has an embedded SIM card that stores a unique *international mobile subscriber identification* (IMSI) number and a secret key for authentication and encryption. In 1993, cdmaOne was introduced as another 2G standard based on CDMA. The third generation (3G) entered service in 2003, providing a faster data rate by adopting a spread-spectrum technology. The fourth generation (4G) was rolled out in 2009, followed by the fifth generation (5G) in 2019. While SS7 has been used as a core signalling protocol in 2G and 3G, it has been replaced by Diameter in 4G and 5G. (The newer 5G Core network uses HTTP/2 signalling). For interoperability, 4G and 5G still need to interconnect with previous-generation networks before they are phased out. Today, SS7 is still supported by nearly all wireless carriers [29].

Voice-over-IP (VoIP). In VoIP, voice data are digitised, compressed, and routed over the Internet [20]. Currently, the dominant signalling protocol in VoIP is the Session Initiation Protocol (SIP). SIP borrows many of its syntax and semantics from HTTP (hypertext transfer protocol), but it also inherits many weaknesses of HTTP [171]. For example, the SIP header is unprotected, which makes it trivial to spoof a caller's identity. A primary benefit of VoIP is that the running cost is low as no dedicated wires are needed. Unlimited local and long-distance calls are often included in a bundle, which greatly saves costs

for users, but at the same time also enables spammers and scammers to do robocalling (with a spoofed caller ID) at little cost [160]. The interconnection between VoIP and other networks (PSTN and cellular) is achieved through gateways, which are responsible for the conversion of different signalling types.

3.3.2 Caller ID spoofing

There are two pieces of information to identify a caller: a caller ID (phone number) and an optional caller name. A spoofing attack involves modifying the phone number, name, or both. Calling line identification (CLI) was first introduced in 1987 in the US as a telephone service to allow displaying of the caller's number on the receiver's phone. Later, Caller Name (CNAM) was introduced as a separate service to allow the calling party to specify a name associated with the caller ID. Typically, a caller name is limited to 15 characters, including alphanumerics, commas, and spaces. Special characters (e.g., &, @, etc) are not allowed for ordinary users but are permitted for business users. In the US, the originating carrier does not normally send the caller's name when initiating a call. Instead, the carrier registers the caller name in shared CNAM databases. It is the responsibility of the terminating carrier to look up the CNAM databases (paying a "dip" fee) based on the received number and deliver the retrieved caller name to the callee's phone. When access to CNAM databases is not available (e.g., outside the US), the caller name is sent directly to the called party together with the caller ID; in the case of VoIP, it is included in the SIP "From" header.

Modifying the caller ID/name is always possible in a telecommunication system. As part of CNAM, a user can freely modify the caller name (subject to basic checks such as the length). Modifying the caller ID is less straightforward. In PSTN, the telephone user is authenticated to the local switch through a dedicated wire. The caller ID is generated by the local switch, and an ordinary user cannot modify it. However, if the calling party is behind a Private Branch eXchange (PBX) which is connected to a local switch via Primary Rate Interface (PRI), PBX can modify the caller ID [41]. The modified caller ID is usually accepted by the switch without validation and passed on to the rest of the network. In a cellular network, the IMSI number stored in the SIM card is used by the switching centre to identify the caller ID of a cellular user [135]. Similar to PSTN, an attacker has a limited chance to modify the caller ID. In VoIP, the caller ID, together with the optional caller name, is specified as part

of the SIP “From” header. However, the header is unauthenticated. Hence, a user can arbitrarily change the caller ID and the caller name in the header. Once the modified header is permitted by the originating carrier, it will pass the subsequent networks without validation.

Carriers often allow users to use a different caller ID than what they are assigned with [113]. According to the Truth in Caller ID Act of 2009, modifying the caller ID is permitted, unless it is done “with the intent to defraud, cause harm, or wrongfully obtain something of value” [101]. There are legitimate reasons for modifying the default caller ID [113], e.g., to provide a toll-free number for the receiver to call back, or to display a central phone number for all outgoing calls from an organisation. We do not regard these cases as spoofing attacks and will distinguish them from illegitimate spoofing based on whether the caller possesses the claimed phone number (details in Section 3.4).

3.3.3 Dual-tone multi-frequency

Dual-tone multi-frequency (DTMF) is one important element in our solution, so we explain it in more detail here. DTMF is a telecommunication signalling system that was first invented by the Bell System in 1963 [165], and then standardised and adopted globally. It uses a pair of the 8 predefined frequencies (hence called dual-tone) to encode one of the 16 phone button presses, including digits (0-9), *, #, A, B, C and D. DTMF is especially useful for enabling users to provide input of short digits over a phone line, e.g., to enter a PIN in telephone banking, or navigate menus in an interactive voice response (IVR) system.

DTMF signals can be transmitted through the voice channel as part of an audio stream (in-band) or a separate control path (out-of-band), depending on the underlying telecom system (see Figure 3.1). Gateways are responsible for the seamless transmission of DTMF tones across different phone networks, handling the conversation between in-band and out-of-band signals wherever necessary.

Circuit switched network. In a PSTN network (analogue lines or ISDN), DTMF tones are typically transmitted as in-band signals in the range of human voice frequencies (300 – 3400 Hz). However, in a cellular network like GSM, the data transmission rate is significantly slower than in a wired network. To achieve a comparable voice quality in cellular networks, the voice data are heavily compressed before transmission. Widely used compression algorithms

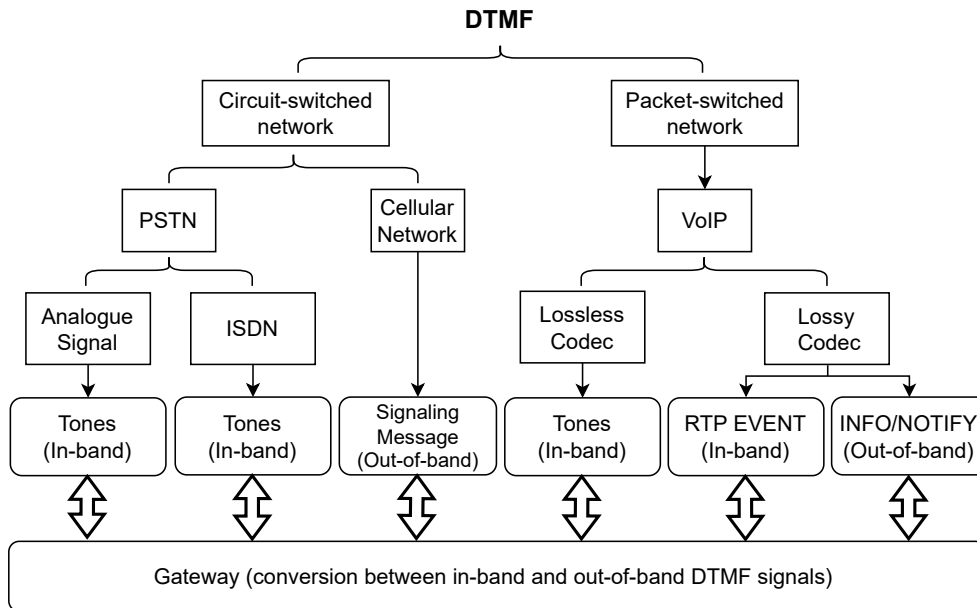


Figure 3.1: Overview of DTMF transmission

remove frequencies in the audio data that are insensitive to human ears so a phone conversation is not affected, but the loss of certain frequencies impacts the decoding of DTMF tones. To ensure reliable transmission of DTMF in a cellular network, DTMF data are transmitted *out-of-band* as signalling messages via a control channel separate from the voice communication [165].

Packet switched network. In a VoIP network, DTMF tones can be transmitted in-band together with the media stream if no compression or a lossless codec (e.g., G.711) is used. When a lossy codec is used (e.g., G.729), there are two ways to transmit DTMF: 1) in-band as part of the media stream but in a special Real-time Transport (RTP) **Event** packet based on RFC 4733; 2) out-of-band in a SIP **INFO** (RFC 6086) or **NOTIFY** (RFC 3265) message. Either way, only the digital values of DTMF (not the analogue tones) are transmitted through the networks. When needed, the analogue sound of the DTMF tones is played locally on the phone to inform users about sending/receiving DTMF.

3.4 Our proposed system

Figure 3.2 shows an overview of our solution in heterogeneous networks. Our proof-of-concept implementations only require updating the software on the user’s phone; implementation details for landline (in conjunction with a trueCall

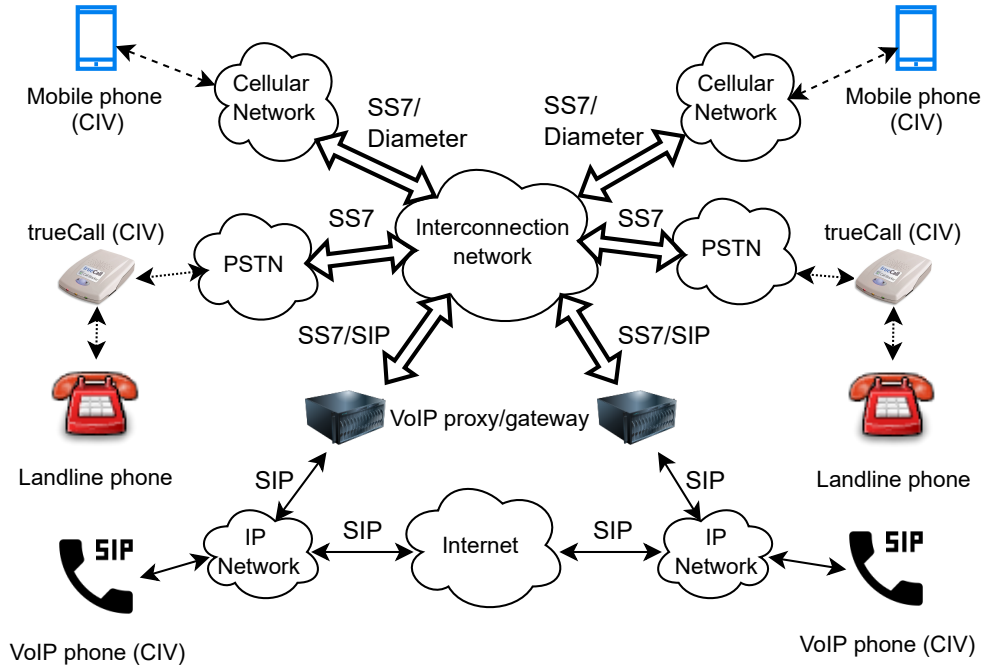


Figure 3.2: CIV for heterogeneous telecommunication networks

box), cellular and VoIP phones will be explained in Section 3.5. In Section 3.7.4, we will explain how to optimise the performance of CIV by integrating it into the networks.

3.4.1 Threat model

In our threat model, we assume that the attacker is able to arbitrarily modify the caller ID/name when initiating a call. The modified caller ID and name are permitted by the originating carrier and pass through subsequent networks. However, we assume the attacker is not able to intercept calls in the telecommunication system. We note that a powerful adversary can intercept calls through the Law Enforcement Monitoring Facility (LEMF), SIM swap, and SS7 hacking [20], but this is beyond the capability of ordinary telephone scammers behind the number spoofing attacks.

3.4.2 Protocol description

We name the caller ‘Alice’, the callee ‘Bob’, and the spoofing attacker ‘Eve’. Eve tries to impersonate Alice by spoofing Alice’s caller ID as his outbound number. Here, we focus on the spoofing of the caller ID (number) not the caller

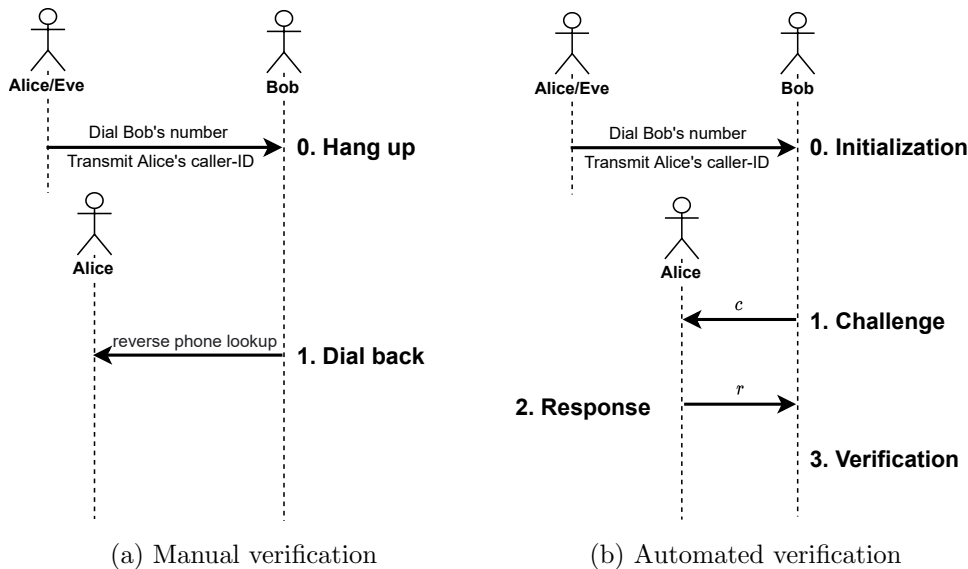


Figure 3.3: Intuition behind CIV

name, since the latter can be addressed by caller ID filtering [13] or reverse number look-up [14].

In CIV, the authentication of the caller ID is based on the possession of the claimed phone number. The intuition follows from how people verify the caller ID of an incoming call in real life by manually calling back the number. As an example, suppose that Bob receives a call displaying his bank’s telephone number, which matches exactly the number shown on the back of Bob’s bank card. But the number is actually spoofed². Ofcom advises that Bob should hang up and call back the bank’s contact number [117]. The manual call-back verification is slow, and tedious and may incur a charge for Bob. The goal of CIV is to turn this manual verification into an automated one (see Fig. 3.3) with minimum delay and cost.

In CIV, we assume that Alice (caller) actively wants to have her caller ID verified. The rationale is that she wants Bob to see the *verified* status of her caller ID so that Bob is more likely to answer the call. The cost to Alice is a possibly longer duration of a call, needed for carrying out the verification process. We assume that Alice is willing to pay for this cost. Here, the cost


²In many countries, the phone numbers printed on the back of bank cards are for *inbound* calls only: customers use them to contact banks but banks *never* use them to call customers. Ofcom includes these numbers in a “do not originate” (DNO) list and requests telecom operators to block them at the network level for *outbound* calls. However, complying with DNO is not (yet) a legal requirement. Blocking these calls by operators is not guaranteed.

and Alice’s incentive are aligned. In the initial call flow, Alice dials Bob’s number and transmits her caller ID/name. In this flow, Alice indicates support for CIV, e.g., by adding a flag in the caller name. In practice, this flag may be a special character or a string of characters added during the registration of the caller name in CNAM databases (in our SIP-based prototype of CIV, we indicate support for CIV by transmitting a caller name appended with ‘*’; for other prototypes, we assume that a flag has been added to CNAM).


Upon receiving Alice’s call, CIV on Bob’s phone holds the incoming call, and performs the following challenge and response protocol to verify the caller ID.

1. **Challenge** – Bob’s CIV calls back the number on the caller ID display and transmits a random n -digit number as a challenge c . In our design, we choose $n = 4$ (see Section 3.6 for evaluation of other values). If Alice is a genuine caller, she will receive c .
2. **Response** – To prove the possession of the purported caller ID, Alice’s CIV simply sends the same n -digit number as the response r to Bob on a separate channel.
3. **Verification** – After receiving r , Bob’s CIV checks if $r \stackrel{?}{=} c$. If they are equal, the caller ID is “verified”; otherwise, it is “not verified”.

After the above challenge-response process, Bob’s phone starts ringing, displaying the caller ID together with the verification status. Figure 3.4 shows an illustration of the CIV user interface on a mobile device. For a landline phone without a display, we play an audio message to inform the user if the caller ID has been verified when they pick up the phone (we have implemented this by using a trueCall box which we will explain later). We emphasise that even if the caller ID verification fails, CIV still connects the call (Figure 3.4.b). In other words, CIV never blocks a call; it only adds additional information about the *verified* status of the incoming caller ID. The caller will likely experience a longer delay in connecting the call, but the delay is not perceivable by the callee. We have implemented this operation mode for all three types of phones (landline, cellular and VoIP). It is possible to implement an alternative operation mode, in which Bob’s phone rings as soon as it receives a call and performs verification in parallel. This is to support an “emergency call” scenario (e.g., for 911) as described by Mustafa et al. [113]. While the implementation of this “emergency call” mode is possible in CIV, one might



(a) Incoming call is verified



(b) Incoming call fails to be verified

Figure 3.4: Illustration of the CIV user interface

question if an emergency service such as 911 really needs this verification process since they already have privileged access to all the call detail records and can trace any incoming call (if needed). For this reason, we do not propose this emergency scenario as any main operation mode in CIV.

Our protocol leverages the call-back session as a trusted channel to reach

the owner of the purported caller ID (the attacker cannot intercept the call based on the threat model in Section 3.4.1). Note that CIV on the callee’s phone only starts the verification process if the caller indicates support for CIV, since completing this process requires the cooperation of the caller. The user may further configure CIV on the phone to perform verification only when the displayed caller ID is a domestic number or a non-premium number.

The main challenge in realising CIV is how to transmit the challenge and response in heterogeneous networks reliably, with minimum delay and cost. As explained in Section 3.3, caller ID spoofing has always been possible at a system level in all telecom networks. Here, we propose to leverage the facility of number spoofing to build a defensive mechanism to combat caller ID spoofing attacks, i.e., spoofing against spoofing. More concretely, Bob uses c as his (spoofed) caller ID and makes an *abandoned* call to Alice. Here, there is no call charge to Bob. Alice’s CIV receives a *missed* call and recognises that it is a verification call (based on the format of the number but we can also make it explicit, e.g., by including the information in the accompanying caller name). The challenge c is extracted from the caller ID. The missed call allows Bob to send a 4-digit message through telephony networks without cost. Alice may use the same spoofing method to send back the response, but we propose to transmit the response via the initial call using DTMF since it is much quicker (see Section 3.6 for experimental results).

3.4.3 Distinguishing legitimate/illegitimate spoofing

An effective caller ID authentication mechanism should be able to distinguish a legitimate modification of a caller ID from an illegitimate one (i.e., a spoofing attack) [113]. By design, STIR/SHAKEN does not make this distinction and relies on the carrier’s ‘word of mouth’ in the attestation. CIV distinguishes these cases based on whether the caller possesses the phone number, hence being able to respond to a challenge sent to that number. As an example, a VoIP phone user wants to modify the caller ID to his mobile phone number. To support CIV, he simply needs to configure the mobile phone to forward the verification call containing the challenge to the VoIP phone. In another example, a caller is behind a PBX and wants to use a single outgoing phone number for the organisation. In this case, PBX needs to keep the state of currently active outbound calls. It can assign a *random* index (say three digits for up to 999 simultaneous calls) for each call and include the index in the caller name




Figure 3.5: Equipment setup

of an outbound call so that when a verification call (containing the challenge and the index) arrives, it can forward the challenge to the corresponding caller. The CIV on the caller’s phone will process the challenge automatically.

3.5 Prototypes

In this section, we present proof-of-concept CIV prototypes for landline, cellular and VoIP phones across heterogeneous networks (SIP/SS7).

