Doctoral theses at NTNU, 2024:226

Pia Bauspieß

Post-Quantum Secure Biometric Systems

Doctoral thesis

NTNU
Norwegian University of Science and Technology
Thesis for the Degree of Philosophiae Doctor
Faculty of Information Technology and Electrical Engineering
Dept. of Information Security and Communication Technology

NTNU
Norwegian University of Science and Technology

Pia Bauspieß

Post-Quantum Secure Biometric Systems

Thesis for the Degree of Philosophiae Doctor

Trondheim, June 2024

Norwegian University of Science and Technology
Faculty of Information Technology and Electrical Engineering
Dept. of Information Security and Communication Technology

NTNU
Norwegian University of
Science and Technology

NTNU
Norwegian University of Science and Technology

Thesis for the Degree of Philosophiae Doctor

Faculty of Information Technology and Electrical Engineering
Dept. of Information Security and Communication Technology

ISBN 978-82-326-8044-3 (printed ver.)
ISBN 978-82-326-8043-6 (electronic ver.)
ISSN 1503-8181 (printed ver.)
ISSN 2703-8084 (online ver.)

Doctoral theses at NTNU, 2024:226

Printed by NTNU Grafisk senter

Acknowledgements

I would like to express my immense gratitude and appreciation to my supervisors Anamaria Costache and Patrick Bours. Your experience, valuable advice, kind words, and encouragement have guided me throughout this thesis. I was fortunate to learn so much from you, both professionally and personally, and I will take what I have learned from you with me for my future career and life. Thank you for always having my back, and for always believing in me. I feel incredibly lucky that I was able to work with you.

I am also extremely grateful to Katrien De Moor, Thomas Zinner, Mona Nordaune, Jascha Kolberg, Tjerand Silde, and Eirin Bar for their incredible support and encouragement. Thank you for validating my experiences, for listening, and for being by my side in difficult situations. You have made an immeasurable difference to my experience during these last years and months.

I would like to thank my co-authors and students for the collaborations that have resulted in publications included in this thesis. Through the collaborations and discussions with you, I was able to dive deeply into the research topics that I am passionate about. I would also like to thank the ATHENE Research Center and Hochschule Darmstadt for the material support of my PhD.

I would like to thank the PhD committee, Iluminada Baturone, Hilder Vitor Lima Pereira, and Katrien De Moor, for their willingness to evaluate this thesis.

Thank you to everyone who offered their time and expertise to proofread my thesis or made helpful suggestions.

To my partner, my family, my Norwegian family, and my friends: You are and have been my rocks, and I would not have been able to begin or complete this journey without you. Thank you for your love, your patience, and your unconditional support. You mean the world to me.

Abstract

Biometrics, which is defined as the automated recognition of individuals based on their behavioural and biological characteristics, can be applied to increase the trust and usability of digital interactions. Due to their uniqueness and persistence, biometric characteristics allow for a reliable connection between an individual and their digital identity.

However, these same properties of uniqueness and persistence also give rise to privacy concerns. Therefore, this thesis investigates the cryptographic protection of biometric systems. If such protection is built from classical cryptography, it has two main shortcomings: firstly, it is limited in the type of computations that can be performed on the encrypted data, and secondly, it does not provide long-term protection against threats posed by quantum computers.

Fully homomorphic encryption can mitigate both of the aforementioned concerns. Based on hard lattice problems, it falls into a new category of mathematical problems that are believed to withstand known quantum computing attacks, called post-quantum cryptography. However, its practical efficiency remains an open challenge. Therefore, this thesis studies the efficiency of biometric systems under fully homomorphic encryption.

In addition, this thesis addresses how biometric characteristics can be used to facilitate cryptographic key exchange, where a shared key for encrypted communication between a client and a server is computed correctly if and only if the biometric verification was successful. As with biometric comparisons in the encrypted domain, the security of such schemes against quantum computing threats needs to be considered in order to achieve a lasting protection of the sensitive biometric information.

Finally, the security of biometric information protection against malicious adversaries, which can deviate arbitrarily from a given protocol, and its impact on the computational efficiency are investigated in this thesis.

Sammendrag på norsk

Biometri, som er definert som automatisert gjenkjennelse av individer basert på deres adferdsmessige og biologiske egenskaper, kan brukes for å øke tilliten og brukervennligheten i digitale interaksjoner. På grunn av deres unikhet og robusthet, sørger biometriske egenskaper for en pålitelig forbindelse mellom et individ og deres digitale identitet.

Imidlertid gir de samme egenskapene av unikhet og robusthet også opphav til personvernbekymringer. Derfor undersøker denne avhandlingen den kryptografiske beskyttelsen av biometriske systemer. Dersom slik beskyttelse er bygget fra klassisk kryptografi, har den to hovedmangler: for det første er den begrenset av mulige beregninger som kan utføres på de krypterte dataene, og for det andre gir den ikke langsiktig beskyttelse mot trusler fra kvantedatamaskiner.

Fullstendig homomorfisk kryptering kan løse begge de nevnte problemene. Basert på vanskelige problemer fra gitre, faller den inn i en ny kategori av matematiske problemer som antas å motstå kjente angrep fra kvantedatamaskiner, kalt kvantesikker kryptografi. Den praktiske effektiviteten er imidlertid en utfordring. Derfor studerer denne avhandlingen effektiviteten til biometriske systemer under fullstendig homomorfisk kryptering.

I tillegg tar denne avhandlingen for seg hvordan biometriske egenskaper kan benyttes til å forenkle kryptografisk nøkkelutveksling, der en delt nøkkel for kryptert kommunikasjon mellom en klient og en server beregnes riktig hvis og bare hvis den biometriske verifiseringen var vellykket. Som med andre biometriske sammenligninger i det krypterte domenet, må sikkerheten til slike protokoller beskyttes mot kvantedatamaskiner for å oppnå en varig beskyttelse av den sensitive biometriske informasjonen.

Til slutt undersøkes sikkerheten til biometrisk informasjonsbeskyttelse mot ondsinnede motstandere, som kan avvike vilkårlig fra en gitt protokoll, og dets innvirkning på beregningseffektiviteten til slike protokoller.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Biometrics, or *biometric recognition*, is defined as the automated recognition of individuals based on their behavioral and biological characteristics [152]. Examples of biological biometric characteristics that can be used for automated recognition are facial images or the ridge line patterns in fingerprints. Examples of behavioral biometric characteristics include keyboard timing patterns, hand-written signatures, or speaker recognition. From such characteristics, distinguishing and repeatable *biometric features* can be extracted and used for comparison, determining the success or failure of the biometric recognition transaction.

Biometric recognition can be used in different modes, three of which are relevant to this thesis: verification, identification, and key establishment. In a *verification* transaction, a biometric probe is compared against a previously stored biometric reference using a one-to-one comparison of the probe and reference features. The outcome of this comparison is a comparison score, which can be compared against a predefined threshold. The verification transaction produces a binary output depending on the threshold comparison, resulting in either an accept or a reject decision for the biometric probe in question. If a similarity measure is used, comparison scores above the threshold are accepted. If a dissimilarity measure such as a distance function is used, comparison scores below the threshold are accepted. Biometric verification has become an established part of our digital infrastructure, from smartphone access to automated border control, where a stored reference, e.g., a fingerprint pattern on a smartphone or a face image in a passport, is compared against a freshly captured probe.

In an *identification* transaction, a biometric probe is compared against a refer-

ence database and a one-to-many comparison is performed. All references that yield comparison scores above the given threshold (for a similarity measure) are collected in a candidate list. Depending on the decision policy, only the reference identifier corresponding to the best candidate (i.e., the reference with the highest comparison score based on a similarity measure) is returned. Alternatively, more candidates may be returned. If no candidate was found within the reference database, the identification transaction was unsuccessful. Biometric identification can be used to prevent duplicate issuance of unique citizenship or tax identifiers [256], or in law enforcement to perform checks against databases of previously convicted criminals.

The third mode of biometric recognition used in this thesis is *biometrics-based key establishment*. This application goes beyond the binary outcome of an accept or reject decision by deriving cryptographic keys from biometric characteristics which can be used for encrypted communication between a client and a server. In this transaction, the server holds a key derived from a biometric reference, and the client generates a key from a biometric probe. The shared key between the client and the server should be the same if and only if the verification for the probe resulted in an accept decision, and different if the probe is rejected. Cryptographic keys generated in this way can be used to establish secure channels for messaging applications and further applications relying on encrypted communication [260].

In all three transactions described above, the uniqueness and persistence of biometric characteristics enable a user-friendly and reliable way of connecting a human user to their digital identity. Compared to passwords or cryptographic tokens, biometric authentication therefore provides an additional level of trust and usability to the service it is applied to [237].

However, these same properties of uniqueness and persistence have also raised security and privacy concerns [154]. If biometric features, e.g., a fingerprint pattern, are stolen, they can be used to conduct impersonation attacks on the data subject in question. In addition, the information may be used to link the subject across different applications or derive personal information from the obtained biometric sample such as the subject's ethnic origin or sensitive medical information. The importance of the protection of such personal identifiable information has been explicitly recognized by the European Union's *General Data Protection Regulation* (GDPR) [109], which has increased awareness and regulatory requirements around digital authentication, including biometric recognition.

Contrary to knowledge-based authentication mechanisms, biometric characteristics cannot be revoked or easily replaced following such an attack. It is there-

fore important to ensure the protection of biometric data used for automated recognition throughout different applications, particularly whenever biometric data is stored remotely. As biometric characteristics allow for an accurate identification of individuals over several decades [166], this protection should be designed with long-term security against potential future attacks in mind.

To ensure the security of biometric information against the threats discussed above, the international ISO/IEC 24745 [149] standard on *Biometric Information Protection* (BIP) gives three requirements: the unlinkability, renewability, and irreversibility of protected *biometric templates*, i.e., extracted feature representations. Therefore, the term *Biometric Template Protection* (BTP) can be used interchangeably with BIP. *Unlinkability* refers to the cross-application linkage attack discussed above. It requires that protected templates created from the same source cannot be linked by an attacker, implying that the space of possible protected representations for a single template needs to be large and provide indistinguishability. *Renewability* addresses the revocation of protected templates, requiring that new protected templates can be created from the same source such that they cannot be tied to stolen biometric information. *Irreversibility* ensures the confidentiality and privacy of the biometric data, as an attacker should not be able to reconstruct original biometric samples given only the protected templates. Unprotected templates, although an abstract representation of the original biometric data, have been shown to be vulnerable against such sample reconstruction attacks [59, 119, 185].

One approach to fulfil all three requirements is computation on encrypted data. If the biometric templates stored and processed within each transaction are encrypted, an attacker cannot recover them, given that the applied encryption scheme achieves confidentiality. Encryption schemes that achieve indistinguishably against chosen-plaintext attacks, or IND-CPA security, ensure that encrypted biometric templates cannot be linked across applications. This is due to the fact that IND-CPA security means that an attacker cannot efficiently distinguish between an encryption of the bit zero and an encryption of the bit one. Extending this property to biometric templates, an attacker is not able to efficiently distinguish an encryption of a given template from an encryption of a different template. In addition, cryptographic schemes with non-deterministic properties do not allow an attacker to efficiently distinguish an encryption of a given template from a second encryption of the same template.

However, the challenge of computing on encrypted data is the meaningful evaluation of biometric comparison functions on the encrypted templates, which cannot be achieved with arbitrary encryption schemes. Instead, a subcategory of encryption schemes with *homomorphic* properties are required, which allow for operations on ciphertexts that translate directly to operations on the under-

lying plaintext. Informally speaking, a homomorphism describes a structure-preserving map between two algebraic structures of the same type, e.g., between two sets. In the case of encryption, arithmetic operations such as additions and multiplications can be preserved between the plaintext space and the ciphertext space of an encryption scheme, and such schemes are referred to as *Homomorphic Encryption* (HE) schemes. For example, a homomorphic multiplication allows for the multiplication of two ciphertexts that results in the product of the two underlying plaintexts after decryption. More generally, a homomorphic evaluation of a given function describes its computation on ciphertexts such that the value of the function is returned in a plaintext after decryption. Thereby, biometric comparison functions can be evaluated homomorphically without revealing private information to the party performing the computation. As the outcome, only the comparison score is revealed after decryption, from which an accept or reject decision over the comparison trial can be derived.

Homomorphic properties of public-key encryption schemes have been studied since the seminal work by Rivest, Shamir, and Adleman [224], whose RSA encryption scheme allows for homomorphic multiplications of ciphertexts [223]. Later, HE schemes such as the ElGamal [102] and Pailler [203] schemes were applied to BIP specifically [26, 126, 233], allowing for a homomorphic evaluation of distance metrics such as the Euclidean distance that can be used for biometric comparisons.

The aforementioned HE schemes are public-key encryption schemes that are based on the security of the factorization problem [224] or the discrete logarithm problem [87]. As such, they have two significant shortcomings that make them inapplicable for the scope of this thesis: Firstly, they are limited in the type and number of homomorphic operations they support, and secondly, they do not provide long-term security in the face of quantum computing.

Even though no quantum computer of sufficient size to break current encryption has been publicly announced, the European Union estimates that practical attacks could be expected as soon as 2035 [107]. The most important threat posed by quantum computing that impacts public-key encryption is Shor's algorithm [238]. Proposed in the 1994, this algorithm allows to solve both the factorization and discrete logarithm problems in polynomial-time, rendering them unusable for cryptography. In other words, Shor's algorithm implemented on a quantum computer allows for an efficient computation of the decryption function without knowledge of the cryptographic secret key. The encrypted data, e.g., biometric templates, can thereby be accessed by an attacker and further exploited for personal information or impersonation attacks. Given the persistence of biometric characteristics [166], security against such attacks needs to be considered today. As a viable attack, an attacker could intercept protected

biometric templates today, and reverse them decades later using a quantum computer. This attack can be considered feasible, as quantum computing may become more attainable and inexpensive in the future.

In addition to Shor's algorithm, a second algorithm threatens symmetric cryptography when implemented on a quantum computer, namely Grover's algorithm [132]. Published in 1996, this algorithm allows for a square-root speed-up on unstructured search problems and impacts the security of hash functions and symmetric encryption schemes [77]. Contrary to Shor's algorithm however, the square-root search time improvement provided by Grover's algorithm can be mitigated by doubling the security parameters of symmetric schemes, which still allows for an efficient computation of their encryption and decryption functions. Therefore, symmetric cryptography is not threatened by quantum computing to the same degree as asymmetric cryptography.

Given the threats posed by Shor's and Grover's algorithm, the most recent version of ISO/IEC 24745 also emphasizes the long-term protection of biometric data against attacks using quantum computers. For BIP schemes relying on asymmetric cryptography, this security can be achieved through *Post-Quantum Cryptography* (PQC), a term established for public-key cryptography that is believed to withstand known attacks implemented on quantum computers [238]. In comparison, cryptographic schemes that base their security on the factorization or discrete logarithm problems are referred to as classically secure. The United States *National Institute of Standards and Technology* (NIST) has shepherded a PQC standardization competition which investigates suitable schemes and encourages cryptanalysis on the latter [10]. Two categories of cryptographic schemes have emerged as the most promising within this competition: lattice-based [204] and code-based [262] cryptography.

Out of these two categories, lattice-based cryptography is particularly interesting and relevant for this thesis, as it provides post-quantum security and homomorphic properties that exceed the capabilities of classically secure HE. *Lattices*, which can informally be defined as discrete subgroups of real vector spaces, allow for the construction of hard mathematical problems that are assumed to withstand attacks implemented on quantum computers. Hard problems based on lattices have been studied since the 1990s [140]. In addition to their believed post-quantum security, lattices enable cryptographic constructions that are not known to be feasible with classical cryptography. A challenge that persisted with classically secure HE schemes was the combination of different arithmetic operations such as additions and multiplications, and the number of consecutive multiplications on a single ciphertexts, i.e., the multiplicative depth of an arithmetic circuit. Only in 2009, the breakthrough construction of *Fully Homomorphic Encryption* (FHE) by Gentry [121] paved the way for the evaluation of

arbitrary arithmetic functions over encrypted data.

Since Gentry's initial construction of FHE, a number of improved schemes with different properties and foci have been developed [50, 51, 52, 65, 69, 98, 112]. As their computation efficiency has improved over recent years, their application to privacy-preserving computation in fields such as BIP has received increasing interest [44, 103, 104, 169, 171, 243, 270]. More recently, the homomorphic properties of key establishment algorithms analyzed in the NIST PQC competition have also been successfully applied to BIP [19, 226]. Even though these algorithms have not been designed with homomorphic encryption in mind, they inherit limited homomorphic properties from their mathematical constructions based on lattices or error-correcting codes which are sufficient to evaluate simple distance metrics between biometric templates. Closely related to homomorphic encryption, further techniques that allow for computation on encrypted data such as *Secure Multiparty Computation* (MPC) and *Functional Encryption* (FE) have also been applied to BIP instantiated with lattices-based primitives [30, 66].

Aside from privacy-preserving biometric comparisons, key generation based on biometric authentication has received increasing research interest [48, 106, 260]. The advantages of biometrics-based key establishment compared to password-based protocols follow from the motivation for biometric authentication, yielding a more secure and user-friendly digital infrastructure. As lattice- and code-based schemes are used in these works, they also contribute to the development of post-quantum secure biometric systems.

While many advances have been made in the evolving research field of BIP, many critical challenges remain, which can be categorized into challenges regarding the security and challenges regarding the efficiency of BIP schemes. In terms of security, the need for post-quantum protection of biometric data is motivated through the threat posed by Shor's algorithm [238]. Additionally, the majority of BIP schemes with post-quantum security have only been constructed under the semi-honest adversary model, where all parties are assumed to follow the given protocol. However, this assumption cannot be considered a realistic model of real-world adversaries, which may behave maliciously [28].

In terms of the computational efficiency, the compatibility of biometric feature representations with privacy-preserving computation techniques is a critical component. Captures of biometric characteristics are inherently noisy. For example, biometric features may change through ageing processes [131, 166], changing the appearance of the same individual to a certain degree. Even if the same features are captured within short time intervals, the feature extraction process will result in slightly different mathematical representations of the same instance.

Therefore, biometric features cannot be compared using absolute equity, but require a similarity measure that is tolerant to noisy representations. Historically, these representations and similarity measures have been developed to obtain the highest recognition accuracy for the targeted biometric modality, but not necessarily with cryptographic protection in mind. The encoding of biometric features into plaintext spaces of cryptographic schemes can therefore be non-trivial [154].

In addition to the challenges for single biometric modalities, combinations of different biometric modalities and their feature representations require further considerations. This approach is referred to as *multibiometrics* and can be applied to increase the trust in the authentication through requiring the authentication of multiple biometric instances of the same individual. Different approaches have been explored for multibiometric in the encrypted domain which address the challenges of consolidating different feature representations [243].

Finally, improving the efficiency of large-scale biometric identification under encryption remains an open problem [154]. As a one-to-many search against a potentially large reference database is computed in a biometric identification transaction, the computational workload of this transaction increases linearly with the size of the database. Given the cost of cryptographic operations for a single comparison discussed above, this can render identification searches in the encrypted domain infeasible [91].

Two main approaches have been explored to facilitate a workload reduction for identification transactions on large biometric databases: feature transformation and preselection [92]. Altering the representation of extracted biometric features, or *feature transformation*, can be applied to reduce the size or form of the biometric feature representation in a way that reduces the cost of a single biometric comparison, and thereby decreases the overall cost of the database search [92]. This approach has been explored for FHE-encrypted databases with efficient execution times for mid-sized galleries of 1 to 5 million subjects [31, 103]. However, the feature transformation approach continues to scale linearly with the database size.

Therefore, *preselection*, i.e., selecting a subset of the reference database, may be applied to reduce the search space and remove biometric comparisons that can be considered unlikely to result in a positive identification outcome [92]. In terms of the application of FHE, preselection can be applied on top of an encrypted database, aiming to produce a subset of encrypted feature vectors to be considered for the full and expensive homomorphic evaluation of the comparison function. However, any auxiliary information required for preselection needs to be evaluated for privacy-sensitive information leakage, and pro-

tected accordingly [33]. Notably, feature transformation and preselection are not mutually exclusive. Indeed, some feature transformation approaches for FHE-encrypted databases have been shown to be compatible with preselection [33], while others apply an encoding of the biometric feature representation that does not trivially allow for a meaningful preselection [103].

1.1 Research Questions

The first research question targets one of the broadest challenges within BIP, namely the compatibility of biometric feature representations with privacy-preserving computation techniques. This first research question also includes a comparison between different privacy-preserving computation techniques for BIP, where the focus of this thesis lies on solutions with post-quantum security. For such solutions, the trade-off between security and efficiency is investigated.

In some cases, the biometric feature representation aligns with the cryptographic scheme. This is for example true for fixed-length binary representations that be compared using a simple distance metric such as the Hamming distance, as many cryptographic schemes allow for binary plaintext representations [170]. With the rise of deep-learning based feature extraction [191] however, biometric features from different modalities can be represented as real-valued vectors of a fixed dimension and may require quantisation to be compatible with cryptographic schemes that operate on integer or binary plaintext spaces [96]. The impact of such changes must be evaluated with regard to the biometric performance.

A particular challenge arises in the case where the biometric feature representation is not expressed in fixed-length vectors, but unordered sets of variable cardinality. This is the case for minutiae-based fingerprint representations, where *minutiae* are significant points in the fingerprint ridge line pattern given through the ridges and valleys in the skin that can be captured using ink and paper, or by different digital fingerprint capture devices [151]. The captured location and number of such minutiae can vary between captures even for the same biometric instance, i.e., a finger. Comparison algorithms for such variable-length feature representations are more complex than distance metrics evaluated on fixed-length vectors, which impacts the computational workload. The feature representation and its encoding have been shown to have a significant impact on the computational workload, and remain one of the main challenges for secure, efficient, and accurate biometric systems [103, 104, 202].

Research Question 1

Which privacy-preserving computation techniques are best suited for biometric information protection?

- How do different approaches to biometric information protection compare in terms of their security and efficiency?
- How can biometric features be represented to aid different encoding mechanisms used in privacy-preserving computation techniques?

The second research question builds upon the first with a focus on the efficiency of biometric identification schemes, where a large number of biometric comparisons are computed within the encrypted domain. In particular, we consider BIP schemes where the biometric references stored in the database are encrypted using FHE. In terms of workload reduction approaches, the compatibility of feature transformation and preselection approaches is an important factor to achieve efficiency. At the same time, the security of the overall transaction with regard to the ISO/IEC 24745 requirements as well as security against quantum adversaries should not be impaired by the application of workload reduction. This is particularly important for preselection approaches using auxiliary indexing data, which require the same level of protection as the reference database.

Research Question 2

How can computational workload reduction be applied to improve the efficiency of FHE-based biometric identification systems?

- How can computational workload reduction for biometric identification be applied in the homomorphically encrypted domain?
- How can the trade-off between computational workload reduction and efficient encryption be optimized?

The remaining two research questions focus on the security of BIP schemes. In particular, the security against semi-honest and malicious adversaries is investigated within the third research question. Semi-honest adversaries are assumed to adhere to the given protocol, which can only be considered a realistic behaviour within a controlled environment. In contrast to semi-honest adversaries, malicious adversaries may deviate arbitrarily from a given protocol, which can be considered a significantly more realistic and challenging scenario. We discuss this further in Section 4.2.

In addition to these cryptographic security models, the ISO/IEC 30136 [148] on the performance testing of BIP schemes has defined the *full disclosure model*. In this model, an adversary is assumed to have access to all secrets used within the protocol, which includes cryptographic secret keys. For FHE-based BIP schemes, this implies that the adversary is assumed to have gained access to the FHE secret key used for the encryption of the reference database. Having gained access to this key, the reference database could be decrypted, and the adversary would obtain the unprotected biometric templates. Notably, this is an attack scenario that is not covered by the definition of malicious security, as an adversary with access to the secret key of a cryptographic scheme is typi-

cally considered to win the security game trivially [165]. In particular, malicious security does not imply security under the full disclosure model.

Even though the full disclosure model does not align with standard cryptographic assumptions, it has recently gained increasing interest within the BIP research community, and is therefore investigated in this research question in addition to the established semi-honest and malicious adversary models. Even though the focus of this research question is on security, the computational workload remains important, as is has been shown that security against malicious adversaries can render post-quantum secure cryptographic primitives infeasible [12, 71].

Research Question 3

Can biometric information protection based on homomorphic encryption be secured against malicious adversaries in a feasible manner?

- Is it possible to secure biometric systems under the full disclosure model defined in ISO/IEC 30136 using only homomorphic encryption?
- Is it possible to efficiently secure biometric systems against malicious adversaries?

Finally, this thesis focuses on post-quantum secure BIP schemes. Therefore, both semi-honest and malicious adversaries with access to quantum computers are considered in addition to standardized security models such as the full disclosure model discussed above. The last research question can therefore be considered an overarching goal for the research contributions presented in this thesis.

Research Question 4

How can biometric systems be secured against quantum adversaries?

- Which quantum adversary models need to be considered for biometric systems?
- How can the computational workload of post-quantum biometric systems be optimized?

1.2 Overview of Contributions

In this section, an overview of the contributions to the open research problems described above is given. The contributions are presented based on six publications or manuscripts that are currently in peer review. For each paper, the contribution to the research question is highlighted. An overview over the contributions of the papers to the research questions is given in Table 1.1.

Contribution to Research Questions	RQ1	RQ2	RQ3	RQ4
Paper A [38]	●	○	○	◐
Paper B [36]	◐	●	○	●
Paper C [34]	◐	○	●	◐
Paper D [32]	●	○	○	◐
Paper E [37]	●	○	●	◐
Paper F [35]	●	●	●	●

Table 1.1: Overview of the contributions of each paper to each research question. Filled circles indicate that the paper contributes to the research question, half-filled circles indicate that the research question is addressed partially in the paper, but is not the focus of the contribution, and empty circles indicate that the paper does not directly contribute to the research question.

Paper A: On the Feasibility of Fully Homomorphic Encryption of Minutiae-Based Fingerprint Representations

Fingerprint recognition has been established as a reliable mode of biometric authentication. It is traditionally based on minutiae, which are significant points in the fingerprint ridge pattern [151]. Related work has previously only studied post-quantum protection for fixed-length fingerprint representations [169], which have shown lower recognition accuracy compared to minutiae-based approaches [153]. However, variable-length fingerprint representations had previously only been combined with classically secure cryptography [20, 125]. The aim of this work was therefore to explore biometric template protection with post-quantum security for variable-length fingerprint representations.

The key challenge of fingerprint recognition is the construction of accurate and efficient comparison functions between two minutiae templates. The difficulty of this challenge increases with the application of FHE, as operations on FHE-

encrypted data are limited. FHE schemes with the best amortized computation times can only handle additions, multiplications, and rotations of vectorized data [51, 52, 65, 112]. FHE schemes with programmable bootstrapping additionally support the efficient homomorphic evaluation of look-up tables, but do not have similarly good amortization, and can only efficiently handle smaller message precision [69]. Despite the challenge of limited efficient homomorphic operations, FHE provides the desired post-quantum protection, as the scheme applied in this work is based on hard lattice problems [65].

In this work, we evaluated the Minutia Cylinder Code (MCC) [58] comparison algorithm on encrypted fingerprint templates. This algorithm was chosen due to its rotation-invariance and the limited number of conditional statements that need to be evaluated during the comparison. For the FHE scheme, Cheon-Kim-Kim-Song (CKKS) [65] was chosen. CKKS allows for operations on fixed-point data, such that the MCC fingerprint comparison algorithm did not need to be altered or quantized significantly, thereby maintaining the biometric performance of the unencrypted system. However, not all operations of the MCC algorithm could be expressed in the encrypted domain. In particular, conditional operations had to be computed on plaintext data, as they cannot be efficiently realized under FHE [147]. Nevertheless, the proposed approach still provides protection to the encrypted fingerprint templates. For example, when the difference between two minutiae angles is decrypted and compared against a threshold, we argue that it is still challenging to derive the single minutiae angles purely based on their difference.