3.5.1 Overview

We have implemented prototypes of CIV for all three types of phones: landline, cellular and VoIP phones (see Figure 3.5). Our proof-of-concept prototypes are done under two levels of constraints. The first is on an infrastructural level: we have no cooperation from telecom providers. Therefore we can only update the software on the user’s phone. The second is on a platform level: we are constrained to work with only the available APIs provided by different phone development platforms, as explained below.

1. **VoIP platform.** We use the Ozeki VoIP Software Development Kit to develop a Windows-based SIP phone that implements CIV. The SIP phone works with a commercial third-party VoIP server, as well with our own VoIP servers, which we set up by using the open-source Asterisk software. Our VoIP servers are connected with the public SS7/SIP networks through SIP trunking.
2. **Android platform.** We use Android phones (Nexus 5) and the third-party phone API available on Android (6.0.1) to build a phone app that implements CIV. The Android API supports the call waiting function (which allows placing a call on hold to engage in another call) but does not support transmitting DTMF in a call.
3. **trueCall box.** We use a third-party nuisance call-blocking device, called trueCall [13], to implement CIV and connect the modified trueCall box to a landline phone. This allows us to control calls to a landline phone without having to modify its firmware. In contrast to Android, trueCall supports sending DTMF during a call but does not support the call-waiting function.

Figure 3.6 shows an overview of the CIV prototypes developed under various constraints. In general, there are two ways to transmit the challenge/response: using 1) spoofed CLI; 2) DTMF. The first method essentially uses CLI as a side channel but requires access to the facility of modifying the caller line identity: i.e., caller ID/name. This is only possible on our SIP platform. It is possible to modify the CLI for PSTN and cellular phones, but this needs to be done at the local switches or switching centres. The second method transmits data through DTMF. Without any local platform constraint (like on our SIP platform), this can be done efficiently by adding only one call setup. However, with the constraints of the trueCall/Android platforms, we need to add two call setups. This does not stop the proof-of-concept demonstration of CIV, but it increases the delay (see Section 3.6 for details).

3.5.2 CLI-based prototypes

These prototypes assume access to the facility of modifying the CLI (Case 1 in Figure 3.6). Based on our SIP platform, we have done two prototypes, ‘CLI/CLI’ and ‘CLI/DTMF’, which use two different methods to send ‘challenge/response’




Figure 3.6: Overview of CIV prototypes

respectively. (‘DTMF/CLI’ is another possibility but is not recommended as we explain below.)

Figure 3.7 summarises the flows in the ‘CLI/DTMF’ implementation using SIP phones. In the initial call setup, Alice (caller) indicates support for CIV, e.g., by appending a flag in the caller name. When CNAM databases are used for registering caller names, a special flag for a caller name indicates that an associated phone number is ready to be verified. In Step 1, when Bob receives an initial call from Alice, CIV on Bob’s phone answers the call and puts it on hold. In Step 2, it then makes a verification call to the displayed incoming call number using a spoofed CLI (the challenge c). The verification call is immediately abandoned by Bob’s CIV, and the initial call is taken off hold. The CIV on Alice’s phone can distinguish it from ordinary missed calls and retrieve the challenge c (random 4 digits that are received as the CLI). This process is handled transparently (users do not need to see the missed call). Finally, in Step 3, CIV on Alice’s phone sends a response r (same 4 digits) through DTMF via the initial channel established in Step 1. When receiving the response r , CIV on Bob’s phone checks if it is equal to the challenge c . It concludes that the caller ID is verified if they are equal, and unverified otherwise. At this point, CIV starts ringing with the display of a caller ID along with the verification status.

Due to the limited resources assigned to the individual subscriber, when CIV working on the end user’s phone, only one CIV can be processed at any given time. Thus, the challenge secret is unnecessary to identify itself³. When working with PBX, the identity of the challenge in CIV is necessary. We have discussed this scenario in Section 3.4.3.

In the design, we use the CLI spoofed verification call to pass the challenge secret, instead of adding an extra field in the SIP INVITE message. We found that a lot of SIP service providers prefer to use back-to-back user agent (B2BUA) to act as an intermediary between two or more SIP endpoints in a communication session. It sits in the middle of the SIP signalling path, receiving SIP messages from one endpoint, processing them, and then forwarding them to the other endpoint(s). Using B2BUA has several benefits, including complete control over the SIP session, and satisfying support for the protocol interworking. As the B2BUA maintains separate signalling and media channels with each endpoint, the only information pass from one end to the other is the caller’s CLI.

An alternative implementation is to use ‘CLI/CLI’ to send ‘challenge/response’ respectively. The first two flows are the same as ‘CLI/DTMF’ in Figure 3.7. However, in Step 3, the response r is sent to Bob as the modified CLI in an abandoned call, instead of using DTMF via the initial call. We have implemented this alternative ‘CLI/CLI’ approach, however, using ‘CLI/CLI’ incurs a longer delay than ‘CLI/DTMF’ because sending r through CLI requires a full call setup (see experimental results in Section 3.6).

In theory, it is also possible to use ‘DTMF/CLI’ to send ‘challenger/response’ respectively. However, this option is not recommended due to the possible call charge to the callee (for sending the challenge through DTMF). Furthermore, it has no performance advantages over other options.

Modifying CLI. In our CLI-based prototypes, we need access to the facility of modifying CLI. On an Asterisk VoIP server, this can be easily done by invoking the `Set(CALLERID())` method to modify the default caller ID and name for an outgoing call in a configuration file. This is normally how the CLI spoofing is carried out from a VoIP platform. However, the modification is statically defined in a configuration file while we need to dynamically modify the CLI during a call. Also, for the proof-of-concept prototypes, we want to limit ourselves to modifying the client phone only (not the server). We use

³For example, concatenating the last three digits of Bob’s number to the spoofed CLI.

Ozeki VoIP SDK to build a client softphone to communicate with a VoIP server based on the standard SIP protocol. The Ozeki SDK does not support modifying the caller ID directly from the client, but it allows modifying the caller name by setting a new value for the `CallerDisplay` property. With access to this function, we set a special (*) flag in the caller name in the initial call to indicate support for CIV, and embed the challenge/response digits in the modified caller name.

Transmitting DTMF. Ozeki supports transmitting DTMF in-band as RTP EVENTS. We invoke the `StartDTMFSignal()` and `StopDTMFSignal()` functions to send DTMF tones as part of the media stream. In telecommunication terminology, the duration of a DTMF tone is called the ‘mark’ time while the gap between two consecutive DTMF tones is called the ‘space’ time. Telecommunication standards require the mark and space to be at least 40 ms (RFC 4733); in our SIP prototypes, we set both to be 50 ms. We note that the Ozeki SDK does not allow sending DTMF through an *out-of-band* channel from the client. However, if CIV is implemented on the SIP server, sending DTMF *out-of-band* will be possible, e.g., as an INFO or NOTIFY message.

3.5.3 DTMF-based prototypes

In our existing telephony systems, the end phones may not have access to the facility of modifying the CLI, or the local carrier may not permit such modification. In this case, it is still possible to implement CIV by using DTMF to send both the challenge and the response.

No local platform constraints

We assume that the user’s phone supports the *call waiting* function (hence it can hold an incoming call) and is able to transmit DTMF during a call. We use SIP phones to implement a prototype for this case (i.e., Case 2 in Figure 3.6).

Figure 3.8 summarises the implementation of this prototype which uses ‘DTMF/DTMF’ to send ‘challenge/response’ respectively. Upon receiving an initial call from Alice with a flag in the caller name, CIV on Bob’s phone holds the call and meanwhile starts a *verification call* to send a challenge c through DTMF. Once CIV on Alice’s phone receives the challenge, it hangs up the verification call and sends a response r using DTMF through the initial call. When CIV on Bob’s phone receives r and determines that it is equal to the challenge c , it confirms Alice’s caller ID as *verified*; otherwise, it is *unverified*.

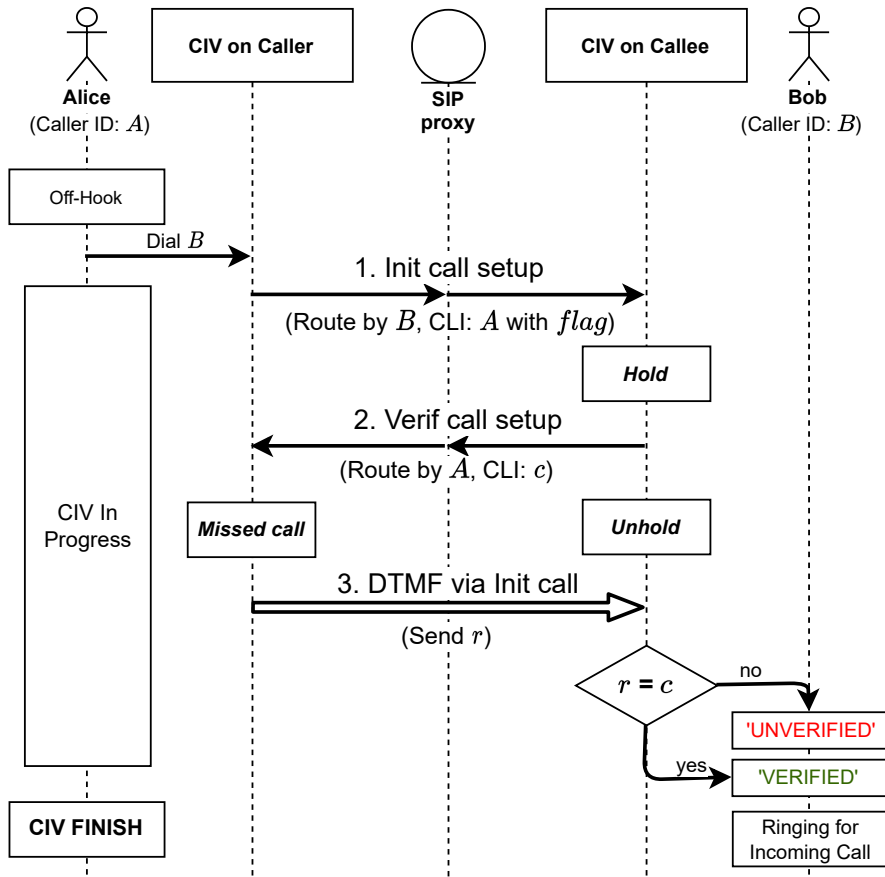


Figure 3.7: CIV using CLI/DTMF to send challenge/response

Finally, Bob’s phone starts ringing, with a display of the caller ID along with the verification status. This prototype is reasonably efficient as it needs only one additional call setup to the challenge (the response is sent through an existing initial call channel rather than a new call).

With Local platform constraints

For non-SIP phones (landline/cellular) used in our prototyping, there are certain platform constraints (Case 3 in Figure 3.6). In particular, the trueCall box does not support the call waiting function (while it supports sending in-call DTMF). The Android third-party phone API does not support sending in-call DTMF (while it supports the call waiting function).

trueCall. We need a way to implement the CIV protocol on an analogue phone, however, we cannot modify its firmware. To address this issue, we use trueCall, which is a commercial call-blocking device, designed to protect

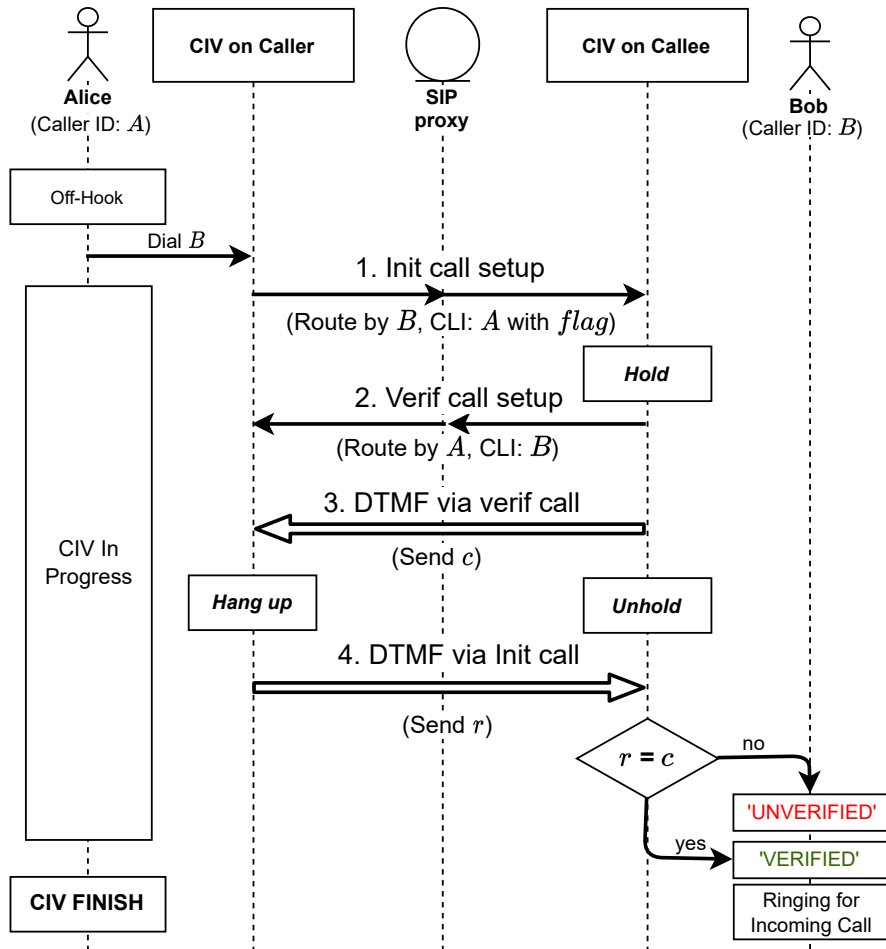


Figure 3.8: Using DTMF/DTMF to send challenge/response

elderly and vulnerable people from nuisance and scam calls. The hardware box contains a micro-controller that performs various call-control functions (e.g., off-hook, hang-up, ringing). We are able to modify the software in the trueCall box to implement CIV, but the hardware has its limitations – specifically, it does not support the call-waiting function. (Supporting the call waiting function in trueCall is possible but it needs an extra chip in the hardware.)

In order to implement CIV on the trueCall hardware three call setups are needed. After Bob's CIV (implemented in trueCall) receives the initial call (Step 1), it cannot hold the call. Instead, it terminates the call and then starts a verification call (Step 2) to send the DTMF challenge. Alice's CIV will need to start a new call (Step 3) to send the DTMF response. Overall, three call setups are required. Since the challenge and response process is handled by

CIV, it is transparent to users except that the caller will experience a longer wait for the extra call setup (details in Section 5).

Android. We use the Android third-party phone API (6.0.1) to build a CIV-enabled phone app for Android phones (Nexus 5). There are several implementation challenges that we need to overcome. First, the third-party API on Android only allows us to control one call at a time, but in CIV, we need to handle the verification call in parallel to the initial call. To overcome this limitation, we use Java reflection to invoke the hidden system service `TELEPHONY_SERVICE` to access the internal interface `ITelephony` in run-time. This allows us to hold the initial call using the system API while performing the verification call using the third-party API. Second, Android only allows a user to send DTMF manually by pressing keys during a call but does not support doing this programmatically. We sidestep this limitation by appending the DTMF digits as an extension after the phone number separated by a comma: e.g., ‘5555555555,1234’ where ‘,’ indicates a short pause (2 seconds on an Android phone). During dialling, Android first calls ‘5555555555’, waits for 2 seconds, and sends ‘1234’ through DTMF automatically. This method is commonly supported on mobile phones, allowing a user to directly reach an extension number behind a PBX without having to talk to an operator. We use the same method to send the challenge and response through DTMF, but this means that we need to specify the DTMF values upon dialling rather than during the call. As a result, we will need three call setups in the implementation. Finally, the third-party API does not support the automatic recognition of DTMF tones. To overcome this limitation, we set `AudioManager` to the `MODE_IN_CALL` mode, which allows the DTMF tones to be played via the speaker. Once CIV receives the DTMF tones through the microphone, it converts them into digits by analysing the frequencies based on the Fast Fourier Transform (FFT). This allows us to build a proof-of-concept prototype of CIV on Android phones, but the speed of recognising DTMF is significantly limited.

For CIV between two Android phones, we cannot send the response using DTMF via the initial call as previously done for SIP phones (see Figure 3.8). Instead, we use a new call to send the response (by appending the 4 digits to the phone number as an extension during dialling) through DTMF. This means that we need three call setups as opposed to the 2 call setups. However, for a SIP phone calling an Android phone, the implementation is not affected by the Android’s in-call DTMF limitation, and it can still be done in 2 call setups. We will present a detailed performance evaluation in Section 3.6.

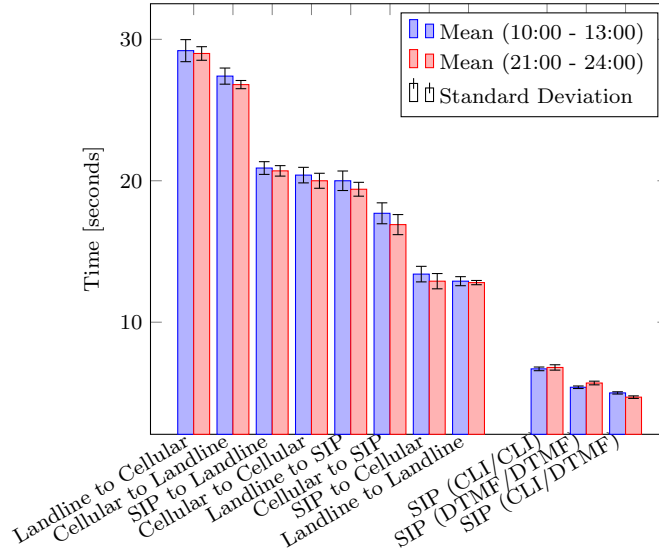


Figure 3.9: Delays in CIV between different phones

3.6 Evaluations

We evaluate the overhead of running CIV for calls between landline, cellular and VoIP phones. For each call scenario, we measure the delays of the CIV protocol during the day (10:00-13:00) and at night (21:00-24:00). We take 15 measurements at each time slot and present the average in Figure 3.9. The breakdowns of the measurements at night are presented in Figure 3.10 (the breakdowns during the day are basically the same). The breakdowns comprise four components: 1) setting up a verification call; 2) transmitting the challenge; 3) setting up a response call (if needed); and 4) transmitting the response.

3.6.1 SIP platforms

First, we evaluate the performance of CIV prototypes based on SIP platforms (see Case 1 and 2 in Figure 3.6). As shown in Figure 3.9, SIP ‘CLI/DTMF’ incurs the lowest 4.7 sec latency, followed by 5.7 sec in SIP ‘DTMF/DTMF’ and 6.8 sec in SIP ‘CLI/CLI’ during 21:00-24:00. The latency measurements during 10:00-13:00 are approximately the same. ‘CLI/DTMF’ requires one call setup to send the challenge through the (modified) CLI, and uses the existing initial call channel to send the response through DTMF. It is quicker than ‘DTMF/DTMF’ as it does not have the cost of transmitting the DTMF challenge (since the challenge is embedded in CLI; see Figure 3.10 (c)). Among

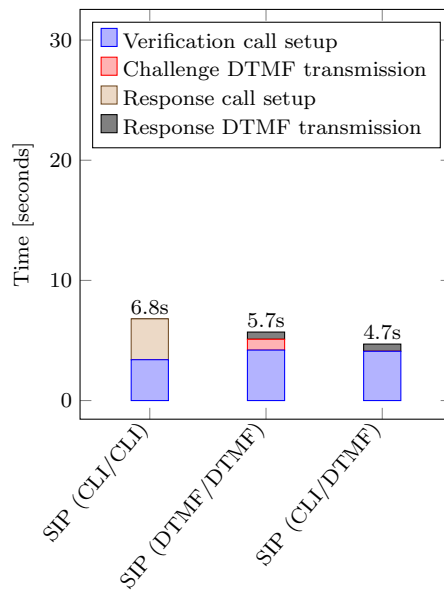



Figure 3.10: Cost breakdowns in CIV

these prototypes, ‘CLI/CLI’ incurs the longest delay because sending the response through CLI requires a call setup, which involves a significant cost. The other two prototypes use an *existing* initial call to send the DTMF response and hence are free from this call setup cost.