In terms of computational performance, the evaluation revealed that a large number of FHE operations needed to be performed in order to evaluate the MCC comparison algorithm. In our approach, the comparison between each minutiae pair corresponded to the cost of one verification computed on fixed-length fingerprint representations, with some additional operations. The overall number of minutiae comparisons is the product of the number of minutiae of the reference and probe template. On the MCYT [199] database, which the proposed protocol was evaluated on, the median number of minutiae per template was 35, which yields an average number of cylinder comparisons for one verification of $35 \cdot 35 = 1225$. Thereby, one verification transaction took over three hours on commodity hardware.

Notably, this work was the first to evaluate biometric template protection with post-quantum security for variable-length fingerprint representations. The biometric performance of the unencrypted computations were not compromised by the applied FHE scheme, which additionally provided long-term protection through post-quantum security. However, the challenge of evaluating conditional statements on encrypted data and the high complexity of minutiae-based

fingerprint comparison algorithms leads to an unacceptable computational cost for real-world applications. This is a motivation for future work on faster FHE computations, as well as fixed-length fingerprint representations with high accuracy. In this area, deep-learning based fingerprint representations are currently emerging [104], which lend themselves ideally to the problem explored in this work.

Contribution of Paper A

Paper A addresses RQ1 on biometric feature representations and their compatibility with post-quantum secure template protection approaches, in particular FHE. The paper shows a clear trade-off between security and efficiency with regard to variable-length fingerprint representations. Based on the results reported in Paper A, it can be argued that variable-length representations cannot be efficiently compared under FHE at the moment. Related work has shown fixed-length representations achieve high efficiency even under post-quantum secure encryption [169]. However, these representations have been shown to lack high accuracy. With advances in both FHE and fixed-length fingerprint representations, this trade-off could be mitigated in the future. Through the post-quantum security of the applied FHE scheme, Paper A also partly addresses RQ4, showing that security against semi-honest quantum adversaries can be achieved when comparing variable-length feature representations, even though this protection could not be achieved with computational efficiency in this particular approach.

Paper B: HEBI: Homomorphically Encrypted Biometric Indexing

The computational workload of biometric identification in the encrypted domain is often a hindering factor in their real-world application [154]. Therefore, indexing has become a popular approach to preselection for biometric identification [90, 137, 195, 200, 205, 229, 246, 261]. While such approaches have been shown to reduce the computational workload, they also introduce additional data to the biometric transaction, which we refer to as auxiliary indexing data. In previous works, this indexing data had not been encrypted, but it was considered that it did not reveal sensitive information about the underlying data subjects clustered under one index [200].

In this work, we showed that this assumption does not hold true in general, and conducted an experimental analysis on a recently proposed scheme. Our experiments showed that soft-biometric features of the subjects assigned to the same cluster can be reconstructed with high fidelity. Significantly, the analysis showed

that underrepresented ethnicities are particularly vulnerable to such attacks. Our results clearly showed that the encryption of the reference database alone does not suffice to grant complete protection of the stored biometric data.

As a solution, we presented the HEBI protocol, which utilizes *Public-Key Encryption with Keyword Search* (PEKS) [45] to encrypt the auxiliary indexing data. The PEKS scheme applied in our work is based on lattices and can therefore be considered to have post-quantum security [39]. Additionally, the reference database is encrypted with lattice-based FHE [65]. If FHE with plaintexts in fixed-point representation is applied as in our work, the biometric performance is not impaired by adding encryption. This also holds true for the PEKS-encryption of the indexing data, as long as it can be represented as binary vectors.

Using our proposed HEBI protocol, the computational cost of a cluster retrieval based on PEKS was evaluated at 0.12 milliseconds per cluster, while post-quantum security was achieved both for the indexing data as well as the biometric templates. In a relative comparison, the experimental evaluation showed that the cost of protected preselection was less than 8% of the overall transaction cost in terms of execution time, which can be considered a negligible overhead. Compared to the FHE baseline system without indexing, HEBI reduced the computational workload down to 3% in terms of execution time.

Contribution of Paper B

Paper B addresses RQ2 and RQ4. With regard to RQ2, Paper B investigates computational workload reduction for biometric identification through indexing, while maintaining protection through all steps of the transaction. The high computational cost of FHE can be mitigated through indexing. At the same time, the PEKS scheme used for protected indexing does not introduce a significant workload of its own, making it applicable to workload reduction. The workload of post-quantum secure schemes such as lattice-based PEKS and FHE is the focus of RQ4. Here, Paper B shows that auxiliary indexing data can efficiently be secured against semi-honest quantum adversaries, building on an FHE-protected reference database. Furthermore, Paper B party addresses RQ1, as it explores PEKS as a privacy-preserving computation technique and shows that PEKS can be applied for biometric identification efficiently and securely, as long as biometric indexes can be represented as binary vectors.

Paper C: MT-PRO: Multibiometric Template Protection Based On Homomorphic Transciphering

FHE typically requires a two-server architecture, where one server holds the protected reference database encrypted with the FHE public key, and a second server holds the FHE secret key [270]. The security model then demands that the two servers do not collude. If they did, the reference database could be decrypted and the biometric templates would no longer be protected.

The aim of this work was to design a post-quantum secure biometric template protection scheme based on FHE that is secure even if the non-collusion assumption is violated. In other words, the reference database should remain protected even if an attacker gains access to the FHE secret key. Related work on this challenging problem has only achieved this through combining FHE with cancelable biometrics, which leads to a loss in accuracy [201]. In addition, we designed our solution to be multimodal, i.e., compatible with multiple different biometric modalities, in order to increase the recognition accuracy of the overall system and show the agility of our approach to different real-world applications.

To address the aforementioned research gap, we proposed the MT-PRO protocol which utilized *Homomorphic Transciphering* (HT) [71]. HT is a cryptographic technique designed to improve the storage and communication requirements of homomorphic encryption. It allows for the transformation of a symmetric ciphertext into an FHE ciphertext through a homomorphic evaluation of the symmetric scheme's decryption circuit. In our protocol, we apply this technique as follows. During enrolment, the biometric references are encrypted with the symmetric encryption scheme. Then, during a verification transaction, the client sends a symmetrically encrypted probe together with a homomorphic encryption of the symmetric secret key to the server, which transciphers both the probe and reference templates and computes the desired comparison function homomorphically.

In terms of the overall system security, MT-PRO achieves security against offline attacks on the reference database. Indeed, in this approach, an attacker that obtains the FHE secret key cannot decrypt the symmetrically encrypted biometric reference database during the offline phase. However, this security does not hold for the verification phase of the system, where both probe and reference templates are available as FHE ciphertexts. We however argue that a database in storage is most vulnerable to attacks by unauthorized external parties, rather than the computations performed on in-house servers, which can be assumed to adhere to the given protocol.

In addition, we argue that MT-PRO partly fulfils the challenging full-disclosure

security model of ISO/IEC 30136 [148], which demands that biometric templates must remain protected even if an attacker learns all system secrets. It is important to note that the symmetric key must remain secret in this case. We ensure this by assigning this key to the client and offering different options for key management, including password-authenticated key exchange.

While MT-PRO does not impair the biometric performance of the unencrypted baseline system, the computational cost is impractically high due to workload of HT framework. The implementation of MT-PRO is based on the framework provided by [71]. We made it publicly available to aid the reproducibility of our work. While the high computational workload can be partly attributed to missing code optimizations, MT-PRO cannot be accelerated without significant efficiency improvements on the cryptographic primitives. This challenge remains for future work.

Contribution of Paper C

Paper C is the main investigation into RQ3, with connections to RQ4 on quantum adversary models. The main objective of both Paper C and RQ3 is a feasible solution to the protection of FHE-based biometric template protection against malicious adversaries. In the case of Paper C, the solution provides the desired security, but not the computational efficiency to make the approach feasible. In connection to RQ4, this work also investigates different security model in a quantum adversary setting, and discusses the full-disclosure model of ISO/IEC 30136 [148], which is partly fulfilled through the protection against offline attacks. In addition, Paper C partly addresses RQ1 by showing that homomorphic transciphering can be applied to multibiometric BIP, even if the efficiency of this approach is not practical yet.

Paper D: Type^2: A Secure and Seamless Biometric Two-Factor Authentication Protocol Using Keystroke Dynamics

Password-based user authentication is widely deployed. However, it comes with the drawback of easy forgeability and offline attacks on password databases, in particular in the case where many users chose the same password string [237]. The aim of this work was to add a seamless two-factor authentication based on keystroke-biometric features to a standard password authentication setup.

The advantages of keystroke dynamics for the purpose of two-factor authentication are their seamless capture during the password typing, which does not require the user to access a second device or tool. While the accuracy of keystroke

authentication is generally speaking lower compared to physiological biometrics such as face recognition, it is a valuable addition to password-based authentication, used as a strengthening factor rather than a stand-alone authentication [167]. Furthermore, the keystroke authentication builds on a fixed-length known password, which is the least challenging of keystroke authentication applications with the highest performance.

Previous works on this topic lacked either accuracy preservation [6], preservation of computational performance [182], or post-quantum security [235]. We applied the CKKS [65] encryption scheme requiring no quantization or rounding on the keystroke timing vectors, which achieves post-quantum security. In terms of comparison functions, we evaluated 14 established keystroke anomaly detectors published alongside the publicly available CMU keystroke dataset [167], aiding the reproducibility of our work. Different protocols for *Password-Authenticated Key Exchange* (PAKE) [155] can be used in our proposed two-factor authentication protocol.

As our main contribution, we analysed the keystroke anomaly detectors with regards to their compatibility with FHE. Not all function components could be computed under encryption, or had feasible computational workload. Taking these considerations into account, we provided a security analysis of the adaptions necessary to accommodate the FHE application, and categorized the detectors into vector-based distance metrics, detectors based on matrix multiplication, and detectors requiring the evaluation of conditional statements. We gave a comprehensive overview of the cost of FHE operations for each detector, and proceeded with the detectors using vector-based distance metrics for the experimental evaluation.

In the experimental evaluation, we showed that variants of the Manhattan and Euclidean distance could be evaluated in the encrypted domain in less than 130 milliseconds for keystroke dynamic features, making our protocol efficient for real-world applications. As the FHE computations directly correspond to the plaintext operations, we inherited the biometric performance of the original work using the CMU keystroke dataset [167]. Finally, we concluded our work with a security analysis, showing that our proposed protocol fulfils the ISO/IEC 24745 [149] requirements for biometric template protection systems.

Contribution of Paper D

Paper D addresses RQ1 with its investigation into the efficiency of FHE protection for keystroke dynamic features. Our evaluation showed that keystroke timing vectors do not need to be altered in representation in order to achieve real-time efficiency. This shows that RQ1 should be evaluated

for each biometric modality individually, as Paper A on fingerprint verification resulted in the opposite conclusion. Furthermore, Paper D provides post-quantum security alongside a discussion of its computational cost and is therefore also related to RQ4.

Paper E: BRAKE: Biometric Resilient Authenticated Key Exchange

Cryptographic key exchange is an established component in authenticated communication, where cryptographic keys are exchanged between devices to facilitate encrypted communication [197]. However, there are many applications where the authentication of individuals is required, such as financial or legal transactions. To this end, we proposed *Biometric Resilient Authenticated Key Exchange* (BRAKE), using biometric verification to derive cryptographic key material, while the underlying biometric data remain protected.

Previous and concurrent works on the topic have either been computationally inefficient or limited in biometric feature representations they support, often limited to fixed-length binary inputs [48, 106, 260]. Our construction on the other hand builds on a recently proposed PAKE protocol [155] using an *Oblivious Pseudo-Random Function* (OPRF) [116]. In our work, we extended this protocol to allow for biometric verification instead of password verification.

A focus of this work was the security of our proposed BRAKE protocol against offline attacks. While offline attacks are harmful for password databases, a password or token can be exchanged easily. This is not true for biometric characteristics, which can be persistent over an individual's lifetime [166]. Therefore, additional considerations must be made to prevent offline attacks. To this end, we modified an established approach to biometric key generation known as *Fuzzy Vaults* [251]. In our construction, we remove the checksum stored alongside the protected biometric features, and replace it with an OPRF evaluation. Thereby, an attacker needs to interact with the system for every brute-force guess, allowing a rate-limiting on brute-force attacks and effectively preventing offline attacks. Significantly, the key material generated in our protocol is not derived from the biometric features directly, i.e., no biometric features are stored inside the key material. A brute-force attack on the public key would therefore not yield any information on the underlying biometric features.

Our protocol can be instantiated both with classical security, using Elliptic Curve Diffie-Hellman (ECDH) [87, 114] primitives, as well as post-quantum security using lattice-based primitives [12, 49]. The main challenge of the lattice-based instantiation was the OPRF, for which only one lattice-based instantiation ex-

isted in previous works at the time of submission [12]. However, the OPRF in question was designed with verifiability, resulting in an infeasible computational workload. Therefore, we adapted the OPRF construction such that it could be applied in a semi-honest setting with real-time efficiency. However, the efficiency of an efficient verifiable lattice-based OPRF goes beyond the scope of this contribution and hence is left for future work.

We implemented our BRAKE protocol and evaluated the biometric and computational performance on publicly available datasets. Our protocol is achieves efficient transaction times for the cryptographic key exchange of under one second on commodity hardware from the biometric capture to the completed key exchange, including communication cost.