3.6.2 Other platforms

Prototypes built on trueCall and Android are limited by various platform constraints (see Case 3 in Figure 3.6). These constraints do not prevent the proof-of-concept implementation of CIV, but they increase the latency.

Android. As shown in Figure 3.9, when a cellular phone calls a cellular phone, the total latency of CIV is about 20 seconds. This delay is due to two main factors. First, the third-party API that we use to build a CIV-enabled phone app does not support sending DTMF during a call. We overcome this by appending DTMF digits as an extension of a phone number upon dialling, but this requires two call setups (for sending the challenge and response respectively). Second, as the third-party API does not support automatically recognising DTMF, we had to overcome this by playing out the audio sound of the DTMF tones via a speaker and decoding them into digits. It is worth noting that when a SIP phone calls a cellular phone, our proof-of-concept implementation is not affected by the first factor; the caller (SIP) is able to send back the DTMF response through the initial call because the SIP platform supports sending DTMF during the call. This removes the cost of “response call setup” (see Figure 3.10 (a) for a breakdown), but the overall delay of ‘SIP to Cellular’ is still dominated by the slow recognition of DTMF tones.

trueCall. As shown in Figure 3.9, when a landline phone calls a landline phone, the total delay of CIV is about 13 sec. The trueCall supports transmitting in-band DTMF and the automatic recognition of DTMF during a call. But the performance is limited by the lack of the *call waiting* function in the trueCall hardware. As a result, we need a full call setup to send the response, which is a significant cost component (see Figure 3.10 (b)). This affects all experiments that involve a landline phone with a trueCall box. Also, since the landline phone is connected to a PSTN network via an analogue line, transmitting DTMF between the landline phone and other networks incurs more delays. The worst case is when CIV is run between a landline phone and an Android phone; the total latency is 29 sec. This is because the hardware limitation of trueCall is compounded by the slow recognition of DTMF tones in our Android prototype. In Section 3.7, we will discuss how the performance can be substantially improved once the underlying platform constraints are removed.

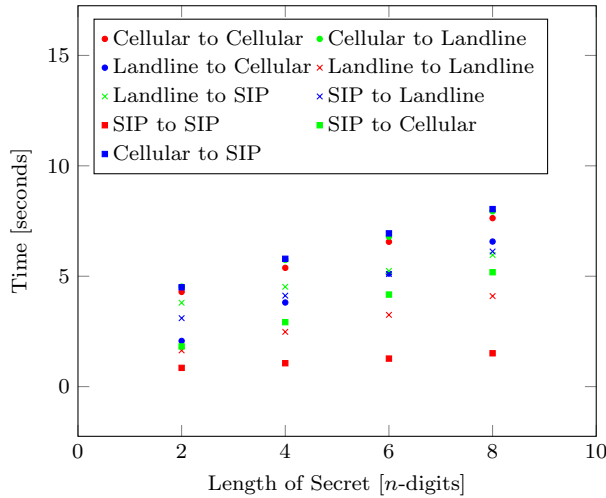


Figure 3.11: DTMF transmission time versus number of digits

3.6.3 Lengths of the challenge and response

In the design of CIV, we choose $n = 4$ for the number of digits in the challenge, as a reasonable trade-off between security and performance. Figure 3.11 shows the variation of the DTMF transmission time for different values of n . As expected, in all cases, the transmission time increases linearly with n . We would like to draw attention to the ‘SIP to SIP’ measurements since they involve only the transmission of digital DTMF values (not analogue tones). This is the trend as analogue phone lines are being phased out. As shown in the figure, the ‘SIP to SIP’ measurements have a relatively flat slope; increasing n hardly increases the delay. This shows the flexible scope of choosing a bigger n with little performance degradation.

3.7 Discussion

In this section, we discuss various aspects of the CIV system, including security, deployment, and limitations.

3.7.1 Downgrading attack

In the CIV design, the caller needs to indicate support for CIV in the initial call. This can be achieved by appending a special CIV flag in the caller name. When the caller name for a phone number is registered in a CNAM database, the CIV flag is also saved there. Hence, when a terminating network retrieves

the caller name from CNAM databases for an incoming call, it recognises that the caller supports CIV. Alternatively, the caller name may be sent in the call along with the caller ID.

This design supports callers who *actively* want to be verified, e.g., to increase the likelihood that the called party will answer the call. However, an attacker may launch a downgrading attack to bypass CIV: spoofing a caller ID and transmitting a false caller name without any CIV flag. If the caller ID/name is registered in CNAM databases, the terminating network can look up the databases, and use the retrieved caller name to overwrite the received caller name. However, CNAM databases are not used everywhere. Also, the terminating network might simply use the received caller name instead of looking up CNAM databases (the latter involves paying a ‘dip’ fee). Hence, it is possible that even though a caller ID is registered as CIV-enabled, an attacker may spoof that caller ID without invoking the CIV process.

But the above downgrading attack has limited effects. In the worst case, the called party’s phone rings with a display of a caller ID, but the CIV program shows a warning “caller not verified”. This is actually an intended outcome for the CIV design. A successful attack should involve spoofing a number that the attacker does not own and showing “caller verified” on the callee’s phone. This is unlikely as we explain below. We note that it is possible for a user to save a phone number as a local contact on the phone with a CIV flag. This ensures that the CIV program is always invoked when receiving a call with a display of that number.

3.7.2 Attacking the challenge-response protocol

The main idea of CIV lies in how we define “authentication”: based on if the caller can prove the possession of a phone number by answering a challenge sent to that number. In practice, when a user registers a phone number with a system, the system checks if the provided number is legitimate by calling the number or sending an SMS code. In our design, we automate the verification process by sending a random 4-digit code as part of the CLI (preferred) or through DTMF. In our threat model (Section 3.4.1), we assume that a spoofing attacker has no control over intercepting calls in telecommunication systems. Hence, if they do not own the phone number, they cannot receive the challenge. This leaves them with the only option of guessing the challenge and forging a response. With the 4-digit code in the challenge, the chance of forging a

valid response is $10^{-4} = 0.01\%$ (which can be further reduced by increasing the number of digits). This does not completely eliminate the theoretical attack, but it massively reduces the success probability, hence serving as a deterrent. We note that CIV *always* connects a call regardless of the verification result. Hence, repeated failures of guessing will alert the user.

3.7.3 Denial-of-service (DoS)

In CIV, a malicious caller may spoof an arbitrary caller ID when making a call. The goal is not to pass the CIV verification, but to leverage the callee to call back to the spoofed number, e.g., to cause a nuisance. We address this threat in two ways. First, CIV is invoked only when the caller indicates support for CIV in the initial call. A user can further configure CIV to perform verification only for certain numbers (e.g., domestic, non-premium numbers). Second, we propose using ‘CLI/DTMF’ as our recommended implementation of sending challenge/response (see Figure 3.7). Since the callee uses a modified CLI to transmit the 4-digit challenge in an *abandoned* call, there is no cost to the callee⁴. If the verification call reaches an unsuspecting phone user who has no CIV installed, it will be shown as a silent *missed call* with a non-dialable four-digit number. If the user’s phone has CIV installed, CIV can recognise unsolicited verification calls and automatically filter them out. We note that the attacker can always make silent *missed calls* to a target user directly, but the attack is deterred since the attacker can be traced by checking the call detail record (CDR) at the telephone exchanges.

The attacker also can leverage CIV to make a reflected DoS attack against a user. But the attack is deterred due to the two reasons. First, since the attacker can still be traced by correlating the two related CDRs, delivering a large scale reflect DoS attack without being detected is infeasible. Second, in the telephony network, the number of caller IDs binding to each phone line is a pre-setup in the exchange office, which is out of the control of individual attacker. Normal setup is that each phone line is assigned only one caller ID. Thus, tens or hundreds of phone lines are required to deploy an effective reflect DoS attack, making the attack costly and impractical.

Any such DoS attack actually has a limited effect: in the worst case, it leaves a silent *missed call* with a non-dialable number, but the call can be easily filtered by the user’s phone (say by installing a CIV program).

⁴Incurring no cost to the callee was highlighted as important during our meetings with telecom providers to gather design requirements.

3.7.4 Stages of deployment and optimisation

For calls between landline, cellular and VoIP phones, we have explored the most practical way to implement CIV for each call scenario within constraints; all this is done without cooperation from network providers. A unified and optimal implementation of CIV is possible by integrating CIV into the networks in a three-stage deployment.

Stage 1 (short term). This stage involves proof-of-concept demonstrations of CIV on the users' phones as we have done in this work. Without any cooperation from network providers, we show CIV can be implemented for all three different telecom networks (PSTN, cellular and VoIP) without modifying the existing infrastructure. This presents probably the best that can be done at the user's end under various platform constraints. Some of the constraints can be easily removed, e.g., by adding call-waiting and DTMF functions on some phones. Other constraints are more fundamental, e.g., access to the CLI modification function. Understanding these constraints lays the foundation for the next stage of CIV deployment.

Stage 2 (medium term). This stage involves integrating CIV into the *terminating network*. This removes one of the most important constraints in implementing the 'CLI/DTMF' method as described in Figure 3.7: namely, access to the facility of modifying CLI for sending the challenge. We note that the integration of CIV can be done *independently* by any terminating network without cooperation from other providers. When receiving a call with the CIV flag, the terminating network performs the verification process on behalf of its subscriber and relays the caller ID as well as the verification result to the subscriber's phone.

Stage 3 (long term). This stage is a natural evolution of Stage 2. A network provider can extend their service in Stage 2 to perform CIV on behalf of the caller as well. Again, this needs no cooperation from other providers. When more networks implement CIV and support both the caller and the callee, it may reach a point which networks commonly support CIV and there is no need to install CIV on the user's device anymore. CIV then essentially becomes a new service in the telecom cloud (as an enhanced version of CLI with verified caller ID), which users can subscribe to. The challenge-response process is done between the switches of the two communicating carriers within the telecom cloud.

Example. trueCall [13] has already been integrated into the UK telecom

network so it is available as a telephone service that users can subscribe to. Users do not need to install any trueCall box at their end, since the service is virtually available in the cloud. This is realised by adding a software “hook” in the cloud to handle calls on behalf of the subscribed trueCall users. We expect that CIV will follow a similar approach for the integration into telecom clouds.

3.7.5 Limitations

Our current work on CIV has several limitations. First, the verification delay for landline and cellular phones is high (12–29 seconds). This is mainly because we do not have access to the facility to modify CLI for the landline and cellular phones (only the switches in the network can modify CLI). Hence, we are not able to use the most efficient ‘CLI/DTMF’ method to implement the challenge-response process. Second, our SIP prototype based on ‘CLI/DTMF’ reports 4.7 seconds delay, which is closer to being practical, but there is still room for improvement. As shown in Figure 3.10 (c), this delay is dominated by the ‘verification call setup’, which involves not only routing the call but also, more importantly, allocating resources along the call path to prepare for the ensuing telephone conversation when the call is answered. However, in our case, the verification call is only to transmit a short challenge, not intending for a conversation. Hence, the call setup using INVITE takes longer than necessary. This delay can be substantially reduced by using a different signalling mechanism (e.g., out-of-band INFO or NOTIFY messages), but this needs to be done between SIP providers rather than from SIP phones. Third, the current implementation of CIV for SIP adds an extra flow of INVITE signalling for the verification call, which may add a burden on some networks and affect the termination success rate. As with the previous one, this can be addressed by using a different (out-of-band) signalling method between SIP providers to transmit the challenge. Finally, as a deterministic system, CIV has zero failure rate by design. In the reality, it is not. Due to the uncontrollable factors, like the volume of the network traffic or the interference from the platform, our test only got 35% successful rate during the busy hours on the platform of Android, while the number increased to 90% during the off-peak hours on the same platform. It should be possible to overcome all these limitations by integrating CIV into the telecom cloud (Stage 3 deployment), which we plan to investigate in further research.

3.8 Conclusion

We propose CIV, a new solution to authenticate caller ID in heterogeneous telephone networks without a public key infrastructure. CIV authenticates the caller ID through a challenge-response protocol; it distinguishes legitimate and illegitimate spoofing based on if the caller owns the phone number; it supports both SS7 and SIP; and it has been implemented on all three types of phone systems and tested across heterogeneous networks to demonstrate feasibility. Contrary to the common belief by the FCC and regulators in some other countries that STIR/SHAKEN is the only solution, our work shows that alternatives exist and that they can be far more cost-effective than STIR/SHAKEN. We hope this will encourage more research into bottom-up solutions to address caller ID spoofing without relying on any trusted third party or a PKI.

Chapter 4

Counterfeiting No More - Polymer Substrate Fingerprinting

In this chapter, we propose polymer substrate fingerprinting (PSF), a PUF-based authentication method to tackle the problem of banknote counterfeiting. We first review the existing solutions, as well as their limitations. Then we introduce the design of our proposal, followed by the evaluation data collected from extensive experiments. The evaluation shows that the entropy in the PSF is extremely high that the fingerprint extracted from the stochastic pattern in coating layer can identify every authentic polymer banknote circulated globally. Even the adversaries possess the dedicated equipment and ink as used by authorities, counterfeiting banknotes remains infeasible.

4.1 Introduction

Despite the increasing volume of transactions made by credit cards and electronic payment methods, banknotes still play a crucial role in our society. In many countries, such as the US, the UK, Canada, Australia, and the European Union, the demand for cash continues to grow with the value of banknotes in circulation increasing each year typically by a factor of 5 to 10 percent [43]. Globally, there are over 500 billion banknotes in circulation. According to a report by McKinsey & Company [25], over the past years, although the share of the world's transactions carried out in cash has fallen, banknotes remain one of the most widely used payment instruments in the world.

Counterfeiting, or the forgery of banknotes, has been a major threat to the society and economy. Since most banknotes cost little to produce, a successful forgery is virtually all profit. People who fall victim to this crime are essentially robbed. Their losses cannot be reimbursed as doing so will facilitate the circulation of counterfeits and encourage illegal activities. Widespread counterfeiting can severely undermine the value of the currency, and disrupt the economic development [142].

In general, anti-counterfeiting methods challenge the forger in two main aspects: the substrate, and the printing. Traditional banknotes use a paper substrate made of cotton and linen. Compared with the bond paper made of wooden particles, the cotton/linen paper is substantially more expensive and more durable. When used for banknotes, it also contains various security features which are introduced during the manufacturing process, such as watermark, embossed metallic thread and other unique features. The printing is another aspect that gives banks an edge against counterfeiting. It requires specialised equipment and ink which are prohibitively expensive for counterfeiters. One of the most important printing techniques is the so-called intaglio (gravure) printing, which gives the raised print and the unique texture feel of a banknote [142].

The latest development in banknotes is to print them on polymer: a thin, flexible plastic [142]. The new polymer substrate not only supports traditional security printing as employed for paper notes, but also allows enhanced security features, such as see-through window and foil patch. This makes them harder to counterfeit than paper notes. Since the first introduction in Australia in 1988, they have become the trend for printed currency and have been adopted by more than fifty countries. In the UK, Bank of England first issued polymer £5 and £10 in 2016 and 2017, respectively. It has started replacing £20 with polymer notes since 2020.

The introduction of polymer banknotes has evidently reduced counterfeiting. For example, after Australia fully replaced paper banknotes with polymer series in 1996, the rate of counterfeits fell noticeably from 16 ppm (parts per million - the number of counterfeits per million genuine banknotes in circulation) in 1996 to only 3 ppm in 2000 [24].

However, counterfeiters have been catching up. After 2000, the counterfeiting rate in Australia gradually increased, and reached above 25 ppm in 2015 [24]. As the quantity of counterfeits increases, so does the quality. While the first recorded counterfeits on polymer were detected in 1997, they were

printed on a paper substrate and used techniques only to simulate the feel of polymer. Around 2010, polymer counterfeits began to appear by using advanced technologies that enabled counterfeiters to print large volumes of counterfeits on a plastic film. This shows that the initial advantage of bringing a new polymer technology to fight against counterfeiting is reducing.

Although polymer banknotes have many existing anti-counterfeiting features, one fundamental limitation for the security assurance of those features is that they critically rely on the difficulty for counterfeiters to obtain the same or equivalent printing equipment and ink. As shown by the example of [70], professional counterfeiters often exploit weaknesses in the supply chain for the manufacturing of banknotes and obtain from worldwide suppliers essentially the same or equivalent printing equipment and ink as used for printing genuine notes. Their chance of success can be significantly boosted when the operation is backed by a state government. For example, many high-quality counterfeits of the US\$100 bill, known as “superdollars”, are allegedly made by countries that are antagonistic toward the USA. Some of the counterfeits are of such high quality that, according to Europol, they “are just U.S. dollars not made by the U.S. government” [172]. In face of such professional counterfeiters backed by a state government, existing security features of a banknote can be easily bypassed.

To maintain one step ahead of forgers, we propose a new anti-counterfeiting technique called Polymer Substrate Fingerprinting (PSF). In contrast to existing banknote security features which require delicate design and printing, our technique exploits the stochastic nature of the polymer substrate manufacturing process. It works by analysing the random translucent patterns of the polymer substrate when it is back-lit. These patterns are caused by stochastic printing and the randomly dispersed impurities in the ink during the opacity coating procedure. They naturally occur during the banknote production, and cannot be precisely controlled or duplicated. We show these patterns can be reliably captured by a commodity film scanner and processed into a compact fingerprint to uniquely and reliably identify each banknote.

4.2 Related Work

A number of researchers have proposed to analyse the physical properties of a banknote. Vila et al. [166] were among the first to propose analysing the infrared spectrum of a banknote to determine if it is genuine or not. They

proposed to examine selected areas of the banknote by using an infrared spectrometer, together with an attenuated total reflectance (ATR) microscope. Their dataset consisted of 18 randomly selected genuine notes of € 50 and € 100 denominations, and 5 counterfeit notes of € 50 and € 100 denominations, provided by the Spanish Police. Although their experiments showed distinguishing features in the infrared spectra between the genuine and counterfeit notes, this result critically relied on the specific counterfeit samples used in the study. Sonnex et al. [144] proposed a similar method based on infrared spectroscopy. Their dataset contained 27 counterfeit £20 notes from the Northamptonshire Police. Their study revealed a lack of contrast in infrared spectra between ink and paper among the forgeries. Hence, the authors proposed to use a simple and portable infrared device to search for spectral difference as the first line of defence, and in case of ambiguity, use a more expensive infrared microscope to map selected areas of printing in contrast to the background paper. Their study has the same limitation as [166] in that the result was only applicable to the specific counterfeit samples used in the experiment.

Some researchers proposed to analyse the ink composition to distinguish legitimate banknotes from counterfeits. Rusanov et al. [134] applied Mössbauer spectroscopy to analyse the chemical composition of the ink used in both genuine and counterfeit banknotes. They examined 54 authentic \$100 US banknotes chosen at random, and 13 forged notes which were provided by a bank in Bulgaria. The authors suggested that the absence of certain elements in the pigment (e.g., green dye sextet) could be used to distinguish counterfeit banknotes from authentic ones. Jara et al. [92] conducted a similar study to analyse the chemical composition of the ink in real and fake banknotes, by using an X-ray fluorescence spectrometer instead. Almeida et al. [52] proposed to apply Raman spectroscopy and chemometric tools to analyse the characterisation of the ink in a banknote. They examined 60 counterfeit banknotes provided by the Brazilian police, and a further set of 28 lab-made fake samples prepared by scanning authentic bills and printing copies on laser and ink-jet printers. Based on the difference in the Raman spectra of the chalcographic ink, the authors proposed to use a Partial Least Square for Discriminant Analysis (PLS-DA) classifier to first distinguish the counterfeits from the originals, and in case of detecting a counterfeit, use a second classifier to identify which type of the printer was used in making the counterfeits. The performance of their solution critically depends on the counterfeit samples used in training the classifier.

Some researchers proposed to use an imaging device to capture the visual difference between genuine notes and counterfeits. Yeh et al. [174] proposed to analyse the luminance histograms of the captured image of a banknote and apply multiple-kernel support vector machines (SVM) to distinguish the counterfeits from the genuine notes. The authors used a dataset of 70 genuine Taiwanese banknotes and 29 counterfeits. Berenguel et al. [27] proposed a similar technique to detect counterfeits by analysing the background texture printing. A surface picture of a given banknote is taken by using a flatbed scanner and converted to grey-scale. Histogram features of the grey-scale image are extracted as input to a linear SVM classifier to determine if the note is real or not. The authors used a lab-made dataset of forgeries by scanning genuine euro bills and then printed counterfeits with an HP LaserJet printer. In a follow-up paper [39], Berenguel et al. proposed a different classification method, but still used the same procedure to generate lab-made counterfeit samples for evaluation. All of these papers have a common limitation that the results are only valid for the specific counterfeit samples used in the study.

Anti-counterfeiting of banknotes is closely related to anti-counterfeiting of documents, since paper documents such as certificates, cheques and contracts face the same counterfeiting problem as banknotes. Clarkson et al. [42] proposed a method to authenticate a paper document based on the unevenness of its surface. Based on the observation that the fuzz-mat surface of a paper document has a unique 3-D texture structure, they proposed to use a flatbed scanner to scan the target document multiple times at 4 different orientations. Based on the measurements, they created a 3-D image of the paper surface texture, and split the image into small patches for feature extraction. This process created a feature vector of 3200 bits as a paper fingerprint. Experiments showed that their method was able to distinguish a genuine document from a forged one. However, one drawback of their method is that it requires repeated scans, and is time-consuming. In our system, we extract features from a 2D image, and require only one scan (or one snapshot using a camera). Sharma et al. [139] proposed a similar technique based on analysing the speckle patterns when light reflects on the paper surface. This follows an earlier work by Buchanan et al. [36], which used a laser to capture the speckle patterns. Recently, Toreini et al. [151] proposed a new fingerprinting technique, which captures the unique features of a paper document using transmissive light instead of reflective light. They showed that using the transmissive light was able to capture richer features in the textural patterns than using the reflective light, and hence

achieve better performance than previous works [36, 42, 139].

Our polymer substrate fingerprinting technique is inspired by the previous research in anti-counterfeiting of banknote and paper, but it is different in a few ways. First of all, we do not require any dataset of forgeries for training. Instead of merely classifying a banknote into a binary result of “real” or “forgery” like in [27, 39, 52, 92, 134, 144, 166, 174], our technique extracts a unique fingerprint from a physical banknote. The authentication of a banknote starts with a null hypothesis that it is a “forgery” until this hypothesis is compellingly rejected by statistics. The 900-bit entropy in the extracted fingerprints is higher than previous works [151], and lays a solid foundation for building a large-scale authentication system for both online and offline applications. Second, we are the first to propose utilising the imperfections in the opacity coating layer of a polymer banknote to build an anti-counterfeiting system. Besides the theoretical design, we have developed a complete data acquisition and processing method, built a concrete proof-of-concept prototype, collected an extensive dataset and conducted experiments with both empirical and theoretical analysis to demonstrate the feasibility of our proposed solution.

Compared to existing counterfeit detection methods that heavily rely on the security features of banknotes, we shift the problem from proving the real/fake banknote to solving the biometric/PUF problem, which identify the uniqueness of individual banknote. With our method, even using the same equipment to repeat the same producing process, the produced banknotes are still distinguishable. Thus, ours has a distinctive advantage that even if the attacker has acquired the same printing equipment and ink as used for printing genuine banknotes, and counterfeiting remains hard.

4.3 Production of Polymer Banknote

In this section, we briefly explain the production process of polymer banknotes. The stochastic nature of this process, especially during opacity coating, forms the basis for our proposed anti-counterfeiting solution.

While the world’s first banknote printed on clear plastic film was issued in Australia in 1988, this was the result of nearly twenty years of research and development. The major breakthrough in the field was the invention of a special type of plastic called biaxially-oriented polypropylene (BOPP), which after being covered with opacity coating allows quality printing of all of the security features that are printed on traditional paper notes [127]. The use of

BOPP makes the polymer banknote highly durable, as well as being waterproof and dirt-resistant.

A polymer note starts as clear plastic sheets, which are melted down at a high temperature (around 166 °C) and then blown into a large bubble of several storeys high. The walls of the bubble are pressed together and cooled to form a laminated polymer film. A layer of opacity coating will be added to allow printing security features on the polymer film.

The opacity coating process applies white ink to the film to make it opaque, except for areas that are left clear as see-through windows. The see-through window is a security feature applied on the polymer note as it forces a forger to use clear plastic film as the substrate, which requires more advanced printing equipment than a paper substrate.

The technique used for opacity coating is called *gravure* printing. Figure 4.1 shows an overview of the process. The substrate is pressed against the inked cylinder on a rotary press between a backing roller and a gravure roller. The cylinder is etched with small cells on the edge which hold the ink fetched from a liquid pool. When the cylinder is partially immersed in the liquid pool, it picks up ink to fill its recessed cells on each rotation of the press. A flexible blade (also known as the “doctor blade”) is used to remove any excess ink from the printing cylinder, leaving ink only in the cells.

At a microscopic view, the opaque ink layer after the gravure printing process is highly non-uniform, showing random variations in the thickness, as shown in Figure 4.1(b). This is due to two main reasons. The first is related to air bubbles. When the ink in the cell is transferred to the substrate under the pressed contact, air menisci penetrate the gap and become air bubbles trapped in the ink [175]. Due to the air bubbles, the ink transferring process is only partially performed. The second reason is related to the solid residues. After the ink is transferred to the substrate, the remaining liquid in the cell evaporates, leaving a solid substance. The substance adhering to the bottom of the cell reduces the volume of the container. As a result of a combined effect of air bubbles and solid residues, the opaque ink layer is highly uneven. The uneven coating layer causes the polymer substrate to exhibit random translucent patterns when it is back-lit by a light source, which we will demonstrate later. The existence of impurities in the ink adds further randomness to these patterns. All these are the imperfections from the opacity coating process, and they constitute the physical basis for the anti-counterfeiting technique that we propose in this paper.




Figure 4.1: Schematic of the Gravure printing process

After the white ink coating, the polymer substrate is ready for the subsequent printing of security features. Our technique does not rely on any of the printed security features, however we describe the process here for completeness. Security printing involves several layers of printing applied in sequence. The first is *offset litho*, which uses an offset roller to transfer ink to the polymer substrate and puts the basic pattern of the banknote in place. This is followed by *intaglio printing*, which is used to put the major design elements such as the portrait and narrative elements (e.g., Her Majesty the Queen on a £10 note). The next is *letterpress*, which prints letter and digits including the unique serial number. The subsequent stage is to print special line patterns on a polymer substrate to form diffraction gratings, which typically consist of 12,000 lines per centimetre coated with a thin film of a reflecting metal (e.g., aluminium). Light is diffracted from the lines to give changing colours when viewed from different angles. Next, a protective over-coating ink (clear varnish) is applied on both sides of the note to protect the printed design from dirt and solvent. The tactile features are then applied to assist the visually impaired to identify different denominations. Finally, the printed sheets are guillotined into

individual banknotes. Each banknote is then electronically inspected to ensure their quality fulfils the required standard. More details about the polymer note production can be found in [127].


4.4 Proposed Solution

4.4.1 Feature Area


First of all, we need to identify an area on the polymer banknote for feature extraction. Based on the observation that the opacity coating is a stochastic process, the ideal areas for feature extraction are those that are directly exposed from the opacity coating and not obstructed by the subsequent security printing. Therefore, for £10 notes, we choose an area between the “Ten” hologram and the see-through window as shown in Figure 4.2 (a). To locate the area precisely, we use two auxiliary markers: the pound sign in the see-through window and the silver foil patch contained in the hologram. Both are metallic images made by diffraction grating printing at extremely high precision (around 12,000 lines of thin metal film coated per centimetre). These images are darker than the surroundings. Hence, they can be easily separated from the background. Based on the detected markers, the feature area is automatically located with the same position and dimension. Figure 4.2 (b) displays the snapshots of the same feature area from three different polymer £10 notes when they are back-lit by a light source. These pictures exhibit random translucent patterns, which we will process later. Similarly, we identify and locate the feature areas on a polymer £5 note and a paper £20 note as shown in Figure 4.2 (c) and (d), respectively. Here we choose the paper £20 note as an example for comparison. When back-lit, a paper banknote also shows translucent patterns, but they are caused by the random interleaving of the cotton linens rather than the uneven coating as seen in a polymer note. Although we focus on the anti-counterfeiting for polymer notes, our technique can also be applied to prevent forgery of paper notes. The performance of our fingerprinting technique for these two different substrates will be compared in the evaluation section.

4.4.2 Experiment Setup


To capture the random translucent patterns of the polymer substrate when it is back-lit, we choose an off-the-shelf negative film scanner (Epson V850), as shown in Figure 4.3. The resolution of the scanner is set to 3200 dpi to




(a) Feature area on £10 note ($19.3\text{ mm} \times 5\text{ mm}$)



(b) Feature areas on different £10 notes



(c) Feature area on £5 note ($3.3\text{ mm} \times 13\text{ mm}$)



(d) Feature area on £20 note ($16.2\text{ mm} \times 16.2\text{ mm}$)

Figure 4.2: Feature extraction on different banknotes (the zoomed-in pictures are cropped as a square from the original images for demonstration)

obtain high-resolution images with the help of an embedded back-light. In our experiments, we use a film-frame to hold the banknote. The frame helps to keep the banknote flat and in position during the scanning process.

The primary reason for using a negative-film scanner instead of a more common flatbed scanner is that the former is specifically designed to scan




Figure 4.3: Scanner setup

a film by shining light through it, while using light sensors to capture the image on the other side. This fits precisely our purpose. On the contrary, a flatbed scanner scans an object using reflective light. In the UK and other countries, it is prohibited to scan a banknote in this way as it may allow a casual counterfeiter to produce a fake copy. In fact, the firmware of a flatbed scanner has an embedded function to search for anti-copy patterns, e.g., EURion constellation [115] printed on banknotes. Once the scanner finds such patterns, it will stop the scanning process. By contrast, with the film scanner, when the light shines through the £10 banknote, the EURion pattern has been blended into the background. As a result, the obtained image is extremely “noisy”, and totally unsuitable for counterfeiting. On the other hand, the “noise” or the randomness in the image is exactly what we need for building an *anti-counterfeiting* system.

For the purpose of comparison and evaluation, we also build a second prototype using an off-the-shelf camera (Panasonic DMC-FZ72) and a light-box, as shown in Figure 4.4. A piece of glass is covered on top of the banknote to keep it flat. A light-box brightens up the banknote from the underneath so that the camera can photograph the translucent patterns on the top at a close distance (about 2 cm). The aperture of the camera is fixed at 5.0 and the shutter speed at 1/100. This combination provides sufficient depth of field as well as stability to get a clear and sharp image. The shooting mode is set to “Macro” to capture the details of random patterns in a close-up.




Figure 4.4: Camera setup

4.4.3 Image Processing

After we photograph a back-lit polymer banknote, the image is cropped to contain only the feature area, which is located by the aide of auxiliary markers. The cropped image is further processed by applying 2-D Gabor filters into a compact 2048-bit binary code, which we call a *polymer substrate fingerprint*. Details of this process are explained below.

Gabor Filter Selection

Image analysis with Gabor filters are similar to perception as human visual system. As different patterns under microscope (shown in Figure 4.2) can be seen by naked eye, Gabor filters are more adaptive and suitable for PSF pattern analysis, rather than the low/high-pass or the band-pass filters. Two-dimensional Gabor filters are a common technique used to analyse the textural patterns of an image. They have been commonly employed in biometric applications such as iris and face [30]. A 2-D Gabor filter comprises a sinusoidal wave modulated by a Gaussian envelope. It efficiently detects the edges and textural patterns existing in a 2-D image by capturing features in both frequency and spatial domains. This allows the output of the 2-D Gabor filter to be used in distinguishing whether the two snapshots are originally from the same pattern. In our work, we only need it to work in the spatial domain. In this domain, a 2-D Gabor filter is described as below [30]:

$$\psi(x, y) = \frac{F^2}{\pi\gamma\eta} e^{-F^2[(x'/\gamma)^2 + (y'/\eta)^2]} e^{i2\pi Fx'}$$

$$\text{with :} \tag{4.1}$$

$$x' = x \cos(\theta) + y \sin(\theta)$$

$$y' = -x \sin(\theta) + y \cos(\theta),$$

where F is the central frequency of the sinusoidal wave, θ is the angle between the direction of the wave and the x axis of the spatial domain, e is the natural exponential function, γ and η are the standard deviations of the Gaussian envelope in the direction of the wave and orthogonal to it, respectively. The parameters γ and η represent the shape factors of the Gaussian surface, and are also called the *smoothing* parameters. They determine the selectivity of the filter in the spatial domain.

Different combinations of the Gabor filter parameters are capable to extract different textural features. However, there is no unified way to determine values for these parameters [30], as they depend on particular characteristics of the textural patterns to be extracted [151]. To efficiently select the combination, a matrix called a Gabor filter-bank is created that contains a range of frequencies and orientations of Gabor filters. Each individual frequency in the matrix is called a *scale*, which is calculated from a maximum frequency, known as f_{\max} . For a total of U frequencies, each scale is defined as follows:

$$\text{scale} = \frac{f_{\max}}{\sqrt{2}^{u-1}}, \quad \forall u \in \{1, 2, \dots, U\}. \tag{4.2}$$

For a total number of V orientations, each orientation is calculated as follows:

$$\text{orientation} = \frac{v-1}{V}\pi, \quad \forall v \in \{1, 2, \dots, V\}. \tag{4.3}$$

Suitable parameters for the Gabor filters can be determined by using an iterated process through experiments [30]. Once a suitable set of parameters is found, it can be used for the same type of textural patterns (e.g., using the same set of parameters for processing all human irises in iris recognition).

When choosing the values for the Gabor filter parameters, we have three main considerations. First of all, we consider the *decidability* [51], which measures how far the clustering of samples from the same source is statistically separated from the clustering of samples from different sources. Clearly, the decidability should be sufficiently large. Second, we consider the fractional Hamming distance (HD), which represents the percentage of bits that are different on corresponding bit positions between two binary strings. In the rest

of the paper, we will use HD as a shorthand to refer to fractional Hamming distance. The HD between samples from different polymer notes should ideally centre around 0.5. As we will show in the evaluation, centring around 0.5 will greatly simplify our analysis as the obtained binary fingerprint can be modelled as a series of Bernoulli trials. Third, after the image processing, the obtained polymer fingerprint should contain sufficiently high entropy. A high entropy (say more than 128 bits) will statically guarantee that the chance for a random polymer substrate to successfully pass the verification is negligible. In fact, as we will demonstrate, we are able to achieve much higher entropy (900 bits) in the extracted fingerprints. Based on these requirements and the selection method outlined in [30], we conduct empirical experiments based on 100 samples from a set of randomly chosen £10 banknotes and determine that a suitable set of parameters for extracting the random translucent patterns for a polymer substrate is $f_{\max} = 0.25$, $\gamma = \sqrt{2}$, $\eta = \sqrt{2}$. The values for the scale and orientation that give the best overall performance are $u = 5$ and $U = 6$ for computing the scale (Equation 4.2) and $v = 11$ and $V = 30$ for computing the orientation (Equation 4.3). The procedure of manufacturing the polymer banknotes with different denominations in the UK is essential the same process with the same materials. Thus, the same setting can be used for both £5 and £10 as the textural patterns are of the same kind. For paper notes, we use a different combination: $u = 6$ and $U = 6$ for the scale, and $v = 22$ and $V = 25$ for the orientation. The parameters are slightly different because of the different textural patterns exhibited by a paper note (see Figure 4.2 (d)).

The size of the Gabor filter applied on the scanned polymer banknote is 101×101 (unit: pixel). Given the resolution of the scanner being 3200 dpi, each pixel in the scanned sample corresponds to about $7.94 \mu\text{m}$ ($1/3200$ inch). For 101 pixels, that corresponds to a physical size of $101 \times 7.94 = 802 \mu\text{m}$ on the banknote. According to Equation 4.2, the frequency of the Gaussian envelope applied on the polymer banknote is a quarter of f_{\max} . Therefore, the wavelength is 16 pixels, equating to $127 \mu\text{m}$ as shown in Figure 4.5.

Feature Extraction and Comparison

With the 2-D Gabor filter defined above, we apply it to process the translucent patterns photographed from the feature area of a polymer note into a binary string of 2048 bits, similar to how an iris-code is generated from the textural patterns of an iris image in iris recognition [51]. First of all, a captured




Figure 4.5: Physical dimension of the Gabor filter

photograph is grey-scaled, and a 2-D Gabor filter kernel is applied on the converted image to obtain a matrix of complex numbers. Each pixel in the image is transformed into a complex number. Given an input image $I(x, y)$ of dimensions $X \times Y$ and a bank of discrete Gabor filters $G_{mn}(x, y)$ with $m \in \{1, \dots, M\}$ and $n \in \{1, \dots, N\}$, the complex number matrix $C(x, y)$ is computed for each filter of the bank as follows:

$$C_{mn}(x, y) = \sum_{a=1}^X \sum_{b=1}^Y I(a, b) \overline{G_{mn}(x-a, y-b)}, \quad (4.4)$$

where $\bar{}$ denotes the complex conjugate.

Because the values of adjacent pixels are usually highly correlated, we perform a down-sampling process in order to remove the correlation. Values in every 20th rows and 20th columns are selected to form a new matrix sized 32×32 . All elements in the matrix are complex numbers with real and imaginary parts. Each element is then decoded into 2 bits depending on which quadrant does the complex number falls into. This gives a binary output of $32 \times 32 \times 2 = 2048$ bits, which we call a “polymer substrate fingerprint”.