Contributions of Paper E

Paper E addresses RQ1 through the application of oblivious computation to biometric authentication. Notably, this is the only approach not relying on FHE in this thesis, and therefore an important addition to the evaluation of RQ1. The thorough security analysis and compatibility with different biometric modalities and feature representations complete the evaluation of RQ1. In addition, Paper E addresses the use of verifiable computation as protection against malicious adversaries, which is the objective of RQ3. Finally, Paper E partly addresses RQ4 through the considerations of semi-honest and malicious quantum adversaries, where malicious security was not found to be computationally efficient. Contrary to the FHE-based contributions of this thesis, which inherit post-quantum security from their lattice-based construction, the BRAKE protocol based on the combination of an OPRF and a KEM required a designated post-quantum secure instantiation in addition to the classically secure instantiation.

Paper F: Post-Quantum Secure Biometric Systems: An Overview

The long-term protection of biometric data is important due to their uniqueness and persistence, as recognized by international laws and standards [109, 149, 166]. Therefore, the application of post-quantum cryptography has recently emerged as a research field in the realm of biometric template protection [30, 33, 44, 103]. This work aims to give an introduction and overview of the current literature landscape and highlight current research challenges.

Surveys related to this topic either focused on concrete cryptographic schemes such as FHE alone [75, 267], or considered mostly biometric information protection with classical security [54, 216]. On the other hand, literature reviews exist for various cryptographic contributions [78, 196], but not their application

to biometrics. Therefore, a literature survey dedicated to post-quantum secure biometric systems was missing from the current research landscape.

The survey introduced biometric information protection and relevant post-quantum secure cryptographic schemes. As the main contribution, a comprehensive literature review following a dedicated taxonomy were presented, where a large number of publications were analyzed and compared. A specific focus within the literature review were open-source implementations, including an overview of established libraries for FHE and PQC that can aid the selection of secure parameter choices and reproducibility of future works in post-quantum secure BIP [22, 53, 136, 194, 234, 244].

Utilizing the insights won during the literature review, the survey identified three main open research challenges. Firstly, the analysis revealed that security against malicious quantum adversaries was only investigated in a minority of the reviewed works [3, 18, 34, 66, 192, 201]. To mitigate real-world threats, further research on post-quantum secure BIP under this security model is therefore required. Secondly, the majority of reviewed works focused on FHE, while other cryptographic techniques such as MPC or FE have been applied successfully to post-quantum secure BIP, and have shown better performance in some cases [30]. Finally, the need for open-source implementations and the applicability of synthetic datasets was discussed [158].

Contributions of Paper F

Paper F addresses all four research questions RQ1, RQ2, RQ3, RQ4 through an overview of the research landscape in the area of post-quantum secure BIP. The literature survey and analysis aimed at providing fellow researchers an overview over the emerging research field of post-quantum secure BIP, with an introduction into relevant cryptographic primitives and analysis of key challenges. Thereby, we hoped to increase future work on this important topic and contribute to the long-term protection of biometric data in real-world applications.

1.3 Outline

The remainder of this thesis is structured as follows: Chapter 2 gives background information about biometric performance metrics, post-quantum cryptography, and fully homomorphic encryption. Chapter 3 discusses related works relevant to this thesis, before Chapter 4 presents a summary of the contributions and their limitations as well as future work. The research contributions are presented as published or submitted research articles in Papers A to F.

Chapter 2

Background

This chapter introduces background information relevant to the scientific contributions presented in this thesis. First, biometric evaluation metrics and information protection approaches will be discussed. Then, cryptographic background information will be given for post-quantum cryptography, including fully homomorphic encryption.

2.1 Biometric Evaluation Metrics

Biometric systems can perform two main transactions: *verification*, or a one-to-one comparison of a biometric probe against a biometric reference corresponding to a chosen identity, or biometric claim, and *identification*, or a one-to-many comparison of a biometric probe with an unknown biometric claim against multiple biometric references. In a verification transaction, a binary decision about the acceptance of the biometric claim is revealed, whereas an identification transaction determines if none, one, or more than one of the references yield an accept decision when compared to the probe. In addition to these two transactions, *biometrics-based key establishment* can be derived from a successful verification transaction.

All three types of transactions are preceded by an *enrolment* transaction, where biometric references are stored in a data storage subsystem for further comparison. The transactions of enrolment, verification, and identification have been standardized in the ISO/IEC 19795-1 [152] standard on biometric performance testing and reporting alongside standardized biometric vocabulary and evalu-

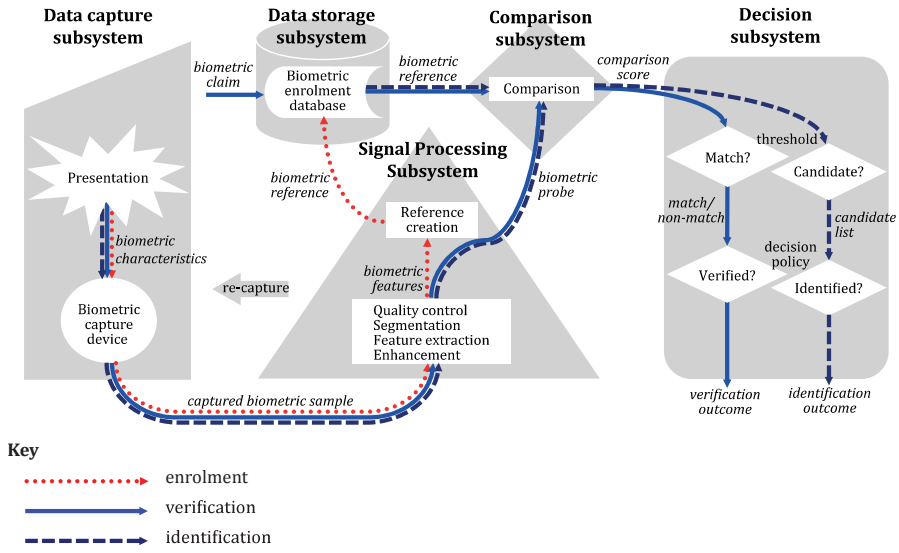

Figure 2.1: ISO/IEC 19795-1 flowchart of biometric system components [152].

ation metrics. Figure 2.1, taken from ISO/IEC 19795-1, gives an overview of a generic biometric system including its subsystems and transactions.

Before we describe the standardized evaluation metrics, we briefly detail the workflow of the transactions depicted in Figure 2.1. All transaction begin with the capture and feature extraction steps. First, a biometric sample is captured by a capture device, e.g., a face image is taken using a camera. From this sample, features are extracted according to the biometric modality. For example, state-of-the-art face recognition is based on deep-learning based feature extractors [191], which have been explored for further modalities such as fingerprint [104] and iris [202] as well. The feature extraction process results in biometric features, which are referred to as *reference templates* if they are stored as references in the data storage subsystem during enrolment. In case of deep-learning based feature extraction, this process yields real-valued vectors of a fixed dimension. It is important to note that feature extraction algorithms and the resulting feature representations may be more complex than in the example of deep-learning based facial features [38].

During a verification transaction, a fresh biometric sample is captured and features are extracted as during the enrolment transaction. This feature represen-

tation is then compared against the reference template corresponding to a biometric claim of the sample using a comparison function with both probe and reference template as the inputs. For example, fixed-length features vectors are often compared using established distance metrics such as the Euclidean distance. Independent of the specific comparison function, this step results in a comparison score, which can be either a similarity score, where a higher score indicates a higher similarity of the probe and reference, or a dissimilarity score, where a lower score indicates higher similarity. Finally, the comparison score is compared against a predefined threshold, and the transaction results in an accept decision if the score is higher than the threshold for similarity metrics, or lower than the threshold for dissimilarity metrics.

In an identification transaction, this procedure is iterated over all stored references in the database, where all references that pass the threshold comparison are added to a candidate list. Depending on the decision policy, only the reference identifier corresponding to the best comparison score is returned, or multiple candidates are revealed.

We now introduce the main standardized evaluation metrics, beginning with verification transactions. In a set of such transactions, the *False-Match Rate* (FMR) describes the percentage of accepted comparison trials where the reference and probe stem from different instances, i.e. different fingers. Such a comparison is referred to as a *non-mated comparison*. If a non-mated comparison results in an accept decision, it constitutes a false positive (*false match*). Colloquially, non-mated comparison trials can be referred to as imposters comparisons. Notably, this phrasing implicates a malicious intent of the data subject, which is not measured by the FMR. Instead, false matches may also occur for honest subjects due to the variance of biometric features representations after the capture process, or their fuzziness.

The corresponding metric to the FMR is the *False Non-Match Rate* (FNMR), indicating the percentage of rejected verification transactions where probe and reference stem from the same instance, and should therefore have been accepted. Such comparisons are referred to as *mated comparisons*. If a mated comparison results in a reject decision, it constitutes a false negative (*false non-match*). Informally, the FNMR can be thought of as a convenience measure of the system, as it expresses the inconvenience of honest users that are rejected by the automated recognition system and may have to repeat the biometric authentication or use a fall-back authentication procedure. The FMR on the other hand determines the biometric security of the system. Depending on the application, it can be argued that the FMR is therefore the more significant measure of the two. Biometric feature extraction and comparison algorithms with a high FMR risk unauthorized access through an acceptance of non-mated comparisons, or imposter compar-

isons. Biometric systems operated in high-risk scenarios such as automated border control are required to achieve a FMR lower than 0.01% [117].

It is important to note that the FMR is a metric of the comparison subsystem and can be generalized for the entire biometric system through the *False Accept Rate* (FAR), which takes into account those instances where no features of quality high enough to be used for biometric comparison could be extracted. This measure is referred to as the *Failure To Acquire Rate* (FTAR), which can be used to define the FAR as

$$FAR = FMR(1 - FTAR). \tag{2.1}$$

While biometric performance is typically reported through the FMR, the FAR has recently been used to define the security of biometric systems compared to previous approaches based on the entropy of biometric feature representations [251]. Even though high entropy can positively impact the recognition accuracy by allowing for more discriminate feature representations, it can only be considered an upper bound for the overall security of a biometric system. Instead, the FAR gives a more realistic and reliable measure of the biometric security under brute-force attacks. For example, the entropy of minutiae-based fingerprint representations has been estimated at 82 bits [215], while [251] derive the FAR security of their concrete fingerprint BIP scheme at around 20 bits. Following the argumentation of [251], we continue to refer to the FAR security as the measure for biometric security. We note however that in the case where the FAR of a given system is not measured or reported, the FMR can be used instead to derive the biometric security of the system. The corresponding metric to the FAR is the *False Reject Rate* (FRR), which is defined as

$$FRR = FTAR + FNMR(1 - FTAR). \tag{2.2}$$

For identification transactions, the corresponding metrics of *False-Positive Identification Rate* (FPIR) and the *False-Negative Identification Rate* (FNIR) apply. As in the verification scenario, the FPIR quantifies the biometric security of the overall system, while the FNIR indicates the convenience. Notably, high biometric performance of identification transactions is significantly more difficult to achieve than in a verification scenario due to the approximation

$$FPIR \approx (1 - FTAR)(1 - (1 - FMR)^N), \tag{2.3}$$

given in the standard [152]. This relation indicates that the FPIR grows exponentially with the number of enrolled subjects N, creating a challenging sce-

nario even for biometric systems that achieve a FMR lower than 0.01%. The approximation can be considered an exact equation in the case where no failure-to-acquire errors were measured and all mated comparison trials resulted in the same FMR at a given threshold, which cannot be assumed to arbitrarily hold true [92].

2.2 Biometric Information Protection

The uniqueness and persistence of biometric characteristics put them at risk of attacks when processed unprotected in the transaction described above [166]. Even though the feature extraction process yields an abstract representation of the original biometric characteristics, this representation alone is not sufficient to hide the data subject's characteristics or identity. Indeed, sample reconstruction attacks have been successfully conducted for the most widely used biometric modalities, with the most notable works including [185] for face, [119] for iris, and [59] for fingerprint. Reconstructed biometric samples can be used for more effective presentation attacks or leak personal information such as medical conditions or ethnic origin, which are protected under international laws [109].

Therefore, additional protection mechanisms need to be applied after the feature extraction process to ensure the privacy of biometric information. The international ISO/IEC 24745 [149] standard has defined three main requirements for *Biometric Information Protection* (BIP) which all protection mechanisms should fulfil. These requirements are the *unlinkability*, *renewability*, and *irreversibility* of biometric reference templates stored and transmitted within the biometric system, as well as probe feature vectors handled during verification and identification transactions. In particular, these requirements must be upheld during the comparison step of the transaction, requiring solutions that allow for the evaluation of comparison functions on protected templates. Before we describe such approaches, we define the ISO/IEC 24745 requirements in more detail.

Unlinkability requires that an attacker cannot link two protected biometric templates or their corresponding data subjects stored in different applications. To fulfil this requirement, the applied protection mechanism needs to allow for the creation of a large number of protected templates, such that several protected templates created from the same input, i.e., the same reference template, cannot be distinguished from protected templates created from different inputs.

Renewability is closely linked to unlinkability, requiring that multiple unlinkable protected templates can be created without the need to re-enrol a subject. In particular, should a protected template or its source leak, renewability ensures

that this biometric instance can still be securely enrolled into the system. In other words, knowledge of a biometric template does not allow an attacker to determine whether the corresponding subject is enrolled in an application, or gain additional information about further templates enrolled for this subject. In other words, given two protected templates, an attacker should not be able to efficiently determine if both protected templates were created from the same template, or from different templates.

Irreversibility describes the protection of the biometric data against unauthorized use. Concretely, it should be impossible for an attacker to retrieve original samples given only protected templates, implying the confidentiality of biometric templates. As discussed above, this requirement should be upheld both during the storage and comparison stages. In particular, irreversibility is not fulfilled if the protection mechanism needs to be removed in order to compute a comparison between a probe and one or multiple references, as they would be vulnerable to reconstruction attacks at this point in the pipeline [185].