The similarity between two polymer substrate fingerprints, denoted as f_1 and f_2 , is measured by computing a fractional Hamming distance based on an XOR \oplus operation, as below:

$$HD = \frac{|f_1 \oplus f_2|}{2048} \quad (4.5)$$

This is similar to how iris-codes are compared in iris recognition [51], however, in the case of the iris, there is a 2048-bit mask vector in addition to a 2048-bit iris-code. The purpose of the mask is to filter out unreliable bit positions caused by artefacts such as eyelids and eyelashes from the HD

Table 4.1: Summary of Datasets

Group	Condition	Equipment	Denom.	$N \times S^1$
Benchmark	Favourable	Scanner	£10	100 x 10
Robustness Test	Rotation			100 x 10
	Scribbling			100 x 10
	Soaking			20 x 5
	Folding			20 x 5
Variation Test	Equipment	Camera Set		100 x 10
	Denomination	Scanner	£5	100 x 10
	Substrate		£20	100 x 10

computation. In our system, we carefully select a feature area that is not obstructed or interfered by artefacts such as holograms and other printed security features. This removes the need for a mask. Hence, the stored data is only half the size of an iris-code. In the ideal case, the HD between any two fingerprints extracted from the same banknote should be close to 0, and the HD between fingerprints extracted from different banknotes should be close to 0.5. In the sections below, we will systematically evaluate the HD comparison results.

4.5 Datasets

We collect an extensive set of samples from the UK banknotes of different denominations, under different conditions. In total, we have collected 8 datasets containing 6,200 sample images, taken from 340 different banknotes, including 140 £10 notes, 100 £5 notes and 100 £20 notes. These datasets can be divided into three groups: benchmark, robustness test and variation test, as summarised in Table 4.1.

4.5.1 Benchmark

The dataset in the benchmark group is collected in a favourable condition. It consists of 100 different £10 polymer notes with 10 image samples for each note, making it a total of 1,000 samples. Each sample banknote is sandwiched between two pieces of thin clear glasses in the aligned frame during the scanning process. The use of the frame helps constrain the banknote in the correct orientation.

¹Number of banknotes(N) times number of samples for each banknote(S)

4.5.2 Robustness Test

Rotation. As part of the robustness test, we rotate the banknote and use the auxiliary markers to automatically re-orient the image before processing. This is done in Matlab. We collect 10 samples per each banknote from the same £10 banknotes used in the benchmark set, after rotating each note by a different angle varying from -10° to 10° . Testing rotation within this range is sufficient for our purpose as in practice errors of mismatch occur by only a small rotation angle. In the two dimensional space, given coordinates of two points (x_1, y_1) and (x_2, y_2) in a Cartesian coordinate system, the orientation angle α is calculated below:

$$\alpha = \tan^{-1} \frac{|x_1 - x_2|}{|y_1 - y_2|} \quad (4.6)$$

The angle of the rotation α for each sample image is computed based on the centres of the two auxiliary markers. Then the image is rotated accordingly.

Scribbling. Under the Currency and Banknotes Act 1928 in the UK, it is prohibited to scribble on the surface of banknotes as that may deface the notes. Therefore, we use hairs and fibres attached to the surface of each £10 banknote used in the benchmark dataset to mimic the same effect of scribbling when the banknote is photographed.

Soaking. Sometimes a banknote may drop into water by accident, or get wet (e.g, by rain) during the daily usage. Because every polymer banknote is protected by a over-coating layer (varnish) as part of the production process, a polymer note is water-resistant by design. Nonetheless, we use twenty randomly selected £10 to conduct a soaking test, with one sample for each note taken before the test and four samples taken after the test. These banknotes are soaked in water for 2 minutes and then dried naturally on a flat surface for 30 minutes before they are scanned and processed.

Folding. In daily life, banknotes are often folded before being put in a wallet. We conduct a test to study the effect of folding on our method. Initially, we take a set of randomly selected £10 notes, fold each note in half and store them in a daily used wallet for three days. Afterwards, the folded notes are flattened with the images of the feature area taken. Next, we fold each banknote twice along the long side to make it more compact for storage in the wallet. The double folded banknotes are put in the wallet for another three days before they are flattened and scanned. The folding dataset consists of 100 sample images taken from twenty £10 with one sample of the original

note, two samples after folding once, and another two samples after folding twice.

4.5.3 Variation Test

Alternative Equipment. Instead of a film scanner, we use a camera and a light-box to photograph the same 100 £10 notes used in the benchmark set with 10 images for each note. Film scanners and cameras are two different types of optical imaging devices, using different physical mechanisms. A film scanner obtains an image by moving a bar of light sensors alongside the surface of a flat film with a light shining on the opposite side of the film, while a camera flashes an array of light sensors in one go. Despite having a slow developing speed, the scanner tends to capture a high-quality edge-to-edge image. The reason is that it has a relatively simple optical structure with only one flat protective screen being laid on top of the sensor, while for a camera, light needs to pass through 4 to 7 lenses before reaching the sensors. The polymer fingerprints obtained from using these two different devices will be compared in the evaluation section.

Different Denominations. The £5 and £10 banknotes use essentially the same polymer substrate. Under the microscopic view, we observe similar random translucent patterns in the opacity coating layer for both £5 and £10 notes. To study the variation between these notes of different denominations, we use the film scanner to photograph 100 £5 polymer notes with 10 samples per note, and compare them against the benchmark set. We use the same Gabor filter setting for £5 as used for £10 in the benchmark dataset.

Different Substrates. To study of the variation between a polymer substrate and a paper substrate, we randomly choose 100 £20 paper notes. The paper £20 note in the UK uses a paper substrate made of cotton and linen. We use the same film scanner to image 100 £20 notes with 10 samples per banknote. As we will show in the evaluation, although our technique is designed for the anti-counterfeiting of polymer notes, it can be easily adapted to prevent forgery of paper notes as well.

4.6 Evaluation

4.6.1 Framework

Our polymer substrate fingerprinting technique is closely related to the technology of biometrics which authenticates people based on their inherent physical or behavioural features. Here, we authenticate a polymer banknote based on its inherent physical properties in the polymer substrate. On the other hand, our method is also related to the field of physically unclonable function (PUF), which provides security assurance based on the impossibility to physically clone a physical object. However, biometrics and PUF generally use different evaluation metrics despite that the two are inherently related. Based on earlier work [151], we propose to use a unified framework that combines both biometrics and PUFs metrics for evaluating our polymer substrate fingerprinting system.

Biometrics

A biometric system authenticates people based on their unique physical or behavioural features [51]. The performance of a biometric, especially one that uses HD for comparison, is commonly evaluated in terms of decidability, degree of freedom, and error rates as explained below.

Decidability. In a biometric system, there are two groups of biometric data distributions: the intra-group that refers to the distances between samples from the same subject and the inter-group that refers to the distances between samples from different subjects. In this paper, we use fractional Hamming distance (HD) as an example of the distance metric. Clearly the two distributions should be as further apart as possible. We use the *decidability* metric [51] to measures how far the two distributions are separated. This metric is denoted d' and is computed as below:

$$d' = \frac{|\mu_1 - \mu_2|}{\sqrt{\frac{\sigma_1^2 + \sigma_2^2}{2}}}, \quad (4.7)$$

where σ_1 and σ_2 are the standard deviations of distances between samples from the intra-group and the inter-group, respectively, μ_1 and μ_2 are the mean values from these two groups. $|\cdot|$ denotes the absolute value.

Degree of Freedom. The number of degrees of freedom (*DoF*) is a metric that measures how many independent bits exist in a biometric instance. In our systems, the more degrees of freedom contained in the extracted feature

Table 4.2: Notations used in PUF metrics. (In the benchmark dataset, $S = 100$, $L = 2048$ and $T = 10$)

f	Feature vector
S	Total number of banknotes
s	Index of each banknote ($1 \leq s \leq S$)
L	Bit Length of the feature vector from each banknote
l	Index of each bit position in a feature vector ($1 \leq l \leq L$)
T	Total number of samples measured per banknote
t	Index of each sample ($1 \leq t \leq T$)

vectors, the more statistically unlikely it will be for any two random feature vectors to match, thus the more entropy the vector contains (In this thesis, we use DoF and entropy interchangeably). The DoF is calculated below [51]:

$$N = \frac{\mu(1 - \mu)}{\sigma^2}, \quad (4.8)$$

where μ is the mean of the HD in the inter-group, and σ is the standard deviation of the HD in this group.

Error Rates. In a biometric verification system, there are two types of error rates: a false rejection rate (FRR) and a false acceptance rate (FAR). FRR refers to the probability that a genuine sample is falsely rejected, while FAR refers to the probability that a fake sample is falsely accepted. For practical purposes, both FRR and FAR should be kept as small as possible (ideally 0%). In reality, they vary according to the choice of a threshold. Increasing the threshold can reduce FRR but often at the expense of increasing FAR . Commonly an equal error rate (EER), where the curves of FRR and FAR intersect, is used to indicate the overall error rate performance of a biometric system.

Physical Unclonable Function

Physical Unclonable Function (PUF) is a security primitive built upon the difficulty of replicating the same physical properties of an object or device. Maiti et al. [107] proposed a framework to evaluate the performance of PUF. We adapt their framework as part of the metrics used to evaluate the performance of our system in the following three dimensions: space, time, and device. Notations used in this framework are summarised in Table 4.2.

Space Dimension - Uniformity, Randomness. In the space dimension, we assess how uniform the 0s and 1s are distributed in a feature vector and

how random the binary values are at each bit position of a feature vector.

$$Uniformity(s, t) = \frac{1}{L} \sum_{l=1}^L f_{s,t,l} \quad (4.9)$$

$$Randomness(s) = -\log_2 \max(p_s, 1 - p_s) \\ \text{where } p_s = \frac{1}{TL} \sum_{t=1}^T \sum_{l=1}^L f_{s,t,l} \quad (4.10)$$

Time Dimension – Reliability, Steadiness. In the time dimension, we assess the similarity of samples taken at different times from the same banknote. Reliability measures how consistent a feature vector from a banknote is as compared with other feature vectors taken in different times from the same banknote. Steadiness measures how stable the value at each bit position is among all feature vectors taken from the same banknotes.

$$Reliability(s) = 1 - \frac{2}{T(T-1)L} \sum_{t=1}^{T-1} \sum_{t'=t+1}^T \sum_{l=1}^L (f_{s,t,l} \oplus f_{s,t',l}) \quad (4.11)$$

$$Steadiness(s) = 1 + \frac{1}{L} \sum_{l=1}^L \log_2 \max(p_{s,l}, 1 - p_{s,l}) \\ \text{where } p_{s,l} = \frac{1}{T} \sum_{t=1}^T f_{s,t,l} \quad (4.12)$$

Device Dimension – Uniqueness, Bit-Aliasing. In the device dimension, we consider the diversity of the feature vectors taken from different banknotes. Uniqueness measures how distinguishable a feature vector is from other feature vectors extracted from different banknotes. Bit-aliasing measures how likely different banknotes are to produce identical values at the same bit positions in the feature vector.

$$Uniqueness(s) = \frac{2}{T^2 S(S-1)L} \cdot \sum_{t=1}^T \sum_{\substack{s'=1 \\ s' \neq s}}^S \sum_{t'=1}^T \sum_{l=1}^L (f_{s,t,l} \oplus f_{s',t',l}) \quad (4.13)$$

$$Bit-Aliasing(l) = \frac{1}{ST} \sum_{s=1}^S \sum_{t=1}^T f_{s,t,l} \quad (4.14)$$

4.6.2 Results

Benchmark Performance

Based on the benchmark dataset, we compute pair-wise HDs between the feature vectors obtained from the same banknotes (intra-group) and from different banknotes with the same denomination (inter-group). The histograms

for the two groups of HD calculations are plotted in Figure 4.6.

Biometric Metrics. From Figure 4.6, the inter-group and intra-group HD distributions are clearly separated. Based on Equation 4.7, we calculate the *decidability* $d' \approx 29$, which is much larger than the reported $d' = 14$ from iris codes [51]. The higher decidability means more separation between the two distributions, giving us more flexibility when choosing the threshold. One main reason for the higher decidability in our system is that we photograph the random features of a polymer substrate at an extremely close distance (1-2 cm), but this is not possible with the iris scanner as that would be too invasive to a human.

For the inter-group HD distributions, we obtain the mean HD $\mu = 0.500$ with a standard deviation $\sigma = 0.017$. Based on Equation 4.8, we are able to calculate the number of degrees of freedom $N = 900$. To confirm that N accurately reflects the number of degrees of freedom for the actual polymer fingerprints, we plot a binomial distribution curve which models a series of 900 Bernoulli trials (i.e., tossing an unbiased coin) with a probability of 0.5 for each trial. As shown in Figure 4.6, this binomial distribution curve fits perfectly the HD histogram in the inter-group. This corroborates the fact that the obtained 2048-bit fingerprints from the polymer banknotes have 900 degrees of freedom, or in other words 900 bits entropy. By comparison, the number of the degrees of freedom for a 2048-bit iris code is only 249 [51]. Note that the iris textural patterns tend to be correlated along the radial directions [51], which reduces the entropy of the iris codes, while such correlations do not exist in the polymer substrate. This, together with the fact that we can take a close-up of the polymer substrate at an extremely short distance, contributes to the much higher entropy in the extracted polymer substrate fingerprints than in iris-codes. The intra-group HD distributions do not show the same symmetric shape as the inter-group HD distributions as they heavily depend on the noise in the data acquisition. A few noisy samples can result in relatively high intra-group HDs for the same banknotes, leaving a long tail in the distribution.

From Figure 4.6, it is clear that the two groups of distributions are far apart. If we choose an HD value 0.33 as the threshold, the *FRR* and *FAR* will be both 0%. Obviously the *EER* for the overall performance is also fixed at an ideal value 0%.

PUF Metrics. In Section 4.6.1, we have defined a set of metrics to evaluate PUF. Table 4.3 summarises the performance of polymer substrate fingerprints using those metrics due to Maiti et al. [107] along with other

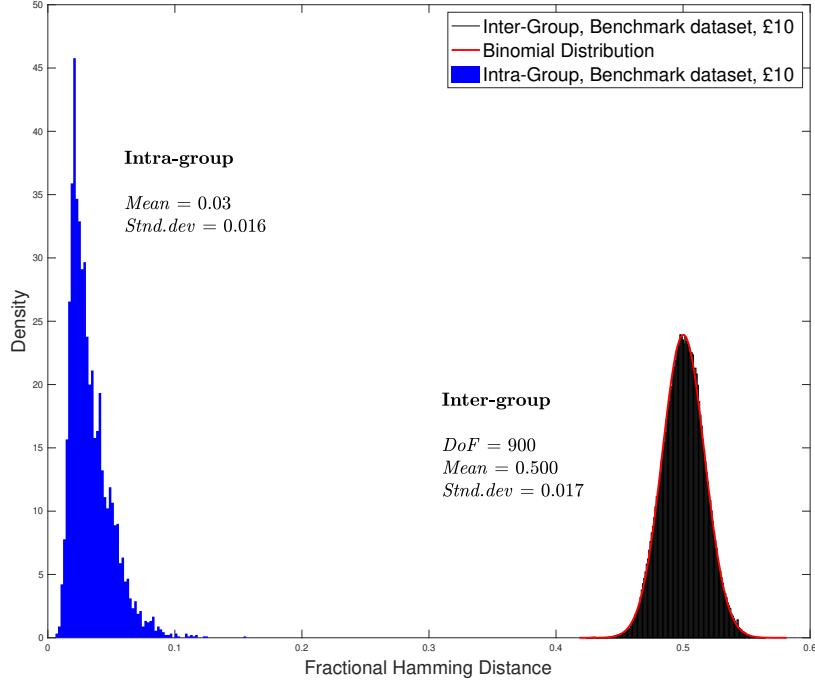


Figure 4.6: HD distributions of Benchmark dataset. Decidability $d' \approx 29$.

Table 4.3: PUF metrics from Benchmark dataset

PUF Metrics	Ideal Value	Bench. Dataset	Paper PUF [151]	Arbiter PUF [107]	Ring Oscillator PUF [107]
Uniformity	0.5	0.500	0.466	0.556	0.505
Randomness	1	0.980	0.907	0.846	0.968
Steadiness	1	0.962	0.945	0.984	0.985
Reliability	1	0.967	0.938	0.997	0.991
Uniqueness	0.5	0.500	0.465	0.072	0.472
bit-Aliasing	0.5	0.500	0.466	0.195	0.505

related PUFs proposed in the past work for comparison. As shown in Table 4.3, our technique achieves results close to the ideal values in each of these metrics. Overall the results also compare favourably in general to the state-of-the-art PUF systems reported in the literature [107, 151].

Robustness Tests

We plot HD histograms for different robustness test cases in Figure 4.7. We explain each case below.

Rotated Dataset. As shown in Figure 4.7a, *rotation* has little effect on




Figure 4.7: HD histograms after robustness tests

the performance as the software is able to automatically re-orient a banknote image based on auxiliary markers before the feature area is processed. Both the intra-group and inter-group distributions remain largely unchanged. As an example, if we choose $HD = 0.33$ as the threshold, the FRR and FAR still maintain at 0%.

Scribbled Dataset. As compared to the benchmark dataset, scribbling on the banknotes shifts the centre of the intra-group distribution to the right (from 0.03 to 0.06), but it has little effect on the inter-group distribution (shown in Figure 4.7b). This means scribbling on a banknote adds noise to the data, but the two groups remain clearly separated. At a threshold of $HD = 0.33$, the FRR and FAR are kept at 0%.

Soaked Dataset. As shown in Figure 4.7c, soaking a banknote has little effect on both the intra-group and inter-group distributions. This is as expected since the polymer banknotes are water-proof by design (due to the application of clear veneer at the outer layer). Given $HD = 0.33$ as the threshold, the FRR

Table 4.4: PUF metrics after robustness tests vs benchmark

PUF Metrics	Ideal Value	Rotated Dataset	Scribbled Dataset	Soaked Dataset	Folded Dataset	Bench. Dataset
Uniformity	0.5	0.499	0.499	0.498	0.501	0.500
Randomness	1	0.981	0.980	0.978	0.978	0.980
Steadiness	1	0.960	0.960	0.971	0.949	0.962
Reliability	1	0.965	0.965	0.972	0.950	0.967
Uniqueness	0.5	0.500	0.500	0.501	0.501	0.500
bit-Aliasing	0.5	0.499	0.499	0.498	0.501	0.500

Table 4.5: PUF metrics after variation tests vs benchmark

PUF Metrics	Ideal Value	Camera Set	Polymer £5	Paper £20	Benchmark Dataset
Uniformity	0.5	0.499	0.499	0.497	0.500
Randomness	1	0.984	0.985	0.983	0.980
Steadiness	1	0.884	0.935	0.978	0.962
Reliability	1	0.900	0.944	0.981	0.967
Uniqueness	0.5	0.500	0.500	0.500	0.500
bit-Aliasing	0.5	0.499	0.499	0.497	0.500

and FAR are still 0%.

Folded Dataset. Figure 4.7d shows that given $HD = 0.33$ as the threshold, the FRR and FAR of the folded dataset distributions are both 0%. The result indicates that folding a banknote in daily usage has limited effect on the separation of the intra-group and inter-group.

PUF Result. As shown in Table 4.4, the robustness tests in our experiments have little effect on the PUF metrics. All the values computed after the robustness tests remain close to the ideal values. This suggests that our technique is reasonably robust against non-ideal daily handling of banknotes. To a large extent, the strong robustness of our method is attributed to the basic design of a polymer note: in particular, the veneer coating at the outer layer protects the printing underneath and makes the polymer note highly durable against rough daily usage.

Variation Tests

In the section, we study the variation of performance under different test conditions, including the use of a different imaging device, a polymer note of a different denomination and a banknote of a different substrate. Fitted curves for the HD histograms under these different conditions are shown in Figure 4.8. Values of the PUF metrics calculated under these conditions are summarised in Table 4.5.

Alternative Equipment. When a camera and a light-box are used




Figure 4.8: Histograms of variation tests

instead of a film scanner to photograph the feature area, the steadiness and the reliability of the obtained feature vector slightly decrease as shown in Table 4.5. This is also reflected in Figure 4.8a, in which the intra-group distribution slightly shifts to the right. As a result, the decidability d' is reduced from 29 in the benchmark set to 22. This is because the camera used in the experiment has a more complex optical path for light reaching the *CMOS* sensors than the film scanner. Furthermore, a close-up taken by a camera in the macro mode under a close distance (about 2 cm) tends to be slightly bent near the edge of the ring [151], which adds noise to the feature extraction. Nonetheless, the distributions of the two groups are still clearly separated with a maximum HD of 0.195 for the intra-group, and a minimum HD of 0.425 for the inter-group. The FRR and FAR remain at 0% when the threshold is set to 0.33.

Different Denominations. As compared to £10, the intra-group HD distributions for £5 shifts to the right as shown in Figure 4.8b, while the inter-group distribution remains basically unchanged. This is mainly because

the £5 polymer note is physically smaller than the £10 polymer note, and the area suitable for feature extraction (i.e., directly exposed from the opacity coating layer without the obstruction of security printing) is also smaller. In our experiment, while the Gabor filter setting is the same, the feature area defined for £5 is only about half of the area for £10 (also see Figure 4.2). While the smaller area has little impact on the PUF metrics (see Table 4.5), it reduces the decidability d' from 29 to 18, and the DoF from 900 to 854.

Different Substrates. Since paper £20 notes are still used in the UK, we test our fingerprinting technique on £20 notes that use paper substrate. We obtain slightly better performance than the benchmark £10 polymer notes. The decidability d' is slightly increased from 29 to 32, while the DoF is increased from 900 to 1043. The FRR and FAR remain at 0% for the threshold of $HD = 0.33$ (as seen in Figure 4.8c).

PUF Result. The PUF metric values are basically the same as the benchmark dataset. The slightly better performance of £20 is related to its inherent textural patterns. As shown in Figure 4.2, a paper £20 banknote also exhibits random translucent patterns, which are caused by the random leaving of the cotton fibre and linen rather than the opacity coating, but the image seems to contain richer textural information than a polymer substrate. This shows that although our technique is designed for the anti-counterfeiting of polymer notes, it can also be adapted to prevent forgery of traditional paper notes.

4.6.3 Limitations

Our current work has a few limitations. First all, the data samples are taken from the UK banknotes only. Given that the manufacturing of polymer notes follows essentially the same process, we believe the results are applicable to banknotes in other countries, but this needs to be confirmed in further research. Second, we have done robustness tests under common cases, but the tests are not exhaustive. Further evaluation may include folding the banknote more than twice, placing the banknote under high temperature (near the melting point), and studying the effect of wearing out after years of usage. Secondly, the features are extracted from different areas on banknotes of different denominations. Hence, the system needs to identify the denomination first, which is doable but adds an extra step in the processing. Defining a standardised feature area for all polymer banknotes will be highly desirable.

Finally, as each banknote can be tracked by its fingerprint, our method weakens the privacy of the banknote, especially when PSF is verified online. Therefore, as the trade-off between security and privacy, we designed an offline application that reveals only the essential information during the authentication (as shown in the section 4.7).

4.7 Anti-counterfeiting Applications

4.7.1 False Acceptance Rate

In the application of counterfeiting detection, we care less about the problem of false rejection rate than of the false acceptance rate. If a genuine banknote fails to pass the initial verification, it can be subjected to further verification. However, in the case of false acceptance, a counterfeit banknote will be erroneously accepted as genuine.

The false rejection rate of a system heavily depends on the data acquisition environment during the verification. On the other hand, the false acceptance rate is essentially determined by the inherent entropy of the data source, and we can theoretically estimate the value as follows. Let P_a be the false acceptance rate of a fingerprint for one-to-one comparison. Based on the 900 degrees of freedom and the binomial distribution fitting in Figure 4.6, we model each fingerprint as the result of performing a series of $N = 900$ Bernoulli trials with the probability $p = 0.5$ of guessing ‘heads’ (or ‘tails’) correctly for each trial. Hence, we can compute $P_a = \sum_{i=0}^m N! / (m!(N - m)!) \cdot p^m \cdot (1 - p)^{N-m}$, where m is the number of successful guesses [51]. Given a threshold $\theta = 0.33$, $m \approx \theta \cdot N$, the results are summarised in Table 4.6. As shown in the table, even if we set the threshold to be $\text{HD} = 0.4$ to give more tolerance to intra-group variations, the false acceptance rate remains negligible.

4.7.2 Online Application

First of all, we propose an online application, which works with an existing *unmodified* banknote. Here, we will leverage the fact that each banknote has a unique serial number, as we will explain below.

We divide an online application into two phrases: registration and verification. During the registration phase, a polymer substrate fingerprint for each newly manufactured polymer banknote is extracted and recorded in a database along with a unique serial number of the banknote. In the verification phase,

Table 4.6: False match for one-to-one comparison

HD threshold	Probability of False Match
0.3	3.5×10^{-34}
0.31	6.0×10^{-31}
0.32	6.7×10^{-28}
0.33	5.0×10^{-25}
0.34	2.5×10^{-22}
0.35	8.2×10^{-20}
0.36	1.9×10^{-17}
0.37	2.9×10^{-15}
0.38	3.0×10^{-13}
0.39	2.2×10^{-11}
0.4	1.1×10^{-9}

a fresh photograph of the feature area is taken and processed into a compact 2048-bit fingerprint. The fingerprint, along with the banknote serial number, is then sent to a remote server through a secure channel (e.g., SSL/TLS). Based on the serial number, the server retrieves the reference fingerprint and compares it with the sample fingerprint against a HD threshold. Finally, the verification result is communicated back to the client through the existing secure channel.

Thanks to the unique serial number, the verification in the online application is based on one-to-one comparison (rather than one-to-many as required in exhaustive search). This is not only extremely fast, but also gives great flexibility in choosing a threshold. (as seen in Section 4.7.1)

4.7.3 Offline Application

An offline application differs from an online one by printing the registration information onto a banknote rather than saving it to a database. However, this adds an extra step of registration to the existing banknote manufacturing process. Figure 4.9 summarises the process of the registration. The feature vector extracted from the translucent patterns of a polymer substrate is digitally signed, along with other contextual information such as the serial number and denomination value. The private signing key is kept securely by the authorities who issue banknotes. In the proof-of-concept implementation, we use ECDSA with 512 bits key length (256-bit security) for digital signing. Encoded in Base64, the total length of the message and the digital signature is approximately 4420 bits, which can be fit into a QR code (version 18) with medium error correction (as shown in Figure 4.9).

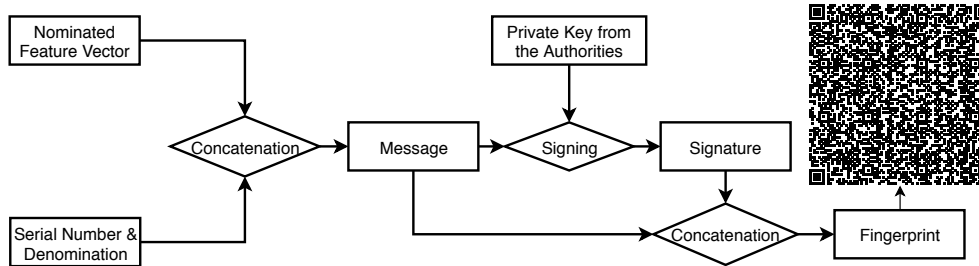


Figure 4.9: The Procedure of Fingerprint Registration. The QR code shown in the diagram is generated from a real £10 polymer banknote.

Fuzzy encryption is a common technique to encrypt a biometric sample such that it can only be decrypted by another biometric sample taken from the same subject [77]. We could apply the same technique to encrypt the fingerprint contained in the QR code, however, in the context of our application, a counterfeiter always has physical access to real banknotes which they wish to counterfeit. Hence, encrypting the fingerprint does not offer real security benefit in our case. For this reason, we simply save the fingerprint in its plain text in the QR code, but we add a digital signature to protect the integrity of data.

During the verification phase, data from the QR code is first read, which contain a reference fingerprint, a digital signature and other data. After the digital signature is verified successfully by using a public key, a fresh image of the feature area is taken and then processed into a sample fingerprint. This fingerprint will be compared with the reference fingerprint against a HD threshold with a binary outcome: *accept* or *reject*. The integrity of the content in the QR code is protected through the digital signature. If an attacker copies the same QR code to a different banknote, the verification will fail as the two fingerprints will not match.

4.8 Security Analysis

4.8.1 Threat Model

We assume that the attacker knows everything about the fingerprinting of banknotes. He knows the full state of the device and exactly which area is used for feature extraction and how the feature vector is computed from the feature area. We further assume that the attacker has access to effectively the same printing material and equipment as used for producing legitimate banknotes.

Under this assumption, the security of the existing polymer banknotes will be easily broken. While the security of existing banknotes is easily broken in this threat model, we aim to provide additional security assurance such that counterfeiting remains difficult.

The attacker has several limitations. First of all, we assume he is unable to obtain the private signing key used by the banknote issuing authority. Furthermore, we assume the adversary is unable to physically clone the same features of a polymer substrate. We emphasise that our security protection is in addition, and orthogonal, to existing security features on a banknote. In reality, a security feature is considered effective if it raises the cost of counterfeiting above the nominal value of the banknote.

4.8.2 Attack on Fingerprints

Like every human being is unique (which forms the basis of “biometrics”), every physical object is unique too (which forms the basis of “physical unclonable function”). Under the microscopic view, every object has distinguished features that can not be exactly duplicated. The same applies to the polymer substrate. Its unevenness in the opacity coating layer reflects the imperfections during the opacity coating, which cannot be avoided. The counterfeiter’s challenge is to make another polymer substrate, which gives the same or sufficiently similar feature vector as that of a genuine banknote so the same digital signature can be reused to legitimise the counterfeit.

We argue that it is hard for the attacker to make another polymer substrate that matches a given 2048-bit feature vector even if he has access to the same printing material and equipment as used by the banknote issuing authority. First of all, the attacker needs to produce a “physical” object that looks and feels like a legitimate polymer note. This is substantially harder than launching spoofing attacks in “biometrics”, e.g., using a gummy finger to deceive a scanner in an unmanned (unsupervised) environment. By contrast, the verification of banknotes is usually “supervised” by nature. Visual inspection by a human is almost always the first line of defence to detect counterfeits, which is followed by the possible use of tools for further confirmation such as special pens, UV light, or in our case, a film scanner. With reference to the gummy-finger attack, this means an attacker has to make a real “finger” that looks and feels like a human finger to pass the human inspection first, before it can deceive the fingerprint scanner. This is substantially more difficult than a conventional

spoofing attack in an unsupervised environment.

The use of the “see-through” window (a security feature of polymer notes) forces the counterfeiter to use a clear plastic film as the substrate. With access to the same printing material and equipment as used by the bank authorities, the attacker will be able to produce polymer substrates that look and feel the same as legitimate notes. However, merely producing another substrate gives only a probability of $p = 5 \times 10^{-25}$ for mismatch, based on an HD threshold of 0.33 (Figure 4.6). In reality, the feature vector of a second polymer substrate does not have to match exactly that of a target substrate. It only needs to be close enough in the Hamming space, say less than an HD distance of 0.33. Based on the degrees of freedom $N = 900$ (mean HD 0.5), and an HD threshold of 0.33, finding a random N -bit string that is within the $w = 0.33 \cdot N = 297$ bits difference to the target string requires the minimum number of attempts N' as estimated below according to the sphere-packing bound [77].

$$\begin{aligned}
 N' &= \frac{2^N}{\sum_{i=0}^w \binom{N}{i}} \\
 &= 4 \times 10^{24}
 \end{aligned}
 \tag{4.15}$$

Note that N' is only a lower bound. The above result implies that if the attacker repeats the same production process, he must produce 4×10^{24} polymer substrates in order to find one that might match a given digitally signed fingerprint. This is clearly infeasible for the attacker.

In the extreme case that the attacker has the ability to collect all of the PSFs in global circulation² (about 5×10^{11}), he still need to try more than 10^{12} times to get a match, which is also unlikely to happen in the real world.

The attacker might improve his chance by adding a custom-built printing step on top of the existing banknote manufacturing process. It is worth noting that printing on a plastic film is much harder than printing on a paper substrate. The novel idea that uses a special plastic film made of BOPP to support high-quality security printing is precisely the key innovation that makes polymer notes possible [142]. However, the film still has to undergo a special opacity coating process to form a polymer substrate, which provides a canvas to allow printing in the subsequent procedure.

As explained earlier, the opacity coating is inherently imperfect, producing a layer of white ink with uneven thicknesses. This leads to random translucent

²Assuming all the banknotes have applied PSF

patterns when the light shines through the substrate. A close-up of the translucent patterns is shown in Figure 4.10. As shown in the picture, the patterns contain randomly distributed bright spots, as well as dark spots (impurities in the ink). The physical dimensions of these features are on the scale of a few micrometres. As a comparison, high-resolution ink-jet printers use very small drops (normally 17 to 50 pL volume of liquid in one droplet [109]) to create different colours or grey levels. With a volume of $v = 17$ pL, assume it forms a perfect semi-sphere once it falls on the substrate to form a printed dot, the diameter of the dot is $d = 2 \cdot \sqrt[3]{v \cdot 2 \cdot 3/4\pi} = 40 \mu\text{m}$. However, in reality, the droplet collides with the substrate at a high-speed, creating a much larger dot with randomly scattering patterns which resemble nothing like a dot under the microscopic view (e.g., see [42]). While an ink-jet printer is physically limited by the size of the nozzle, an attacker might use a laser printer. However, a laser printer has its own physical limitation. Due to the interaction of multiple rolls, a laser printer prints uncontrollable repeated patterns at the microscopic view [96]. As an experiment, we used two high-resolution inkjet (HP Deskjet 2700) and laser (Kyocera TASKalfa 5052ci) printers to print a dot ‘.’ in different font sizes as shown in Figure 4.11. The smallest printed size is at least one order of magnitude larger than the size of the impurities observed in the opacity coating layer (see Figure 4.10). More importantly, the printed dots in Figure 4.11 exhibit random scattering patterns because the printers cannot precisely control the nozzle or the toner at the microscopic level.

Hence, modern printers have physical limits in what they can print at the microscopic level. While on-top printing can increase the opacity level, the attacker also needs to be able to decrease the opacity level, e.g., by removing white ink in the coating layer, so to have the full control of the translucent patterns. This will require the attacker to acquire much more sophisticated printing equipment than what is used by a legitimate state government. While this is theoretically possible, we believe it is extremely unlikely in practice, and we leave it to further research in the future.

4.9 Conclusion

In this paper, we have proposed a new anti-counterfeiting solution for polymer notes based on analysing the imperfections in the opacity coating of a polymer substrate. The imperfect coating process leaves a coating layer of uneven thickness and randomly distributed impurities from the ink. We propose a




Figure 4.10: A close-up of translucent patterns



















Font Size	Inkjet (HP Deskjet 2700)			Laserjet (Kyocera TASKalfa 5052ci)		
1	 0.276 mm	 0.122 mm	 0.256 mm	 0.161 mm	 0.151 mm	 0.131 mm
2	 0.175 mm	 0.142 mm	 0.274 mm	 0.185 mm	 0.198 mm	 0.201 mm
3	 0.245 mm	 0.220 mm	 0.260 mm	 0.246 mm	 0.235 mm	 0.250 mm

Figure 4.11: A printed dot ‘.’ in different font sizes using inkjet and laser printers. The resolution for both printers is 1200 DPI.

method to capture and transform these imperfections into a unique fingerprint. Our experiments show that our solution is able to authenticate banknotes with high accuracy, is extremely scalable, and is robust against rough daily handling of banknotes.

Chapter 5

Conclusion and Future Work

5.1 Summary

In this thesis, we presented two new approaches for authentication. The first approach is the caller ID verification (CIV) protocol, which uses a challenge-response-based one-time password (OTP) authentication method to detect and prevent caller ID spoofing attacks. To the best of our knowledge, this is the only solution that can work across landline, cellular, and IP networks. The second approach is the use of polymer substrate fingerprint (PSF) for authenticating banknotes. Through extensive experimentation, we have demonstrated that PSF is a highly accurate and scalable solution for addressing the issue of banknote counterfeiting.

In Chapter 2, we conducted a comprehensive review of authentication approaches that utilise different factors to verify a user's identity. We also summarised the challenges and benefits associated with each approach, as reported in the literature.

In Chapter 3, we presented the Caller ID Verification (CIV) protocol as a case study for ownership authentication. In the past, caller ID spoofing was not a significant issue in the legacy telephone network because it was expensive and carried a high risk of exposure for the attacker. However, with the widespread use of Voice over IP (VoIP), caller ID can now be easily spoofed by even inexperienced attackers at low cost. We observed that working with DTMF tones, the spoofed caller ID can be leveraged to solve the spoofing problem. The proposed protocol verifies the caller ID through a challenge-response process using the secret carried by the modified caller ID that is transmitted between the authentic caller and the callee. We designed prototypes for three platforms

working on different telecommunication networks, including Android on the cellular network, trueCall box on the PSTN network, and SIP on the IP network. With extensive experiments, we demonstrated that our proposed solution can detect spoofed caller IDs within 5 seconds when the two participants are on the IP network. The longest delay in our experiments was around 29 seconds when the protocol was working across the landline and cellular networks. We argued that this delay can be substantially reduced if constraints on the platforms were removed. As the dominant delay in the protocol was found to be the call setup, CIV can be optimised by simplifying the challenge signal to eliminate the need for a full bidirectional voice call setup in the authentication process.

The second solution proposed in Chapter 4 utilises biometrics-based authentication techniques to address the problem of banknote counterfeiting. With the advancement of printing and image processing technologies, organised forgers are able to produce high-quality counterfeits, which pose a serious threat to society and economy. Our proposed method is based on the observation that the opacity coating process used in the production of polymer banknotes results in uneven thickness in the coating layer and random dispersion of impurities from the ink, creating unique translucent patterns when viewed under back-lighting. The solution is designed to protect polymer banknotes, but it can also be adapted to secure paper banknotes. Due to the complex process of paper sheet manufacture, the unpredictable pattern was found on the paper banknotes as well. We designed two prototypes: one is based on a commercial film scanner and the other is based on a professional camera with a light-box. To evaluate the performance of our prototypes, we collected eight datasets containing 6200 sample images from 340 different banknotes issued in the UK [168, 169]. Both setups showed promising results when assessed using the biometrics evaluation framework. Using a threshold of $HD = 0.33$ for Hamming distance, we achieved a 0% false match rate on £5, £10 polymer banknotes and £20 paper banknotes. Our evaluation using the PUF evaluation framework also produced consistent results. We found that the extracted feature (length of 2048 bits) from the benchmark dataset contained 900 independent bits, which is much higher than the degree of freedom of 249 in the 2048-bit iris code. This high degree of freedom in banknote fingerprints allows our solution to reliably identify individual banknotes in circulation around the world and makes counterfeiting impracticable even with the use of sophisticated printing equipment by the forgers.

5.2 Future work

Future works are suggested as follows.