A final requirement that is stated implicitly in ISO/IEC 24745 is the performance preservation of the transactions under the applied protection mechanism. It is clear that both the biometric performance, i.e., recognition accuracy, and the computational performance should not be decreased significantly through the application of BIP to an unprotected biometric system. If a protection mechanism significantly decreases the recognition accuracy of the overall system, the biometric security in terms of false-accept attacks is mitigated. Therefore, the preservation of the biometric performance is crucial to obtain a secure, as well as a private, biometric system. The computational performance of the unprotected system is typically increased by adding a layer of protection to the original data. Different approaches have achieved significantly different performances, from real-time efficiency with transaction times of under one second to infeasible computational workload of several hours per comparison [38].

Historically, three different classes of approaches have been established to meet the ISO/IEC 24745 requirements for secure and private biometric systems [270]. The first approach is based on *irreversible feature transformations* which can be viewed as part of the feature extraction process. Approaches that fall into this category are cancelable biometrics [214], robust hashing [247], Bloom filter-based approaches [128] or secure sketches [177]. While such approaches have been shown to be computationally efficient, they typically suffer from a decrease in recognition accuracy compared to unprotected system [216]. As discussed above, this negatively impacts the overall system security through an increased FMR. Additionally, they have not in all cases undergone the same level of cryptanalysis or standardization as other cryptographic approaches, and it can therefore be challenging to make reliable and quantifiable guarantees about their se-

curity. Similarly, unlinkability and irreversibility need to be measured experimentally for each approach [127], yielding a high overhead and making it more difficult to compare or reproduce approaches.

The second approach is referred to as *biometric cryptosystems*, in which biometric feature representations are tied to secrets that can be further used as key material in cryptographic protocols [61]. Such schemes utilize error-correcting codes to correct the variable component of biometric captures, and create a stable output associated to a biometric instance. Even though concerns about the security and unlinkability of biometric cryptosystems have been raised in the past [231], new and improved constructions have been proposed that achieve high accuracy, unlinkability, and information-theoretic security against reconstruction attacks [218, 220, 250]. However, their construction continues to be vulnerable to offline attacks, as a checksum of the secret tied to the biometric input is stored alongside the protected template, allowing an attacker to determine whether a brute-force attack succeeded. This attack can be executed offline without the knowledge of the system provider, and can consume unlimited resources outside of the provided infrastructure [37].

The third and final category of BIP approaches can be summarized as *biometrics in the encrypted domain*. While previous works have focused on homomorphic encryption [270], there exist further cryptographic techniques such as *Secure Multiparty Computation* (MPC) [269] and *Functional Encryption* (FE) [47] that have been applied to evaluate comparison functions on encrypted biometric templates. One important advantage of biometric comparisons in the encrypted domain is its limited impact on the recognition accuracy compared to the unprotected systems. For schemes that directly translate operations on encrypted data into the plaintext domain, the outcome of the computation, e.g., a biometric comparison function, it not altered.

However, depending on the feature representation, the quantization of floating point values into integers or bits is required to ensure compatibility with the cryptographic scheme, which can lead to a loss in accuracy. In addition, some schemes operate on approximations of fixed-point data and introduce an additional inaccuracy to the computation [65], which typically does not impact the recognition accuracy significantly [44]. Overall however, the aforementioned cryptographic techniques allow for the evaluation of biometric comparison functions on private input data, i.e., in a way that does not reveal the unprotected templates. Either the comparison score or the decision outcome are revealed as a result of the respective transaction.

The security of these cryptographic schemes is based on hard mathematical problems and supported by a research community that conducts extensive crypt-

analysis on the schemes independent of their application to biometrics or other fields. This increases trust in the biometric protection, as security improvements and the selection of secure parameters are provided by cryptographic experts. The clear limitation of computations on encrypted data is however their computational workload, which can render solutions infeasible for practical use [38, 210]. Particularly in identification transactions, workload reduction approaches that maintain the privacy of the reference database are therefore relevant [33, 103].

2.3 Post-Quantum Cryptography

The protection mechanisms discussed above have recently come under the additional threat of quantum adversaries due to the continued advances in quantum computing. The latter allow for an implementation of two algorithms presented in the 1990s: *Shor's factorization algorithm* [238] and *Grover's search algorithm* [132].

Grover's algorithm implemented on a quantum computer offers a square-root speedup on unstructured search problems. This attack affects symmetric ciphers such as the *Advanced Encryption Standard* (AES) [77] or cryptographic hash functions, as it allows for more efficient brute-force attacks. To counter this attack, the key size of symmetric cryptographic schemes need to be doubled and in this case, may remain computationally feasible.

Shor's algorithm however has a more severe impact, as it allows for efficient attacks on the mathematical problems that asymmetric cryptography has been based upon in the past decades. More concretely, Shor's algorithm allows for a polynomial time solution of the factorization and discrete logarithm problems, which underlie the majority of the currently deployed public-key cryptography. Examples of schemes vulnerable to this attack are RSA [224], Diffie-Hellman [87], ElGamal [102], and Paillier [203]. Increasing the parameter sizes of these schemes to a level secure against Shor's algorithm would render the schemes infeasible in terms of execution time. Therefore, new cryptographic schemes based on different mathematical problems are required which withstand attacks based on Shor's algorithm [10].

To aid the research and development of such new cryptography that is believed to be hard to break even in the presence of a quantum computer, NIST has run a *Post-Quantum Cryptography* (PQC) standardization process [10]. The standardization process focused on public-key cryptography, and considered contributions based on five types of cryptography for post-quantum security: lattice-based, code-based, multivariate, hash-based, and isogeny-based cryptography.

Out of these, lattice-based and code-based schemes have emerged as the most robust and relevant candidates [10].

Lattice-based cryptography is constructed from hard mathematical problems based on lattices, which can informally be defined as discrete subgroups of real vector spaces. Hard problems on lattices have been studied since the 1990s [140], including the *Shortest Vector Problem* (SVP) and *Closest Vector Problem* (CVP) [204]. Regev [222] later proposed the seminal *Learning with Errors* (LWE) problem many modern schemes build upon. Variants of this problem based on polynomial rings, or *Ring-Learning With Errors* (RLWE) [183], facilitate smaller public key sizes than the original LWE problem. Aside from LWE-based schemes, schemes based on the *NTRU problem* [140] are also currently considered within the NIST PQC standardization process. We point the reader to Paper F for definitions of important lattice problems and further details on post-quantum cryptography that has become relevant for BIP.

Code-based cryptography utilizes error-correcting codes and was first proposed in the McEliece cryptosystem in 1978 [189]. Messages in code-based encryption schemes can be encrypted as erroneous codewords, and efficiently decrypted using the decoding algorithm of the chosen code [262]. Three code-based encryption schemes were under consideration of the NIST PQC standardization effort [10]. We point the reader to Paper F for further discussion on the McEliece cryptosystem and its application to BIP.

2.4 Fully Homomorphic Encryption

A class of encryption schemes particularly relevant to this thesis are *Fully Homomorphic Encryption* (FHE) schemes, which allow for computations on encrypted data. Through this property, biometric comparison functions can be evaluated while maintaining template protection, revealing only the comparison score. As the hardness of FHE schemes is based on hard lattice problems, they can be considered to be post-quantum secure [121].

More concretely, a homomorphic encryption scheme is defined through the property that

$$Dec(Enc(x \mathbin{\text{☼}} y)) = Dec(Enc(x) \mathbin{\text{☾}} Enc(y)), \tag{2.4}$$

where $Enc(\cdot)$ denotes the encryption function, $Dec(\cdot)$ denotes the decryption function, ☼ is an operation on plaintext messages x and y, and ☾ is an operation on ciphertexts. The two operations ☼ and ☾ can be the same, i.e., addition of ciphertexts corresponds to an addition of the underlying plaintexts, or

different, e.g., a multiplication of ciphertexts corresponds to an addition of the underlying plaintexts. Homomorphic encryption schemes can be categorized by the number and type of operations they support. In the following, we focus on post-quantum secure homomorphic encryption schemes within the different categories.

Partially Homomorphic Encryption (PHE) schemes allow for the homomorphic evaluation of only one operation, e.g., addition, an unlimited number of times. A number of lattice-based encryption schemes that have additive homomorphic properties have been applied to BIP [170]. They can be applied to evaluate simple comparison functions such as the Hamming distance of two encrypted biometric templates, which can be expressed as a homomorphic addition modulo 2.

In addition to PHE schemes, the lattice-based key encapsulation mechanisms Kyber [49], which was recently standardized by NIST as a result of the PQC standardization competition [10], Saber [101], and the previously discussed code-based encryption scheme McEliece [189] have been applied to BIP using their partially homomorphic properties [19, 226]. Even though these schemes have not been designed with homomorphic evaluations in mind, their underlying primitives allow for homomorphic additions, and can thereby also be applied to post-quantum secure BIP.

Somewhat Homomorphic Encryption (SHE) schemes allow for three operations, e.g., addition, multiplication, and vector rotations, to be evaluated homomorphically, where at least one of the operations is restricted to a limited number of iterative computation. Going above this given limit results in incorrect decryption. A prominent example of an SHE scheme is the lattice-based *Brakerski-Vaikuntanathan* (BV) [52] scheme.

Fully Homomorphic Encryption (FHE) schemes allow for multiple homomorphic operations with an unlimited number of iterative computations. The first construction that fulfils this requirement was Gentry's [121] lattice-based FHE scheme with the seminal concept of *bootstrapping* that allows to maintain correct decryption under iterative computations. In theory, this approach therefore allows for an unlimited number and combination of different homomorphic operations such as addition and multiplication. In practice however, the computational workload of the bootstrapping step can render this construction infeasible.

On a lower abstraction level, FHE schemes consist of the following algorithms.

- $(sk, pk) \leftarrow KeyGen(1^\lambda)$: on input of the security parameter λ, this algorithm generates a secret key sk and public key pk, where pk includes the homomorphic evaluation keys.

- $c_m \leftarrow HomEnc(pk, m)$: on input of the public key pk and a message m, this algorithm outputs a ciphertext c_m.
- $c_{f(m_1,m_2)} \leftarrow HomEval(pk, c_{m_1}, c_{m_2})$: on input of the public key pk and two ciphertexts c_{m_1} and c_{m_2}, this algorithm outputs an encryption $c_{f(m_1,m_2)}$ of the evaluation of a function f on the underlying plaintext messages m_1 and m_2.
- $m' \leftarrow HomDec(sk, c_m)$: on input of the secret key sk and ciphertext c_m, this algorithm outputs a message m'.

Instantiations of the *HomEval* algorithm vary between different FHE schemes, but typically include homomorphic addition, multiplication, and element-wise rotation of encrypted vectors [65].

Two different approaches have been proposed to address the handling of the computational overhead of FHE. In *levelled* FHE schemes, homomorphic operations up until a fixed number of times can be computed without requiring bootstrapping, where the total computational workload increases with each added level. If the allowed limit is exceeded, the costly bootstrapping operation must be computed. Schemes that follow this approach are the *Gentry-Halevi* (GH) [122], *Brakerski-Fan-Vercauteren* (BFV) [50, 112], the *Brakerski-Gentry-Vaikuntanathan* (BGV) [51] and the *Cheon-Kim-Kim-Song* (CKKS) [65] schemes. In contrast to levelled FHE schemes, a different approach is followed by the *TFHE* [69] and *FHEW* [98] schemes, which prioritize the optimization of the bootstrapping operation. Thereby, the number of homomorphic operations becomes unlimited. We refer the reader to Paper F for further insights into the different FHE schemes.

Chapter 3

Related Work

In this chapter, related work to the contributions of this thesis is presented. The chapter follows the structure of the research questions presented in Chapter 1. Related work relevant to each research question is discussed separately for the first two research questions and combined for the last two research questions.

3.1 Feature Representations Under Encryption

The first research question is concerned with the efficiency and security of different BIP approaches. A significant contributing factor to this question is the compatibility of biometric feature representations with privacy-preserving computation techniques, which will be the focus of this section.

Iris features have historically been represented as fixed-length binary vectors, following the seminal work by Daugman [80]. Not only has this representation been shown to yield a high recognition accuracy, but it is also compatible with cryptographic techniques that natively operate on binary inputs. The comparison of iris features in the encrypted domain has therefore been considered soon after their initial proposal by Schoenmakers and Tuyls [233], who used the Paillier [203] encryption scheme to compute Hamming distances between fixed-length binary vectors. This approach has received further interest for iris BIP with classical security [126]. Later, iris features have also been used in combination with the lattice-based homomorphic encryption scheme NTRU [140] by Kolberg et al. [170], the GH [122] scheme by Yasuda et al. [270], the BGV [51] scheme by Torres et al. [253] and Cheon et al. [64], and the BFV [50, 112] scheme

by Morampudi et al. [192], Bassit et al. [29], and Vallabhadas and Sandhya [258], all of which can be considered to be post-quantum secure.

In addition to (F)HE schemes, MPC has also been applied to binary fixed-length iris features. One of the first works to apply Yao's *Garbled Circuits* (GC) protocol [269] were Blanton and Aliasgari [41]. However, their work did not provide post-quantum security. Through applying a lattice-based GC protocol [55], Bauspieß et al. [30] later achieved post-quantum secure iris verification based on MPC.

While all of the works discussed above utilize a fixed-length binary iris feature representation, deep-learning based iris features have recently been explored [202]. The rapid increase in recognition accuracy for deep face recognition [191] was the motivation behind the construction of deep iris features. Furthermore, efficiency improvements could be expected from smaller feature vectors, as the typical length of the Daugman representation is an order of magnitude larger than typical deep-learning based feature vectors [30, 170]. However, the recent initial constructions have not outperformed the traditional Daugman feature representation in terms of recognition accuracy [202].

Deep-learning based feature extraction models have become state-of-the-art for face recognition [86, 191]. They typically produce fixed-length vectors of small floating point numbers centered around zero, which have been computed such that distance metrics can be applied to distinguish mated from non-mated comparisons. However, floating point numbers, even if transformed to fixed-point numbers, do not trivially correspond to the plaintext representations of cryptographic schemes that use binary or integer inputs. Indeed, the only FHE scheme that naturally operates on fixed-point numbers is the CKKS [65] scheme, which allows for an approximate computation on the latter up to a predefined accuracy level. For all other schemes, a quantization or encoding needs to be applied to map deep-learning based features into plaintext spaces.