- In Chapter 3, we discussed the optimisation of the Caller ID Verification (CIV) protocol. As we found that the delay in the protocol was mainly due to the call setup, we plan to build a testbed environment to transmit a simplified challenge signal, similar to the 'ping' command in the IP network, instead of establishing a full bidirectional voice call. This approach is expected to significantly reduce the overall delay of the protocol.
- Additionally, in Chapter 3, we noted that at the end of the CIV process, the caller and callee share a one-time password (OTP). This OTP could be used to encrypt the audio payload in the scenario of one SIP phone calling another SIP phone. However, the 4-digit secret used in the prototypes has inadequate entropy for cryptographic applications. To address this, we plan to use a password-authenticated key exchange (PAKE) protocol to convert the weak OTP into a cryptographically secure password. This stronger password will then be used to encrypt the Secure Real-time Transport Protocol (SRTP) stream after the SIP negotiation. By adopting the OTP and PAKE, the sender and receiver can establish an end-to-end secure communication channel using CIV.
- In Chapter 4, we evaluated our proposed Polymer Substrate Fingerprint (PSF) solution on £5 and £10 polymer banknotes issued in the UK and obtained promising results. We plan to expand our research to include the new £20 and £50 polymer banknotes issued in the UK as well as banknotes issued in other countries.
- Finally, in the experiments described in Chapter 4, we implemented a prototype with the setup of camera. The setup produced less favourable results compared to the other prototype using the film scanner but was still convincing with a 0% false rejection rate and false acceptance rate when the threshold was set to 0.33. This demonstrates the possibility of using non-specialised equipment for data capture. With the widespread adoption of smartphones and the rapid development of embedded CMOS and lens technology, more and more smartphones are capable of taking high-quality images. To make PSF more widely accessible, we plan

to expand its use to smartphones with top-notch quality lens, as the capturing of the detail-rich image with non-distortion is necessary for getting satisfactory results. However, adjustments in light intensity, camera settings, image processing, and Gabor filter parameters may be required due to the differences in design between professional cameras and phone cameras.

Bibliography

- [1] About face id advanced technology. URL <https://support.apple.com/en-gb/HT208108>.
- [2] How to check your banknotes — bank of england. URL <https://www.bankofengland.co.uk/banknotes/counterfeit-banknotes/how-to-check-your-banknotes>.
- [3] Bbb scam alert: "neighbor spoofing" is a common type of phone scam. URL <https://www.bbb.org/article/news-releases/16670-a-new-kind-of-phone-scam-neighbor-spoofing>.
- [4] Scams research 2021 chart pack. URL https://www.ofcom.org.uk/_data/assets/pdf_file/0029/232877/2021-ofcom-scams-survey.pdf.
- [5] Hid activid one-time password (otp) tokens. URL <https://www.hidglobal.com/products/one-time-password-tokens>.
- [6] Polymer banknotes — bank of england. URL <https://www.bankofengland.co.uk/banknotes/polymer-banknotes>.
- [7] The UK's PSTN network will switch off in 2025. URL <https://business.bt.com/why-choose-bt/insights/digital-transformation/uk-pstn-switch-off/>.
- [8] Caller id spoofing scams. URL <https://www.identityforce.com/blog/caller-id-spoofing-scam>.
- [9] Banknote statistics — bank of england, . URL <https://www.bankofengland.co.uk/statistics/banknote>.

- [10] Scam call trends and projections report, fall 2018, . URL [https://ecfsapi.fcc.gov/file/109272058817712/FirstOrion_Scam_Trends_Report_FINAL%20\(002\)%20\(002\).pdf](https://ecfsapi.fcc.gov/file/109272058817712/FirstOrion_Scam_Trends_Report_FINAL%20(002)%20(002).pdf).
- [11] Scam call trends and projections report, summer 2019, . URL http://firstorion.com/wp-content/uploads/2019/07/First-Orion-Scam-Trends-Report_Summer-2019.pdf.
- [12] Voip market size forecast 2021-2027 — industry share analysis report. URL <https://www.gminsights.com/industry-analysis/voice-over-internet-protocol-voip-market>.
- [13] TrueCall. <https://www.truecall.co.uk/>, .
- [14] Truecaller. <https://www.truecaller.com/>, .
- [15] Who’s really calling you? an investigation into the worrying rise of ‘number spoofing’. URL <https://www.which.co.uk/news/2019/10/whos-really-calling-you-an-investigation-into-the-worrying-rise-of-number-spoofing/>.
- [16] Withdrawn banknotes — bank of england. URL <https://www.bankofengland.co.uk/banknotes/withdrawn-banknotes>.
- [17] Mohammed Abuhamad, Ahmed Abusnaina, DaeHun Nyang, and David Mohaisen. Sensor-based continuous authentication of smartphones’ users using behavioral biometrics: A contemporary survey. *IEEE Internet of Things Journal*, 8(1):65–84, 2020.
- [18] Md Liakat Ali, John V Monaco, Charles C Tappert, and Meikang Qiu. Keystroke biometric systems for user authentication. *Journal of Signal Processing Systems*, 86(2):175–190, 2017.
- [19] Israa M Alsaadi. Physiological biometric authentication systems, advantages, disadvantages and future development: A review. *International Journal of Scientific & Technology Research*, 4(12):285–289, 2015.
- [20] Ross Anderson. *Security engineering: a guide to building dependable distributed systems*. John Wiley & Sons, 2020.
- [21] K Aravindhan and RR Karthiga. One time password: A survey. *International Journal of Emerging Trends in Engineering and Development*, 1(3):613–623, 2013.

- [22] BS Archana, Ashika Chandrashekar, Anusha Govind Bangi, BM Sanjana, and Syed Akram. Survey on usable and secure two-factor authentication. In *2017 2nd IEEE International Conference on Recent Trends in Electronics, Information & Communication Technology (RTEICT)*, pages 842–846. IEEE, 2017.
- [23] Gonzalo Bailador, Carmen Sanchez-Avila, Javier Guerra-Casanova, and Alberto de Santos Sierra. Analysis of pattern recognition techniques for in-air signature biometrics. *Pattern Recognition*, 44(10-11):2468–2478, 2011.
- [24] Meika Ball. Recent trends in banknote counterfeiting. *Reserve Bank of Australia Bulletin*, 1, 2019.
- [25] S Bansal, P Bruno, O Denecker, M Goparaju, and M Niederkprn. Global payments 2018: A dynamic industry continues to break new ground. *Global Banking McKinsey*, 2018.
- [26] Luciano Bello, Maximiliano Bertacchini, Carlos Benitez, Juan Carlos Pizzoni, and Marcelo Cipriano. Collection and publication of a fixed text keystroke dynamics dataset. In *XVI Congreso Argentino de Ciencias de la Computación*, 2010.
- [27] Albert Berenguel, Oriol Ramos Terrades, Josep Lladós, and Cristina Cañero. Banknote counterfeit detection through background texture printing analysis. In *2016 12th IAPR Workshop on Document Analysis Systems (DAS)*, pages 66–71. IEEE, 2016.
- [28] Guido Bertoni, Joan Daemen, Seth Hoffert, Michaël Peeters, Gilles Van Assche, and Ronny Van Keer. Farfalle: parallel permutation-based cryptography. *Cryptology ePrint Archive*, 2016.
- [29] Charles Beumier and Thibault Debatty. Attack detection in ss7. In *International Conference on Multimedia Communications, Services and Security*, pages 11–20. Springer, 2022.
- [30] Francesco Bianconi and Antonio Fernández. Evaluation of the effects of gabor filter parameters on texture classification. *Pattern recognition*, 40(12):3325–3335, 2007.

- [31] Ramon Blanco-Gonzalo, Raul Sanchez-Reillo, Oscar Miguel-Hurtado, and Judith Liu-Jimenez. Performance evaluation of handwritten signature recognition in mobile environments. *IET biometrics*, 3(3):139–146, 2014.
- [32] Joseph Bonneau, Cormac Herley, Paul C Van Oorschot, and Frank Stajano. Passwords and the evolution of imperfect authentication. *Communications of the ACM*, 58(7):78–87, 2015.
- [33] Patrick Bours and Hafez Barghouti. Continuous authentication using biometric keystroke dynamics. In *The Norwegian Information Security Conference (NISK)*, volume 2009, 2009.
- [34] Chad Boutin. Nist evaluation shows advance in face recognition software’s capabilities, 2018.
- [35] Colin Boyd. Digital multi-signatures. cryptography and coding (hj beker and fc piper eds.), 1989.
- [36] James DR Buchanan, Russell P Cowburn, Ana-Vanessa Jausovec, Dorothee Petit, Peter Seem, Gang Xiong, Del Atkinson, Kate Fenton, Dan A Allwood, and Matthew T Bryan. Forgery: ‘fingerprinting’ documents and packaging. *Nature*, 436(7050):475, 2005.
- [37] Philippe Bulens, F-X Standaert, and J-J Quisquater. How to strongly link data and its medium: the paper case. *IET Information Security*, 4(3):125–136, 2010.
- [38] Reynel Castrillón, Alejandro Acien, Juan Rafael Orozco-Aroyave, Aythami Morales, JF Vargas, Rubén Vera-Rodríguez, Julian Fierrez, Javier Ortega-Garcia, and A Villegas. Characterization of the handwriting skills as a biomarker for parkinson’s disease. In *2019 14th IEEE International Conference on Automatic Face & Gesture Recognition (FG 2019)*, pages 1–5. IEEE, 2019.
- [39] Albert Berenguel Centeno, Oriol Ramos Terrades, Josep Lladós i Canet, and Cristina Cañero Morales. Evaluation of texture descriptors for validation of counterfeit documents. In *2017 14th IAPR International Conference on Document Analysis and Recognition (ICDAR)*, volume 1, pages 1237–1242. IEEE, 2017.

- [40] Gogineni Krishna Chaitanya and Krovi Raja Sekhar. Verification of pattern unlock and gait behavioural authentication through a machine learning approach. *International Journal of Intelligent Unmanned Systems*, 2021.
- [41] Stanley T Chow, Christophe Gustave, and Dmitri Vinokurov. Authenticating displayed names in telephony. *Bell Labs Technical Journal*, 14(1):267–282, 2009.
- [42] William Clarkson, Tim Weyrich, Adam Finkelstein, Nadia Heninger, J Alex Halderman, and Edward W Felten. Fingerprinting blank paper using commodity scanners. In *2009 30th IEEE Symposium on Security and Privacy*, pages 301–314. IEEE, 2009.
- [43] Victoria Cleland. Insights into the future of cash. <https://tinyurl.com/qsafb51>, July 2017.
- [44] US Congress. Truth in caller id act of 2009, . URL <https://www.gpo.gov/>.
- [45] US Congress. Repack airwaves yielding better access for users of modern service, . URL <https://www.gpo.gov/>.
- [46] Radigya M Correia, Eloilson Domingos, Flavia Tosato, Luiz Felipe M Aquino, André M Fontes, Vagne M Cáo, Paulo R Filgueiras, and Wander-son Romão. Banknote analysis by portable near infrared spectroscopy. *Forensic Chemistry*, 8:57–63, 2018.
- [47] National Research Council et al. *Counterfeit deterrent features for the next-generation currency design*, volume 472. National Academies Press, 1993.
- [48] Joseph Cox. The secret sims used by criminals to spoof any number. URL <https://www.vice.com/en.us/article/n7w9pw/russian-sims-encrypted>.
- [49] Ioannis G Damousis and Savvas Argyropoulos. Four machine learning algorithms for biometrics fusion: A comparative study. *Applied Computational Intelligence and Soft Computing*, 2012, 2012.

- [50] John Daugman. Statistical richness of visual phase information: update on recognizing persons by iris patterns. *International Journal of computer vision*, 45(1):25–38, 2001.
- [51] John Daugman. How iris recognition works. In *The essential guide to image processing*, pages 715–739. Elsevier, 2009.
- [52] Mariana R de Almeida, Deleon N Correa, Werickson FC Rocha, Francisco JO Scafi, and Ronei J Poppi. Discrimination between authentic and counterfeit banknotes using raman spectroscopy and pls-da with uncertainty estimation. *Microchemical Journal*, 109:170–177, 2013.
- [53] Claudio De Stefano, Francesco Fontanella, Donato Impedovo, Giuseppe Pirlo, and Alessandra Scotto di Freca. Handwriting analysis to support neurodegenerative diseases diagnosis: A review. *Pattern Recognition Letters*, 121:37–45, 2019.
- [54] Gerald DeJean and Darko Kirovski. Rf-dna: Radio-frequency certificates of authenticity. In *International Workshop on Cryptographic Hardware and Embedded Systems*, pages 346–363. Springer, 2007.
- [55] Haotian Deng, Weicheng Wang, and Chunyi Peng. Ceive: Combating caller id spoofing on 4g mobile phones via callee-only inference and verification. In *24th Annual International Conference on Mobile Computing and Networking*, pages 369–384, 2018.
- [56] Mohammad Omar Derawi, Claudia Nickel, Patrick Bours, and Christoph Busch. Unobtrusive user-authentication on mobile phones using biometric gait recognition. In *2010 Sixth International Conference on Intelligent Information Hiding and Multimedia Signal Processing*, pages 306–311. IEEE, 2010.
- [57] Peter Drotár, Jiří Mekyska, Irena Rektorová, Lucia Masarová, Zdeněk Smékal, and Marcos Faundez-Zanuy. A new modality for quantitative evaluation of parkinson’s disease: In-air movement. In *13th IEEE international conference on bioinformatics and bioengineering*, pages 1–4. IEEE, 2013.
- [58] Lee Dryburgh and Jeff Hewett. *Signaling System No. 7 (SS7/C7): protocol, architecture, and services*. Cisco press, 2005.

- [59] William G Eckert. *Introduction to forensic sciences*. CRC press, 1996.
- [60] Alexander Eng and Luay A Wahsheh. Look into my eyes: A survey of biometric security. In *2013 10th International Conference on Information Technology: New Generations*, pages 422–427. IEEE, 2013.
- [61] Marcos Faundez-Zanuy, Julian Fierrez, Miguel A Ferrer, Moises Diaz, Ruben Tolosana, and Réjean Plamondon. Handwriting biometrics: Applications and future trends in e-security and e-health. *Cognitive Computation*, 12(5):940–953, 2020.
- [62] FCC. Combating spoofed robocalls with caller id authentication. <https://www.fcc.gov/call-authentication>.
- [63] FCC. Caller id spoofing, 2022. <https://www.fcc.gov/spoofing>.
- [64] UK Finance. Uk payment markets 2019, 2109. URL <https://www.ukfinance.org.uk/sites/default/files/uploads/pdf/UK-Finance-UK-Payment-Markets-Report-2019-SUMMARY.pdf>.
- [65] Javier Galbally, Fernando Alonso-Fernandez, Julian Fierrez, and Javier Ortega-Garcia. A high performance fingerprint liveness detection method based on quality related features. *Future Generation Computer Systems*, 28(1):311–321, 2012.
- [66] Ravi Ganesan. Yaksha: Augmenting kerberos with public key cryptography. In *Proceedings of the Symposium on Network and Distributed System Security*, pages 132–143. IEEE, 1995.
- [67] Ravi Ganesan and Yacov Yacobi. A secure joint signature and key exchange system. *Bellcore TM*, 24531, 1994.
- [68] Blaise Gassend, Dwaine Clarke, Marten Van Dijk, and Srinivas Devadas. Silicon physical random functions. In *Proceedings of the 9th ACM Conference on Computer and Communications Security*, pages 148–160, 2002.
- [69] Luca Ghiani, Gian Luca Marcialis, and Fabio Roli. Fingerprint liveness detection by local phase quantization. In *Proceedings of the 21st international conference on pattern recognition (ICPR2012)*, pages 537–540. IEEE, 2012.

- [70] Tom Gillespie. Money for nothing: The story of the biggest counterfeiter in us history. <https://news.sky.com/story/money-for-nothing-the-story-of-the-biggest-counterfeiter-in-us-history-11942377>.
- [71] Giacomo Giorgi, Andrea Saracino, and Fabio Martinelli. Using recurrent neural networks for continuous authentication through gait analysis. *Pattern Recognition Letters*, 147:157–163, 2021.
- [72] Romain Giot, Mohamad El-Abed, and Christophe Rosenberger. Greyc keystroke: a benchmark for keystroke dynamics biometric systems. In *2009 IEEE 3rd International Conference on Biometrics: Theory, Applications, and Systems*, pages 1–6. IEEE, 2009.
- [73] Jorge Guajardo, Sandeep S Kumar, Geert-Jan Schrijen, and Pim Tuyls. Fpga intrinsic pufs and their use for ip protection. In *International workshop on cryptographic hardware and embedded systems*, pages 63–80. Springer, 2007.
- [74] Jorge Guajardo, Boris Škorić, Pim Tuyls, Sandeep S Kumar, Thijs Bel, Antoon HM Blom, and Geert-Jan Schrijen. Anti-counterfeiting, key distribution, and key storage in an ambient world via physical unclonable functions. *Information Systems Frontiers*, 11(1):19–41, 2009.
- [75] Neil Haller. The s/key one-time password system. Technical report, 1995.
- [76] Ghaith Hammouri, Aykutlu Dana, and Berk Sunar. Cds have fingerprints too. In *International Workshop on Cryptographic Hardware and Embedded Systems*, pages 348–362. Springer, 2009.
- [77] Feng Hao, Ross Anderson, and John Daugman. Combining crypto with biometrics effectively. *IEEE transactions on computers*, 55(9):1081–1088, 2006.
- [78] Ryan Helinski, Dhruva Acharyya, and Jim Plusquellic. A physical unclonable function defined using power distribution system equivalent resistance variations. In *2009 46th ACM/IEEE Design Automation Conference*, pages 676–681. IEEE, 2009.
- [79] Daniel E Holcomb, Wayne P Bursleson, Kevin Fu, et al. Initial sram state as a fingerprint and source of true random numbers for rfid tags. In *Proceedings of the Conference on RFID Security*, volume 7, page 01, 2007.

- [80] Bin-Tsan Hsieh, Hung-Min Sun, and Tzonelih Hwang. On the security of some password authentication protocols. *Informatica*, 14(2):195–204, 2003.
- [81] Yun Huang, Zheng Huang, Haoran Zhao, and Xuejia Lai. A new one-time password method. *IERI Procedia*, 4:32–37, 2013.
- [82] Jing-Jang Hwang and Tzu-Chang Yeh. Improvement on peyravian-zunic’s password authentication schemes. *IEICE Transactions on Communications*, 85(4):823–825, 2002.
- [83] Min-Shiang Hwang, Cheng-Chi Lee, and Yuan-Liang Tang. A simple remote user authentication scheme. *Mathematical and Computer Modelling*, 36(1-2):103–107, 2002.
- [84] IETF. Personal Assertion Token (PaSSporT) Extension for Signature-based Handling of Asserted information using toKENs (SHAKEN), . <https://datatracker.ietf.org/doc/rfc8588/>.
- [85] IETF. Secure Telephone Identity Revisited (STIR), . <https://datatracker.ietf.org/wg/stir/documents/>.
- [86] Tanya Ignatenko, Geert-Jan Schrijen, Boris Skoric, Pim Tuyls, and Frans Willems. Estimating the secrecy-rate of physical unclonable functions with the context-tree weighting method. In *2006 IEEE International Symposium on Information Theory*, pages 499–503. IEEE, 2006.
- [87] Donato Impedovo. Velocity-based signal features for the assessment of parkinsonian handwriting. *IEEE Signal Processing Letters*, 26(4): 632–636, 2019.
- [88] Ronald S Indeck and Marcel W Muller. Method and apparatus for fingerprinting magnetic media, November 15 1994. US Patent 5,365,586.
- [89] Philip G Inglesant and M Angela Sasse. The true cost of unusable password policies: password use in the wild. In *Proceedings of the sigchi conference on human factors in computing systems*, pages 383–392, 2010.
- [90] Ryoichi Isawa and Masakatu Morii. One-time password authentication scheme to solve stolen verifier problem. *Information Processing Society of Japan and The Institute of Electronics, Information and Communication Engineers*, pages 225–228, 2011.