Drozdowski et al. [96] analyzed the quantization and binarization of deep face templates and showed that only minor accuracy losses can be expected from the investigated quantization and binarization techniques. Notably, four quantization intervals are sufficient to maintain an acceptable recognition accuracy, ensuring that feature vectors do not grow to infeasible length under binarization. This and similar encodings have been applied to post-quantum secure face recognition based on homomorphic encryption by many works in the past decade, including [19, 44, 142, 157, 171, 201, 210, 226, 243, 249, 254, 268]. For an in-depth discussion of these works, we refer the reader to Paper F. MPC for face recognition has also been proposed by Sadeghi et al. [228], however not with post-quantum security. Further, less frequently applied biometric modali-

ties such as keystroke dynamics, gait, finger and hand veins, and voice have also been studied under protection by privacy-preserving computation techniques. Whenever these features can be represented as fixed-length floating point vectors, approaches that apply to face recognition may be transferred. Such feature representations can be found in [6, 18, 178, 182] with post-quantum secure protection. For an in-depth review of BIP with classical security, we refer the reader to the surveys by Barni et al. [27] and Bringer et al. [54].

Fingerprint features have historically been based on minutiae, or significant points in the fingerprint ridge pattern. As such, they have been standardized in ISO/IEC 19794-2 [151]. The cryptographic protection of minutiae-based representations was recently recognized as challenge by Engelsma et al. [104], who argue for the need of accurate fixed-length fingerprint representations based on deep-learning or filter-based approaches [153]. The latter allow for a direct binary representation of feature vectors, which was used by Kim et al. [169], who applied the TFHE [69] scheme for their protection, thus achieving post-quantum security. Prior to Paper A, variable-length feature representation had only been protected with classical security, amongst others by Barni et al. [26], Gomez-Barrero et al. [126], and Yang et al. [266]. In addition, MPC with classical security was applied for fingerprint BIP by Liu and Zhao [180], Zhang and Koushanfar [277], and Gilkalaye and Derakhshani [124].

3.2 Efficient and Secure Biometric Identification

The second research question focuses on two challenges regarding biometric identification in the encrypted domain: its security and efficiency. Regarding the latter, Drozdowski et al. [92] gave an overview of workload reduction approaches to biometric identification without BIP, which can be transferred to protected identification transactions. As above, the focus lies on BIP with post-quantum security.

A straightforward approach to the application of FHE to biometric identification was presented by Drozdowski et al. [91]. However, this work did not consider any workload reduction and proved to be computationally infeasible. One category of approaches to improve upon this baseline system is concerned with optimizing the cost of single biometric comparisons, such that the total cost of the exhaustive search is reduced. Bauspieß et al. [31] extended the plaintext packing technique proposed by Boddeti [44] to biometric identification such that multiple biometric comparisons could be computed at the cost of one. In combination with feature dimensionality reduction, this approach yields a quadratic improvement over the baseline system in [91]. The same approach was later

discussed by Ibarrondo et al. [145], who extended their packing approach by a group testing technique that determines the identification outcome.

A different encoding approach that yielded high efficiency was presented by Engelsma et al. [103], outperforming the previously discussed works [31, 91, 145]. In their work, they encoded the feature dimensions into individual ciphertexts instead of using a packed encoding, which allows them to utilize a trade-off between the computational efficiency and communication requirements. While [103] also utilize feature dimensionality reduction, their feature encoding is not trivially compatible with further preselection, contrary to [31] and [145].

Huang and Wang [143] and Bai et al. [23] expanded upon the state-of-the-art through the addition of result-revealing protocols based on classically secure MPC. They argue that this improves the security of the system as no comparison scores are revealed, but only binary identification outcomes. However, the classically secure components of their protocols would need to be exchanged for post-quantum secure instantiations to achieve long-term protection throughout the entire transaction. Notably, the scheme presented in [23] is the most efficient in terms of computational performance, outperforming [103].

An additional improvement to biometric identification can be achieved through preselection, or determining a subset of the enrollment database that is likely to contain the mated reference. Then, the expensive exact comparisons in the encrypted domain only need to be computed on the subset, lowering the computational workload significantly. However, this approach introduces an preselection error in the case where a mated comparison is not included in the selected subset [92].

Different preselection approaches have been explored for FHE-protected reference databases. Drozdowski et al. [95] utilized feature fusion to construct binary search trees through averaging feature vectors iteratively, and traversing the search tree upon an identification transaction. While their approach ensures post-quantum security through the encryption of the fused search vectors, it is inflexible in terms of database changes, and limited in workload reduction potential for large databases. Osorio-Roig et al. [200] mitigated this shortcoming by proposing an efficient and accurate indexing scheme based on short stable indexing strings derived from clustered feature vectors. However, their unprotected indexing is vulnerable to template reconstruction attacks, as Paper B presented in this thesis shows.

Bauspieß et al. [33] proposed an approach to protected preselection using *Public-Key Encryption with Keyword Search* (PEKS) [45]. In their work, an encrypted reference database can be filtered based on soft-biometric attributes such as the

gender or ethnicity of the enrolled subjects. However, their construction relies on a ground-truth assumption of these soft-biometric attributes that cannot be realistically reproduced and is limited to face identification, as soft-biometric descriptors with high correctness are difficult to derive from other biometric modalities. PEKS had previously been applied for biometric verification by Zhang et al. [278]. However, their approach requires strong statistical assumptions to the feature representation that cannot be assumed to be applicable for arbitrary biometric modalities [33]. Zhang et al. [276] further utilized PEKS for biometrics-based key generation, but do not apply their scheme to biometric identification [33].

3.3 Security Models

In this section, we combine the related work relevant for the third and forth research questions on malicious security and post-quantum security. A comprehensive overview of post-quantum secure biometric systems and their respective security models is given in Paper F.

Paper F revealed that the majority of post-quantum secure BIP schemes are considered under the semi-honest adversary model, including the seminal works presented in [44, 103, 253, 270]. However, the semi-honest adversary model does not include realistic capabilities of real-world adversaries which may deviate from the protocol. Different approaches to achieve security against malicious adversaries have therefore been explored for post-quantum secure BIP.

Abidin and Mitrokotsa [3] described an attack against the semi-honest verification scheme presented by Yasuda et al. [270] along with a mitigation of this attack based on PIR [72], achieving protection against a malicious client. Arjona and Baturone [18], Cheon et al. [66], and Morampudi et al. [193] considered malicious servers and achieve security against the latter implicitly through their respective constructions, but without explicit use of verifiable computation. For a more detailed comparison of these works, we refer the reader to Paper F.

In addition to the established semi-honest and malicious adversary models, the ISO/IEC 30136 [148] standard defines the challenging full disclosure model, where an adversary is assumed to receive all secrets used within the BIP scheme. In terms of FHE-based BIP, this translates to the disclosure of the secret decryption key to the adversary, a scenario not typically considered within cryptographic security models. Nevertheless, Otroshi et al. [201] proposed a BIP scheme that combined FHE with a feature transformation approach that remains in place even if the FHE ciphertexts are decrypted. However, the protection of the feature transformation approach cannot be considered post-quantum se-

cure. In addition, security against a loss of the FHE secret key does not imply security against malicious adversaries, as parties deviating from the protocol would not be detected in the protocol by [201]. We refer the reader to Paper F for further discussion on the full disclosure model and its relation to malicious security.

Chapter 4

Conclusion

In this chapter, the contributions of this thesis are summarized and highlighted, and conclusions for each research question are presented. In addition, limitations of the work presented in this thesis are discussed, and finally, open research problems and opportunities for future work are identified.

4.1 Summary of Contributions

We summarize the contributions of this thesis based on the research questions, and reiterate the latter for this purpose.

> **Research Question 1**
>
> Which privacy-preserving computation techniques are best suited for biometric information protection?
>
> - How do different approaches to biometric information protection compare in terms of their security and efficiency?
> - How can biometric features be represented to aid different encoding mechanisms used in privacy-preserving computation techniques?

The main contributions to this first research question were given in Paper A, Paper D, Paper E, and Paper F, while the remaining two Papers B and C contributed partly. Regarding the comparison of privacy-preserving computation techniques for BIP, the focus of this thesis was on BIP schemes with post-quantum security, a comprehensive overview of which was given in Paper F.

Papers B and D showed that FHE can be applied to achieve privacy-preserving biometric verification and identification secure against semi-honest adversaries. Notably, even though Paper B was concerned with the modality face and Paper D with the modality keystroke, both works consider fixed-length feature representations under FHE protection, which aids the efficiency of the presented protocols. In contrast to this, Paper A showed that efficiency could not be maintained for variable-length fingerprint representation, which results in an infeasible computational overhead when compared under FHE. Additionally, Paper C showed that the efficiency of FHE-based BIP schemes cannot generally be expected to be high if adversary models stronger than the semi-honest model are considered. In Paper C, using the cryptographic tool of homomorphic transciphering introduced a significant overhead over purely FHE-based comparisons. Notably, the proof-of-concept implementation this work built upon was not optimized for execution time [71], as the FHE comparisons using the CKKS [65] encryption scheme were evaluated to be more expensive than in related work [171] and Paper B. Given the fixed-length feature representations used in Paper C, improvements on the cryptographic components and implementation of the transciphering framework can be expected to decrease the computational workload in the future.

In Paper E, fuzzy vaults [160] were combined with oblivious evaluations [155], showing the feasibility of the latter for BIP. The feature representation in this case depends on the fuzzy vault constructions, which have been evaluated for different biometric modalities in independent works [218, 220, 251]. While the biometric performance of these schemes can be lower than in FHE-protected BIP schemes due to additional quantization required during the encoding step, Paper E showed that biometrics-authenticated key exchange can be efficiently constructed with both classical and post-quantum security, where the latter was only efficient in the semi-honest adversary model.

Overall, the contributions to the first research question revealed that the feature representation is a significant factor for the efficiency of privacy-preserving computation techniques such as FHE, as it determined the number and type of comparison steps that need to be computed privately. On the other hand, the efficiency of privacy-preserving computation techniques is impacted by the assumed adversarial capacities, where both security against malicious adversaries and security under the full disclosure model can render post-quantum secure BIP schemes infeasible. Finally, it can be concluded that post-quantum secure cryptographic techniques offer long-term protection to biometric systems. However, the efficiency of their application to biometric systems remains an open research problem for stronger adversarial capacities or feature representations with complex comparison functions.

Research Question 2

How can computational workload reduction be applied to improve the efficiency of FHE-based biometric identification systems?

- How can computational workload reduction for biometric identification be applied in the homomorphically encrypted domain?
- How can the trade-off between computational workload reduction and efficient encryption be optimized?

The second research question was mainly addressed in Paper B and Paper F, where the former presented specific workload reduction approach to biometric identification and the latter discusses the state-of-the-art on post-quantum secure biometric identification. In Paper B, preselection in the form of indexing was investigated and the vulnerability of unprotected preselection of a recently proposed scheme was revealed [200]. Mitigating this information leakage, Paper B applied lattice-based PEKS to achieve post-quantum protection throughout the entire identification transaction. Notably, Paper B shows that feature transformation approaches can be combined with preselection, where feature transformation approaches such as [31] can be applied within individual clusters.

Paper F gave an overview and a comparison of further approaches to workload reduction for FHE-based biometric identification systems following both the feature transformation and the preselection paradigm. Out of those approaches, the feature transformation approaches by [103] and [23] perform best among the current literature. However, their feature encoding is not trivially compatible with preselection, which could hinder further efficiency improvements. Ultimately, both feature transformation and preselection introduce a trade-off between computational efficiency any accuracy. While feature transformation approaches rely on feature dimensionality reduction to achieve efficiency [31, 103], preselection introduces an additional error in the case where a mated comparison is not included in the candidate list [95, 200].

Overall, we conclude that computational workload reduction for FHE-based biometric identification systems can be achieved both using feature transformation and preselection approaches. However, it is important to ensure the protection of all data that allows for conclusions about biometric features, such as indexing strings derived from features vectors, in addition to the FHE-protected reference database. Furthermore, feature transformation approaches that can be combined with preselection allow for reducing the workload on the selected subset of the reference database. In such cases, the trade-off between efficient encoding and workload reduction through preselection can be optimized.

Research Question 3

Can biometric information protection based on homomorphic encryption be secured against malicious adversaries in a feasible manner?

- Is it possible to secure biometric systems under the full disclosure model defined in ISO/IEC 30136 using only homomorphic encryption?
- Is it possible to efficiently secure biometric systems against malicious adversaries?

The third research question was addressed in Paper C, Paper E, and Paper F. Paper E showed that malicious security can be achieved through verifiable computation under classical security assumptions, confirming related work [28]. With post-quantum security however, the verifiable computation of the OPRF proved to be infeasible [12]. Paper C followed a different approach to verifiable computation, targeting the ISO/IEC 30136 [148] full disclosure security model. While Paper C achieved post-quantum security under this model, the chosen solution using homomorphic transciphering was not efficient. Even though Papers C and E did not achieve computational efficiency under their respective security model, computationally efficient post-quantum security against malicious adversaries can be achieved, as will be further discussed in Section 4.2.

Paper F identified that there are few BIP approaches that consider malicious security [3, 18, 66, 193, 201] in addition to Papers C and E. While all of these schemes can be considered to be post-quantum secure in the semi-honest adversary model, classically secure primitives were used to achieve security against malicious adversaries in [3]. Further schemes argue that they achieve malicious security implicitly through their construction and do not explicitly apply verifiable computation [18, 66, 193]. Finally, [201] considered the full disclosure model addressed in Paper C, where FHE was combined with a feature transformation approach that cannot be considered to withstand attacks by quantum adversaries and lowers the recognition accuracy of their overall scheme. Contrary to Paper C however, the approach presented in [201] is computationally efficient.

Overall, the security of biometric systems against malicious adversaries remains an important open research problem. In scenarios where adversaries cannot be assumed to behave semi-honestly, stored and processed biometric features need to be protected accordingly. The cryptographic model of security against malicious adversaries remains the most relevant model against these threats, as security under the full disclosure does not imply security against malicious adversaries. Under both adversary models, the efficiency of post-quantum secure solutions remains a challenge.

Research Question 4

How can biometric systems be secured against quantum adversaries?

- Which quantum adversary models need to be considered for biometric systems?
- How can the computational workload of post-quantum biometric systems be optimized?

Post-quantum security of biometric systems was the main focus of this thesis, which is reflected in all contributing papers. The main tool used in this thesis was FHE, which was used in Paper A, Paper B, Paper C, and Paper D, and discussed in the literature review presented in Paper F. Because of the underlying hardness assumptions, FHE can be assumed to be post-quantum secure with appropriate parameter choices [11]. In addition to the protection of the reference database provided by FHE, this thesis highlighted that post-quantum protection needs to be considered for additional steps of the transaction, such as preselection addressed in Paper B and homomorphic transciphering addressed in Paper C. While Paper E builds on the information-theoretic security of error-correcting codes to achieve post-quantum protection of the underlying biometric data, the additional cryptographic components in the presented protocol were also instantiated with lattice-based constructions to achieve post-quantum protection throughout the transaction.

The comprehensive literature survey given in Paper F showed that three security models need to be considered: semi-honest and malicious quantum adversaries as well as adversaries considered in the ISO/IEC 30136 full disclosure model. As discussed above, Paper F revealed that the majority of post-quantum secure BIP schemes operate under the semi-honest adversary model. In terms of the computational workload of post-quantum secure biometric systems, Paper A and Paper F showed that the choice of feature representation impacts the efficiency of single biometric comparisons, whereas feature transformation and preselection approaches can be considered for identification transactions. As discussed above, Paper C and Paper E showed two instances of infeasible post-quantum secure BIP schemes under the full disclosure model and the malicious adversary model, respectively.

Overall, we conclude that biometric systems can and should be protected with post-quantum cryptography to achieve long-term protection. Different privacy-preserving computation techniques such as FHE, MPC, and FE, can be applied to this goal. Regarding relevant adversary models, all models relevant to classically secure cryptography remain relevant for post-quantum secure BIP, where the semi-honest, malicious, and full disclosure model can be considered accord-

ing to the application scenario. As some of these cryptographic operations, in particular FHE operations, can introduce a significant computational overhead to biometric systems, this cost can be mitigated either through efficient feature representations and encoding, or workload reduction techniques applicable to biometric identification.

4.2 Limitations and Future Work

In this Section, the limitations of the contributions presented in this PhD thesis will be discussed, and opportunities for future work are outlined.

We begin with limitations concerning the security of the schemes proposed in Papers A, B, C, D, and E. Out of these schemes, only Papers C and E consider security against adversaries that are not semi-honest. While the focus of the remaining Papers A, B, and D was on different contributions, their semi-honest security can only be assumed to hold true in a controlled setting, e.g., within a governmental or corporate setting. However, adversaries in other scenarios cannot always be assumed to have compelling incentives to behave semi-honestly. This is particularly true in scenarios where the encrypted reference database is outsourced to a third-party service provider. A malicious service provider could tamper with the computation, giving false decisions of biometric transactions without detection. In addition, FHE-based BIP schemes are vulnerable to the hill-climbing attack performed by malicious clients presented by [3]. As [3] presents a solution to their attack based on *Private Information Retrieval* (PIR) [72], further research on post-quantum BIP has considered this attack to be mitigated and has focused on research on FHE-based BIP schemes without the additional PIR protection against hill-climbing attacks [44, 103]. However, the PIR protocol applied in [3] is built from classical security assumptions and does not provide post-quantum security throughout the entire transaction, indicating that further research is required to achieve the desired long-term protection throughout the biometric transactions.

Out of the two papers presented in this thesis that consider stronger adversary models, Papers C and E, only Paper E considers verifiable computation as a means to detect malicious behaviour. As discussed above, security under the malicious adversary model could only be achieved efficiently under classical security assumptions, but not based on lattices. Concretely, the cost of the verifiable computation of the lattice-based OPRF proposed by [12] was impractically high. However, this is not true for all lattice-based verifiable computation, as both [120] and [16] have recently shown efficient verifiable computation for FHE, which can be relevant to FHE-based BIP schemes. Regarding Pa-

per C, which investigated security against the ISO/IEC 30136 [148] full disclosure model, verifiable computation would not have been able to fulfil security under this assumption, as the risk of the disclosure of the FHE secret key cannot be mitigated by verifiable computation. On the other hand, the solution based on homomorphic transciphering presented in Paper C cannot be considered secure against malicious adversaries due to the absence of verifiable computation, revealing a mismatch between the two security models.

Further limitations concern the security of the schemes presented in Papers A and D, both of which involve the evaluation of conditional statements. In Paper A, the fingerprint comparison algorithm for minutiae-based representations requires several conditional statements, such as comparing the two minutiae angles against a given threshold, and proceeding further with computation on this minutiae pair if and only if the distance between the angles falls below the threshold. In Paper D, the conditional statements appear within the computation of the Manhattan distance between two encrypted keystroke dynamic feature vectors, which requires the computation of an absolute value. Even though improvements have been proposed recently [147], the computation of conditional statements under FHE is still impractical. Therefore, these statements were computed after decryption in both schemes, which may impact the privacy of the underlying biometric data. In Paper A, it can be argued that the revealed information does not allow for a reconstruction of the biometric templates. For example, the difference between two minutiae angles does not reveal the individual minutiae angles. However, evaluating the conditional statements outside of the encrypted domain cannot be considered to achieve the same privacy protection as a computation of the entire comparison under encryption. Finally, the computational efficiency of the schemes presented in Papers A and C is a limitation in terms of the practical application of these schemes.

Concluding from the aforementioned limitations, future work on the impact of verifiable FHE [16, 120] for BIP is relevant to achieve security against real-world adversaries. In addition, efficient verifiable computation techniques for post-quantum primitives need to be investigated for oblivious evaluations as applied in Paper E, which would positively impact applications outside of BIP as well. With regard to malicious security, it is interesting to note that only biometric verification has been explored under this model with post-quantum security, as Paper F revealed. The application of verifiable computation or further techniques to achieve malicious security of biometric identification is therefore an interesting open research problem. Similarly to semi-honest FHE-based solutions to biometric identification, the computation workload of the additional malicious security can be expected to require analysis and optimizations. For example, the workload of verifiable computation for identification transactions could be re-

stricted to the candidate list obtained after preselection, or applied to a random subset of comparisons such that a malicious server tampering with the computation will be detected with high probability. Overall, Papers A, C and E showed that the computational workload of post-quantum secure BIP schemes remains an open challenge.

In addition, Paper F showed that a majority of post-quantum secure BIP schemes are based on FHE, and further privacy-preserving computation techniques with post-quantum secure instantiations, such as MPC [55], FE [66], or PSI [63], have not received comparable interest. PSI in particular corresponds to minutiae-based fingerprint recognition, where two sets of minutiae are compared to reveal the minutiae common to both fingerprints. This technique has been applied to fingerprint comparisons [257], but not with post-quantum security. An interesting open problem is the application of these techniques for further biometric modalities, both physiological and behavioral. In addition to the application of further cryptographic techniques, biometric identification could be investigated for a larger variety of biometric modalities. As Paper F revealed, biometric identification with post-quantum security was almost exclusively explored for face recognition. While such constructions can be assumed to translate to other fixed-length feature representations, an interesting open problem is the investigation of the security and efficiency of biometric identification for variable-length feature representations under post-quantum security.

Finally, future work on research areas closely related to BIP can be expected to impact the latter positively. One prominent research area is deep-learning based feature representations, which have become state-of-the-art for face recognition [191], but have also been explored for fingerprint [104] and iris [202]. The current limitation of such feature representations for other modalities than face recognition are their lower biometric performance compared to feature representations that are not based on deep learning, such as the Daugman IrisCode representation [80]. As the biometric performance of these feature extractors increases with further research, their protection under FHE and further schemes that allow for computations on encrypted data becomes computationally feasible, and therefore relevant to BIP. In addition to deep-learning based feature representations, the research field of secure machine learning can similarly impact future work on BIP. As Paper F motivates, machine learning under encryption can be applied to perform deep-learning based feature extraction in the encrypted domain, which allows for an extension of the protection of the underlying biometric features. In state-of-the-art BIP schemes today [44, 103], feature extraction is computed on unencrypted samples, e.g., face images, and only the extracted features are encrypted. Future work can improve the security of this computation through a direct extraction of encrypted feature representations.

References

[1] S. M. Abdullahi, S. Sun, B. Wang, N. Wei, and H. Wang, "Biometric template attacks and recent protection mechanisms: A survey," *Information Fusion*, vol. 103, p. 102144, 2024.

[2] A. Abidin, "On privacy-preserving biometric authentication," in *Proc. Intl. Conf. Information Security and Cryptology*. Springer, 2017, pp. 169–186.

[3] A. Abidin and A. Mitrokotsa, "Security aspects of privacy-preserving biometric authentication based on ideal lattices and Ring-LWE," in *IEEE International Workshop on Information Forensics and Security (WIFS)*. IEEE, 2014, pp. 60–65.

[4] B. Abinaya and S. Santhi, "A survey on genomic data by privacy-preserving techniques perspective," *Computational Biology and Chemistry*, vol. 93, p. 107538, 2021.

[5] A. Acar, H. Aksu, A. S. Uluagac, and M. Conti, "A survey on homomorphic encryption schemes: Theory and implementation," *ACM Computing Surveys*, vol. 51, no. 4, pp. 1–35, 2018.

[6] A. Acar, W. Liu, R. Beyah, K. Akkaya, and A. S. Uluagac, "A privacy-preserving multifactor authentication system," *Security and Privacy*, vol. 2, no. 5, p. e88, 2019.

[7] A. Adler, R. Youmaran, and S. Loyka, "Towards a measure of biometric information," in *Canadian Conference on Electrical and Computer Engineering*. IEEE, February 2006, pp. 210–213.

[8] M. Ajtai, "Generating hard instances of lattice problems," in *Proc. Ann. ACM Symposium on Theory of Computing*, 1996, pp. 99–108.

[9] M. B. Akanbi, R. G. Jimoh, and J. B. Awotunde, "Biocryptosystems for template protection: A survey of fuzzy vault," in *Proc. Information Technology for Education and Development*, 2022, pp. 1–6.

[10] G. Alagic, D. Apon, D. Cooper, Q. Dang, T. Dang, J. Kelsey, J. Lichtinger, C. Miller, D. Moody, R. Peralta, R. Perlner, A. Robinson, and D. Smith-Tone, "Status report on the third round of the NIST post-quantum cryptography standardization process," *US Department of Commerce, NIST*, 2022.

[11] M. Albrecht, M. Chase, H. Chen, J. Ding, S. Goldwasser, S. Gorbunov, S. Halevi, J. Hoffstein, K. Laine, K. Lauter, S. Lokam, D. Micciancio, D. Moody, T. Morrison, A. Sahai, and V. Vaikuntanathan, "Homomorphic encryption security standard," HomomorphicEncryption.org, Toronto, Canada, Tech. Rep., November 2018.

[12] M. R. Albrecht, A. Davidson, A. Deo, and N. P. Smart, "Round-optimal verifiable oblivious pseudorandom functions from ideal lattices," in *IACR International Conference on Public-Key Cryptography*. Springer, 2021, pp. 261–289.

[13] M. R. Albrecht, R. Player, and S. Scott, "On the concrete hardness of learning with errors," Cryptology ePrint Archive, Paper 2015/046, 2015, https://eprint.iacr.org/2015/046. [Online]. Available: https://eprint.iacr.org/2015/046

[14] M. Albrecht, D. J. Bernstein, T. Chou, C. Cid, J. Gilcher, T. Lange, V. Maram, I. v. Maurich, R. Misoczki, R. Niederhagen, K. G. Paterson, E. Persichetti, C. Peters, P. Schwabe, N. Sendrier, J. Szefer, C. J. Tjhai, M. Tomlinson, and W. Wang, "Classic McEliece: conservative code-based cryptography – Round 4 submission," *NIST PQC Competition Round 4*, pp. 1–16, 2022, https://classic.mceliece.org/.

[15] Apple Inc., "About Face ID advanced technology," 2022. [Online]. Available: https://support.apple.com/en-us/HT208108

[16] D. F. Aranha, A. Costache, A. Guimarães, and E. Soria-Vazquez, "Heliopolis: Verifiable computation over homomorphically encrypted data from interactive oracle proofs is practical," *Cryptology ePrint Archive*, 2023.

[17] L. C. F. Araújo, L. H. R. Sucupira, M. G. Lizarraga, L. L. Ling, and J. B. T. Yabu-Uti, "User authentication through typing biometrics features," *IEEE Trans. on Signal Processing*, vol. 53, no. 2, pp. 851–855, 2005.

[18] R. Arjona and I. Baturone, "A post-quantum biometric template protection scheme based on learning parity with noise (LPN) commitments," *IEEE Access*, vol. 8, pp. 182 355–182 365, 2020.

[19] R. Arjona, P. López-González, R. Román, and I. Baturone, "Post-quantum biometric authentication based on homomorphic encryption and classic mceliece," *Applied Sciences*, vol. 13, no. 2, p. 757, 2023.

[20] R. Arjona, M. A. Prada-Delgado, I. Baturone, and A. Ross, "Securing minutia cylinder codes for fingerprints through physically unclonable functions: An exploratory study," in *2018 International Conference on Biometrics (ICB)*. IEEE, 2018, pp. 54–60.