- [91] Anil K Jain, Patrick Flynn, and Arun A Ross. *Handbook of biometrics*. Springer Science & Business Media, 2007.
- [92] MA Zamalloa Jara, C Luízar Obregón, and C Araujo Del Castillo. Exploratory analysis for the identification of false banknotes using portable x-ray fluorescence spectrometer. *Applied Radiation and Isotopes*, 135: 212–218, 2018.
- [93] Khwaja Jawad, Khwaja Mansoor, Ahmed Fraz Baig, Anwar Ghani, and Azmat Naseem. An improved three-factor anonymous authentication protocol for wsn s based iot system using symmetric cryptography. In *2019 International Conference on Communication Technologies (ComTech)*, pages 53–59. IEEE, 2019.
- [94] Paramjit Kaur, Kewal Krishan, Suresh K Sharma, and Tanuj Kanchan. Facial-recognition algorithms: A literature review. *Medicine, Science and the Law*, 60(2):131–139, 2020.
- [95] Kevin S Killourhy and Roy A Maxion. Comparing anomaly-detection algorithms for keystroke dynamics. In *2009 IEEE/IFIP International Conference on Dependable Systems & Networks*, pages 125–134. IEEE, 2009.
- [96] Do-Guk Kim and Heung-Kyu Lee. Colour laser printer identification using halftone texture fingerprint. *Electronics Letters*, 51(13):981–983, 2015.
- [97] Jae-Jung Kim and Seng-Phil Hong. A method of risk assessment for multi-factor authentication. *Journal of Information Processing Systems*, 7(1):187–198, 2011.
- [98] Leslie Lamport. Password authentication with insecure communication. *Communications of the ACM*, 24(11):770–772, 1981.
- [99] Jae W Lee, Daihyun Lim, Blaise Gassend, G Edward Suh, Marten Van Dijk, and Srinivas Devadas. A technique to build a secret key in integrated circuits for identification and authentication applications. In *2004 Symposium on VLSI Circuits. Digest of Technical Papers (IEEE Cat. No. 04CH37525)*, pages 176–179. IEEE, 2004.
- [100] Michael Lesk. Caller id: Whose privacy? *IEEE security & privacy*, 12(2):77–79, 2014.

- [101] Jikai Li, Fernando Faria, Jinsong Chen, and Daan Liang. A mechanism to authenticate caller id. In *World Conference on Information Systems and Technologies*, pages 745–753. Springer, 2017.
- [102] Stan Z Li. *Encyclopedia of Biometrics: I-Z.*, volume 2. Springer Science & Business Media, 2009.
- [103] Daihyun Lim, Jae W Lee, Blaise Gassend, G Edward Suh, Marten Van Dijk, and Srinivas Devadas. Extracting secret keys from integrated circuits. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, 13(10):1200–1205, 2005.
- [104] Keith Lofstrom, W Robert Daasch, and Donald Taylor. Ic identification circuit using device mismatch. In *2000 IEEE International Solid-State Circuits Conference. Digest of Technical Papers (Cat. No. 00CH37056)*, pages 372–373. IEEE, 2000.
- [105] Roel Maes and Ingrid Verbauwhede. Physically unclonable functions: A study on the state of the art and future research directions. *Towards Hardware-Intrinsic Security*, pages 3–37, 2010.
- [106] Roel Maes, Pim Tuyls, and Ingrid Verbauwhede. Intrinsic pufs from flip-flops on reconfigurable devices. In *3rd Benelux workshop on information and system security (WISSec 2008)*, volume 17, page 2008, 2008.
- [107] Abhranil Maiti, Vikash Gunreddy, and Patrick Schaumont. A systematic method to evaluate and compare the performance of physical unclonable functions. In *Embedded systems design with FPGAs*, pages 245–267. Springer, 2013.
- [108] Khwaja Mansoor, Anwar Ghani, Shehzad Ashraf Chaudhry, Shahabuddin Shamsirband, Shahbaz Ahmed Khan Ghayyur, and Amir Mosavi. Securing iot-based rfid systems: A robust authentication protocol using symmetric cryptography. *Sensors*, 19(21):4752, 2019.
- [109] Graham D Martin, Stephen D Hoath, and Ian M Hutchings. Inkjet printing-the physics of manipulating liquid jets and drops. In *Journal of Physics: Conference Series*, volume 105, page 012001. IOP Publishing, 2008.
- [110] Alfred J Menezes, Paul C Van Oorschot, and Scott A Vanstone. *Handbook of applied cryptography*. CRC press, 2018.

- [111] Jugurta R Montalvão Filho and Eduardo O Freire. On the equalization of keystroke timing histograms. *Pattern Recognition Letters*, 27(13): 1440–1446, 2006.
- [112] Ahmadreza Montazerolghaem. Softwarization and virtualization of voip networks. *The Journal of Supercomputing*, pages 1–33, 2022.
- [113] Hossen Mustafa, Wenyuan Xu, Ahmad Reza Sadeghi, and Steffen Schulz. You can call but you can't hide: detecting caller id spoofing attacks. In *44th Annual IEEE/IFIP International Conference on Dependable Systems and Networks*, pages 168–179. IEEE, 2014.
- [114] Hossen Mustafa, Wenyuan Xu, Ahmad-Reza Sadeghi, and Steffen Schulz. End-to-end detection of caller id spoofing attacks. *IEEE Transactions on Dependable and Secure Computing*, 15(3):423–436, 2016.
- [115] Javier Nieves, Igor Ruiz-Agundez, and Pablo G Bringas. Recognizing banknote patterns for protecting economic transactions. In *2010 Workshops on Database and Expert Systems Applications*, pages 247–249. IEEE, 2010.
- [116] Gilbert Notoatmodjo and Clark Thomborson. Passwords and perceptions. In *Proceedings of the Seventh Australasian Conference on Information Security-Volume 98*, pages 71–78. Citeseer, 2009.
- [117] Ofcom. Number spoofing scams, 2022. <https://www.ofcom.org.uk/phones-telecoms-and-internet/advice-for-consumers/scams/phone-spoof-scam>.
- [118] US Government Accountability Office, General Government Division, and United States of America. Counterfeit us currency abroad: Issues and us deterrence efforts. 1996.
- [119] Aleksandr Ometov, Sergey Bezzateev, Niko Mäkitalo, Sergey Andreev, Tommi Mikkonen, and Yevgeni Koucheryavy. Multi-factor authentication: A survey. *Cryptography*, 2(1):1, 2018.
- [120] Michel Owayjan, Amer Dergham, Gerges Haber, Nidal Fakih, Ahmad Hamoush, and Elie Abdo. Face recognition security system. In *New trends in networking, computing, E-learning, systems sciences, and engineering*, pages 343–348. Springer, 2015.

- [121] Ioannis Papavasileiou, Zhi Qiao, Chenyu Zhang, Wenlong Zhang, Jinbo Bi, and Song Han. Gaitcode: Gait-based continuous authentication using multimodal learning and wearable sensors. *Smart Health*, 19:100162, 2021.
- [122] Ravikanth Pappu, Ben Recht, Jason Taylor, and Neil Gershenfeld. Physical one-way functions. *Science*, 297(5589):2026–2030, 2002.
- [123] Sujan TV Parthasaradhi, Reza Derakhshani, Larry A Hornak, and Stephanie AC Schuckers. Time-series detection of perspiration as a liveness test in fingerprint devices. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, 35(3):335–343, 2005.
- [124] The Paypers. Payment methods report 2019, 2109. URL <https://www.europeanpaymentscouncil.eu/sites/default/files/inline-files/Payment%20Methods%20Report%202019%20-%20Innovations%20in%20the%20Way%20We%20Pay.pdf>.
- [125] Radia Perlman, Charlie Kaufman, and Mike Speciner. *Network security: private communication in a public world*. Pearson Education India, 2016.
- [126] Mohammad Peyravian and Nevenko Zunic. Methods for protecting password transmission. *Computers & Security*, 19(5):466–469, 2000.
- [127] Emma L Prime and David H Solomon. Australia’s plastic banknotes: fighting counterfeit currency. *Angewandte Chemie International Edition*, 49(22):3726–3736, 2010.
- [128] Mahmood Azhar Qureshi and Arslan Munir. Puf-ipa: A puf-based identity preserving protocol for internet of things authentication. In *2020 IEEE 17th annual consumer communications & networking conference (CCNC)*, pages 1–7. IEEE, 2020.
- [129] Bradley Reaves, Logan Blue, and Patrick Traynor. Authloop: End-to-end cryptographic authentication for telephony over voice channels. In *25th USENIX Security Symposium*, pages 963–978, 2016.
- [130] Bradley Reaves, Logan Blue, Hadi Abdullah, Luis Vargas, Patrick Traynor, and Thomas Shrimpton. Authenticall: Efficient identity and content authentication for phone calls. In *26th USENIX Security Symposium*, pages 575–592, 2017.

- [131] OSHA Regulations and Regulatory Guidance. Code of federal regulations. *Respiratory Protection*, 1910.
- [132] Research and Markets. *Global VoIP Services Market 2021-2026*. URL https://www.researchandmarkets.com/reports/5457497/global-voip-services-market-2021-2026?utm_source=BW&utm_medium=PressRelease&utm_code=swfs6q&utm_campaign=1607519+-+Global+VoIP+Services+Market+Report+2021%3a+Market+Should+Grow+from+%2485.2+Billion+in+2021+to+%24102.5+Billion+by+2026+&utm_exec=chdo54prd.
- [133] Kenneth Revett, Florin Gorunescu, Marina Gorunescu, Marius Ene, Sergio Magalhaes, and Henrique Santos. A machine learning approach to keystroke dynamics based user authentication. *International Journal of Electronic Security and Digital Forensics*, 1(1):55–70, 2007.
- [134] V Rusanov, K Chakarova, H Winkler, and AX Trautwein. Mössbauer and x-ray fluorescence measurements of authentic and counterfeited banknote pigments. *Dyes and Pigments*, 81(3):254–258, 2009.
- [135] Merve Sahin, Aurélien Francillon, Payas Gupta, and Mustaque Ahamad. Sok: Fraud in telephony networks. In *2017 IEEE European Symposium on Security and Privacy (EuroS&P)*, pages 235–250. IEEE, 2017.
- [136] Manjula Sandirigama, Akihiro Shimizu, and Matu-Tarow Noda. Simple and secure password authentication protocol (sas). *IEICE Transactions on Communications*, 83(6):1363–1365, 2000.
- [137] Sriram Sankaran, S Shivshankar, and K Nimmy. Lhpuf: Lightweight hybrid puf for enhanced security in internet of things. In *2018 IEEE International Symposium on Smart Electronic Systems (iSES)(Formerly iNiS)*, pages 275–278. IEEE, 2018.
- [138] Alireza Shamsoshoara, Ashwija Korenda, Fatemeh Afghah, and Sherali Zeadally. A survey on physical unclonable function (puf)-based security solutions for internet of things. *Computer Networks*, 183:107593, 2020.
- [139] Ashlesh Sharma, Lakshminarayanan Subramanian, and Eric A Brewer. Paperspeckle: microscopic fingerprinting of paper. In *Proceedings of the 18th ACM conference on Computer and communications security*, pages 99–110. ACM, 2011.

- [140] Imani Sherman, Daniel A Delgado, Juan E Gilbert, Jaime Ruiz, and Patrick Traynor. Characterizing user comprehension in the stir/shaken anti-robocall standard. In *49th Research Conference on Communication, Information and Internet Policy*, 2021.
- [141] Shivang Shukla and Sourabh Dave. Comparison of face recognition algorithms & its subsequent impact on side face. In *2016 International Conference on ICT in Business Industry & Government (ICTBIG)*, pages 1–8. IEEE, 2016.
- [142] David Solomon and Tom Spurling. *The Plastic Banknote: from concept to reality*. CSIRO PUBLISHING, 2014.
- [143] Boyeon Song. *RFID authentication protocols using symmetric cryptography*. PhD thesis, PhD thesis, December, 2009.
- [144] Emily Sonnex, Matthew J Almond, John V Baum, and John W Bond. Identification of forged bank of england£ 20 banknotes using ir spectroscopy. *Spectrochimica Acta Part A: Molecular and Biomolecular Spectroscopy*, 118:1158–1163, 2014.
- [145] Ctirad Sousedik and Christoph Busch. Presentation attack detection methods for fingerprint recognition systems: a survey. *Iet Biometrics*, 3 (4):219–233, 2014.
- [146] Douglas R Stinson. *Cryptography: theory and practice*. Chapman and Hall/CRC, 2005.
- [147] Ying Su, Jeremy Holleman, and Brian Otis. A 1.6 pj/bit 96% stable chip-id generating circuit using process variations. In *2007 IEEE International Solid-State Circuits Conference. Digest of Technical Papers*, pages 406–611. IEEE, 2007.
- [148] Wayne C Summers and Edward Bosworth. Password policy: the good, the bad, and the ugly. In *Proceedings of the winter international symposium on Information and communication technologies*, pages 1–6, 2004.
- [149] Noriko Takemura, Yasushi Makihara, Daigo Muramatsu, Tomio Echigo, and Yasushi Yagi. On input/output architectures for convolutional neural network-based cross-view gait recognition. *IEEE Transactions on Circuits and Systems for Video Technology*, 29(9):2708–2719, 2017.

- [150] Ruben Tolosana, Ruben Vera-Rodriguez, Julian Fierrez, and Javier Ortega-Garcia. Exploring recurrent neural networks for on-line handwritten signature biometrics. *Ieee Access*, 6:5128–5138, 2018.
- [151] Ehsan Toreini, Siamak F Shahandashti, and Feng Hao. Texture to the rescue: practical paper fingerprinting based on texture patterns. *ACM Transactions on Privacy and Security (TOPS)*, 20(3):9, 2017.
- [152] TransNexus. Service provider sti fee changes for 2021, 2021. <https://transnexus.com/blog/2021/sti-provider-rate-changes/>.
- [153] TransNexus. Reply comments on shaken extensions and effectiveness, 2022. <https://transnexus.com/blog/2022/shaken-effectiveness-extensions-reply-comments/>.
- [154] TransNexus. Robocalls up sharply in october, 2022. <https://transnexus.com/blog/2022/robocalls-up-sharply-october/>.
- [155] TransNexus. Stir/shaken statistics from october 2022, 2022. <https://transnexus.com/blog/2022/shaken-statistics-october/>.
- [156] Issa Traore. *Continuous Authentication Using Biometrics: Data, Models, and Metrics: Data, Models, and Metrics*. Igi Global, 2011.
- [157] Truecaller. 2022 U.S. Spam & Scam Report, 2022. <https://truecaller.blog/2022/05/24/truecaller-insights-2022-us-spam-scam-report/>.
- [158] Yuh-Min Tseng, Jinn-Ke Jan, and Hung-Yu Chien. On the security of methods for protecting password transmission. *Informatica*, 12(3): 469–476, 2001.
- [159] Takasuke Tsuji, Takashi Kamioka, and Akihiro Shimizu. Simple and secure password authentication protocol, ver. 2 (sas-2). In *ITE Technical Report 26.61*, pages 7–11. The Institute of Image Information and Television Engineers, 2002.
- [160] Huahong Tu, Adam Doupé, Ziming Zhao, and Gail-Joon Ahn. Sok: Everyone hates robocalls: A survey of techniques against telephone spam. In *2016 IEEE Symposium on Security and Privacy (SP)*, pages 320–338, 2016. doi: 10.1109/SP.2016.27.

- [161] Huahong Tu, Adam Doupé, Ziming Zhao, and Gail-Joon Ahn. Toward standardization of authenticated caller id transmission. *IEEE Communications Standards Magazine*, 1(3):30–36, 2017.
- [162] Pim Tuyls, Geert-Jan Schrijen, Boris Škorić, Jan van Geloven, Nynke Verhaegh, and Rob Wolters. Read-proof hardware from protective coatings. In *International Workshop on Cryptographic Hardware and Embedded Systems*, pages 369–383. Springer, 2006.
- [163] Pim Tuyls, Boris Škoric, and Tom Kevenaar. *Security with noisy data: on private biometrics, secure key storage and anti-counterfeiting*. Springer Science & Business Media, 2007.
- [164] Ravitej Uppu, Tom AW Wolterink, Sebastianus A Goorden, Bin Chen, Boris Škorić, Allard P Mosk, and Pepijn WH Pinkse. Asymmetric cryptography with physical unclonable keys. *Quantum Science and Technology*, 4(4):045011, 2019.
- [165] John G Van Bosse and Fabrizio U Devetak. *Signaling in telecommunication networks*, volume 87. John Wiley & Sons, 2006.
- [166] Anna Vila, N Ferrer, J Mantecon, D Breton, and JF Garcia. Development of a fast and non-destructive procedure for characterizing and distinguishing original and fake euro notes. *Analytica Chimica Acta*, 559(2):257–263, 2006.
- [167] Serge Vrijaldenhoven et al. Acoustical physical uncloneable functions. *Philips internal publication PR-TN-2004-300300*, 2004.
- [168] Shen Wang, Ehsan Toreini, and Feng Hao. Anti-Counterfeiting for Polymer Banknotes Based on Polymer Substrate Fingerprinting (datasets archive part 1/2), March 2021. URL <https://doi.org/10.1109/TIFS.2021.3067440>.
- [169] Shen Wang, Ehsan Toreini, and Feng Hao. Anti-Counterfeiting for Polymer Banknotes Based on Polymer Substrate Fingerprinting (datasets archive part 2/2), March 2021. URL <https://doi.org/10.1109/TIFS.2021.3067>.
- [170] Shen Wang, Ehsan Toreini, and Feng Hao. Anti-counterfeiting for polymer banknotes based on polymer substrate fingerprinting. *IEEE Transactions on Information Forensics and Security*, 16:2823–2835, 2021.

- [171] Xinyuan Wang and Ruishan Zhang. Voip security: Vulnerabilities, exploits, and defenses. In *Advances in Computers*, volume 81, pages 1–49. Elsevier, 2011.
- [172] David Wolman. *The end of money: Counterfeiters, preachers, techies, dreamers—and the coming cashless society*. Hachette UK, 2013.
- [173] Wen-Her Yang and Shiuh-Pyng Shieh. Password authentication schemes with smart cards. *Computers & Security*, 18(8):727–733, 1999.
- [174] Chi-Yuan Yeh, Wen-Pin Su, and Shie-Jue Lee. Employing multiple-kernel support vector machines for counterfeit banknote recognition. *Applied Soft Computing*, 11(1):1439–1447, 2011.
- [175] Xiuyan Yin and Satish Kumar. Flow visualization of the liquid emptying process in scaled-up gravure grooves and cells. *Chemical engineering science*, 61(4):1146–1156, 2006.
- [176] YouGov. Don’t call me: Nearly 90% of customers won’t answer the phone anymore, 2019. <https://martech.org/dont-call-me-nearly-90-of-customers-wont-answer-the-phone-anymore-study/>.
- [177] James Yu. An analysis of applying stir/shaken to prevent robocalls. In *Advances in Security, Networks, and Internet of Things*, pages 277–290. Springer, 2021.
- [178] Naser Zaeri. Minutiae-based fingerprint extraction and recognition. *Biometrics*, 2011.
- [179] Moshe Zviran and Zippy Erlich. Identification and authentication: technology and implementation issues. *Communications of the Association for Information Systems*, 17(1):4, 2006.