[21] R. Avanzi, J. Bos, L. Ducas, E. Kiltz, T. Lepoint, V. Lyubashevsky, J. M. Schanck, P. Schwabe, G. Seiler, and D. Stehlé, "Crystals-kyber algorithm specifications and supporting documentation," *NIST PQC Round*, vol. 3, pp. 1–43, 2021.

[22] A. A. Badawi, J. Bates, F. Bergamaschi, D. B. Cousins, S. Erabelli, N. Genise, S. Halevi, H. Hunt, A. Kim, Y. Lee, Z. Liu, D. Micciancio, I. Quah, Y. Polyakov, S. R.V., K. Rohloff, J. Saylor, D. Suponitsky, M. Triplett, V. Vaikuntanathan, and V. Zucca, "OpenFHE: open-source fully homomorphic encryption library," Cryptology ePrint Archive, Paper 2022/915, 2022, https://eprint.iacr.org/2022/915. [Online]. Available: https://eprint.iacr.org/2022/915

[23] J. Bai, X. Zhang, X. Song, H. Shao, Q. Wang, S. Cui, and G. Russello, "Cryptomask: Privacy-preserving face recognition," in *International Conference on Information and Communications Security*. Springer, 2023, pp. 333–350.

[24] A. Banerjee, C. Peikert, and A. Rosen, "Pseudorandom functions and lattices," in *Ann. Intl. Conf. on the Theory and Applications of Cryptographic Techniques*. Springer, 2012, pp. 719–737.

[25] E. Barker, *Digital Signature Standard (DSS)*. Federal Information Processing Standard (NIST FIPS), National Institute of Standards and Technology, Gaithersburg, MD, 2013-07-19 2013.

[26] M. Barni, T. Bianchi, D. Catalano, M. Di Raimondo, R. Labati *et al.*, "A privacy-compliant fingerprint recognition system based on homomorphic encryption and fingercode templates," in *IEEE Intl. Conf. on Biometrics: Theory Applications and Systems (BTAS)*. IEEE, 2010, pp. 1–7.

[27] M. Barni, G. Droandi, and R. Lazzeretti, "Privacy protection in biometric-based recognition systems: A marriage between cryptography and sig-

nal processing," *IEEE Signal Processing Magazine*, vol. 32, no. 5, pp. 66–76, 2015.

[28] A. Bassit, F. Hahn, J. Peeters, T. Kevenaar, R. Veldhuis, and A. Peter, "Fast and accurate likelihood ratio-based biometric verification secure against malicious adversaries," *IEEE Trans. on Information Forensics and Security (TIFS)*, vol. 16, pp. 5045–5060, 2021.

[29] A. Bassit, F. Hahn, R. Veldhuis, and A. Peter, "Hybrid biometric template protection: Resolving the agony of choice between bloom filters and homomorphic encryption," *IET Biometrics*, vol. 11, pp. 430–444, 2022.

[30] P. Bauspieß, J. Kolberg, D. Demmler, J. Krämer, and C. Busch, "Post-quantum secure two-party computation for iris biometric template protection," in *Proc. IEEE Workshop on Information Forensics and Security (WIFS)*, 2020, pp. 1–6.

[31] P. Bauspieß, J. Olafsson, J. Kolberg, P. Drozdowski, C. Rathgeb, and C. Busch, "Improved homomorphically encrypted biometric identification using coefficient packing," in *Proc. Intl. Workshop on Biometrics and Forensics (IWBF)*, 2022.

[32] P. Bauspieß, P. Bours, C. Rathgeb, and C. Busch, "Type^2: A secure and seamless biometric two-factor authentication protocol using keystroke dynamics," in *Norwegian Information Security Conference*, 2023, pp. 1–16.

[33] P. Bauspieß, J. Kolberg, P. Drozdowski, C. Rathgeb, and C. Busch, "Privacy-preserving preselection for protected biometric identification using public-key encryption with keyword search," *IEEE Transactions on Industrial Informatics*, 2022.

[34] P. Bauspieß, C.-M. Zok, A. Costache, C. Rathgeb, J. Kolberg, and C. Busch, "MT-PRO: Multibiometric template protection based on homomorphic transciphering," in *2023 IEEE International Workshop on Information Forensics and Security (WIFS)*. IEEE, 2023, pp. 1–6.

[35] P. Bauspieß, "Post-quantum secure biometric systems: An overview," *under review*, 2024.

[36] P. Bauspieß, M. Grimmer, C. Fougner, D. L. Vasseur, T. T. Stöcklin, C. Rathgeb, J. Kolberg, A. Costache, and C. Busch, "HEBI: Homomorphically encrypted biometric indexing," in *Proc. Intl. Joint Conf. on Biometrics (IJCB)*, September 2023, pp. 1–10.

[37] P. Bauspieß, T. Silde, M. Poljuha, A. Tullot, A. Costache, C. Rathgeb, J. Kolberg, and C. Busch, "BRAKE: Biometric resilient authenticated key exchange," *IEEE Access*, 2024.

[38] P. Bauspieß, L. Vad, H. Myrekrok, A. Costache, J. Kolberg, C. Rathgeb, and C. Busch, "On the feasibility of fully homomorphic encryption of minutiae-based fingerprint representations," in *9th Intl. Conf. on Information Systems Security and Privacy ICISSP*, February 2023, pp. 462–470.

[39] R. Behnia, A. A. Yavuz, and M. O. Ozmen, "High-speed high-security public key encryption with keyword search," in *IFIP Ann. Conf. on Data and Applications Security and Privacy*. Springer, 2017, pp. 365–385.

[40] D. J. Bernstein, T. Chou, T. Lange, I. von Maurich, R. Misoczki, R. Niederhagen, E. Persichetti, C. Peters, P. Schwabe, N. Sendrier *et al.*, "Classic mceliece: conservative code-based cryptography," *US Department of Commerce, NIST*, pp. 1–25, 2017.

[41] M. Blanton and M. Aliasgari, "Secure outsourced computation of iris matching," *Journal of Computer Security*, vol. 20, no. 2-3, pp. 259–305, 2012.

[42] S. Bleha, C. Slivinsky, and B. Hussien, "Computer-access security systems using keystroke dynamics," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 12, no. 12, pp. 1217–1222, 1990.

[43] B. H. Bloom, "Space/time trade-offs in hash coding with allowable errors," *Communications of the ACM*, vol. 13, no. 7, pp. 422–426, 1970.

[44] V. N. Boddeti, "Secure face matching using fully homomorphic encryption," in *Proc. Intl. Conf. on Biometrics Theory, Applications and Systems (BTAS)*. IEEE, 2018, pp. 1–10.

[45] D. Boneh, G. Di Crescenzo, R. Ostrovsky, and G. Persiano, "Public key encryption with keyword search," in *Intl. Conf. on the Theory and Applications of Cryptographic Techniques*. Springer, 2004, pp. 506–522.

[46] D. Boneh and M. Franklin, "Identity-based encryption from the Weil pairing," in *Proc. Annual Intl. Cryptology Conf.* Springer, 2001, pp. 213–229.

[47] D. Boneh, A. Sahai, and B. Waters, "Functional encryption: Definitions and challenges," in *Proc. Theory of Cryptography Conference*. Springer, 2011, pp. 253–273.

[48] J. Bootle, S. Faller, J. Hesse, K. Hostáková, and J. Ottenhues, "Generalized fuzzy password-authenticated key exchange from error correcting codes," in *International Conference on the Theory and Application of Cryptology and Information Security*. Springer, 2023, pp. 110–142.

[49] J. Bos, L. Ducas, E. Kiltz, T. Lepoint, V. Lyubashevsky, J. M. Schanck, P. Schwabe, G. Seiler, and D. Stehlé, "CRYSTALS-Kyber: a CCA-secure module-lattice-based KEM," in *European Symposium on Security and Privacy (EuroS&P)*. IEEE, 2018, pp. 353–367.

[50] Z. Brakerski, "Fully homomorphic encryption without modulus switching from classical GapSVP," in *Annual Cryptology Conference*. Springer, 2012, pp. 868–886.

[51] Z. Brakerski, C. Gentry, and V. Vaikuntanathan, "(Leveled) fully homomorphic encryption without bootstrapping," *ACM Trans. on Computation Theory (TOCT)*, vol. 6, no. 3, pp. 1–36, 2014.

[52] Z. Brakerski and V. Vaikuntanathan, "Fully homomorphic encryption from ring-lwe and security for key dependent messages," in *Annual Cryptology Conference*. Springer, 2011, pp. 505–524.

[53] L. Brenna, I. S. Singh, H. D. Johansen, and D. Johansen, "TFHE-rs: A library for safe and secure remote computing using fully homomorphic encryption and trusted execution environments," *Array*, vol. 13, p. 100118, 2022.

[54] J. Bringer, H. Chabanne, and A. Patey, "Privacy-preserving biometric identification using secure multiparty computation: An overview and recent trends," *IEEE Signal Processing Magazine*, vol. 30, no. 2, pp. 42–52, 2013.

[55] N. Büscher, D. Demmler, N. P. Karvelas, S. Katzenbeisser, J. Krämer, D. Rathee, T. Schneider, and P. Struck, "Secure two-party computation in a quantum world," in *Proc. Int. Conf. on Applied Cryptography and Network Security (ACNS)*, ser. LNCS. Springer, 2020.

[56] R. Canetti, "Universally composable security: A new paradigm for cryptographic protocols," in *Proc. IEEE Symposium on Foundations of Computer Science*. IEEE, 2001, pp. 136–145.

[57] A. Canteaut, S. Carpov, C. Fontaine, J. Fournier, B. Lac, M. Naya-Plasencia, R. Sirdey, and A. Tria, "End-to-end data security for IoT: from a cloud of encryptions to encryption in the cloud," in *Cesar Conf.*, 2017.

[58] R. Cappelli, M. Ferrara, and D. Maltoni, "Minutia cylinder-code: A new representation and matching technique for fingerprint recognition," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, March 2010.

[59] R. Cappelli, A. Lumini, D. Maio, and D. Maltoni, "Fingerprint image reconstruction from standard templates," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 29, no. 9, September 2007.

[60] S. Casacuberta, J. Hesse, and A. Lehmann, "SoK: Oblivious pseudorandom functions," in *2022 IEEE 7th European Symposium on Security and Privacy (EuroS&P)*. IEEE, 2022, pp. 625–646.

[61] A. Cavoukian and A. Stoianov, "Biometric encryption," *Biometric Technology Today*, vol. 15, no. 3, p. 11, 2007.

[62] D. Chaum and T. P. Pedersen, "Wallet databases with observers," in *Annual International Cryptology Conference*. Springer, 1992, pp. 89–105.

[63] H. Chen, K. Laine, and P. Rindal, "Fast private set intersection from homomorphic encryption," in *Proc. ACM SIGSAC Conference on Computer and Communications Security*, 2017, pp. 1243–1255.

[64] J. H. Cheon, H. Chung, M. Kim, and K.-W. Lee, "Ghostshell: Secure biometric authentication using integrity-based homomorphic evaluations," *Cryptology ePrint Archive*, 2016.

[65] J. H. Cheon, A. Kim, M. Kim, and Y. Song, "Homomorphic encryption for arithmetic of approximate numbers," in *Intl. Conf. on the Theory and Appl. of Crypt. and Information Security*. Springer, 2016, pp. 409–437.

[66] J. H. Cheon, D. Kim, D. Kim, J. Lee, J. Shin, and Y. Song, "Lattice-based secure biometric authentication for hamming distance," in *Proc. Australasian Conf. on Information Security and Privacy (ACISP)*. Springer, 2021, pp. 653–672.

[67] I. Chillotti, N. Gama, M. Georgieva, and M. Izabachene, "Faster fully homomorphic encryption: Bootstrapping in less than 0.1 seconds," in *Proc. Intl. Conf. on the Theory and Application of Cryptology and Information Security*. Springer, 2016, pp. 3–33.

[68] I. Chillotti, N. Gama, M. Georgieva, and M. Izabachène, "TFHE: Fast fully homomorphic encryption library," 2016, https://tfhe.github.io/tfhe/.

[69] I. Chillotti, N. Gama, M. Georgieva, and M. Izabachène, "TFHE: fast fully homomorphic encryption over the torus," *Journal of Cryptology*, vol. 33, no. 1, pp. 34–91, 2020.

[70] Chinese Academy of Sciences Institute of Automation, "CASIA Iris Thousand database," available at http://biometrics.idealtest.org/.

[71] J. Cho, J. Ha, S. Kim, B. Lee, J. Lee, J. Lee, D. Moon, and H. Yoon, "Transciphering framework for approximate homomorphic encryption," in *Intl. Conf. on the Theory and Application of Cryptology and Information Security*. Springer, 2021, pp. 640–669.

[72] B. Chor, E. Kushilevitz, O. Goldreich, and M. Sudan, "Private information retrieval," *Journal of the ACM*, vol. 45, no. 6, pp. 965–981, 1998.

[73] T. C. Clancy, N. Kiyavash, and D. J. Lin, "Secure smartcard-based fingerprint authentication," in *Proceedings of the ACM SIGMM Workshop on Biometrics Methods and Applications*, 2003, pp. 45–52.

[74] R. Cramer, I. B. Damgård, and J. B. Nielsen, *Secure multiparty computation and secret sharing*. Cambridge University Press, 2015.

[75] G. Crihan, M. Crăciun, and L. Dumitriu, "A comparative assessment of homomorphic encryption algorithms applied to biometric information," *Inventions*, vol. 8, no. 4, p. 102, 2023.

[76] I. Csiszár, "I-divergence geometry of probability distributions and minimization problems," *The Annals of Probability*, pp. 146–158, 1975.

[77] J. Daemen and V. Rijmen, "AES proposal: Rijndael," 1999.

[78] D.-T. Dam, T.-H. Tran, V.-P. Hoang, C.-K. Pham, and T.-T. Hoang, "A survey of post-quantum cryptography: Start of a new race," *Cryptography*, vol. 7, no. 3, p. 40, 2023.

[79] I. Damgard, M. Geisler, and M. Kroigard, "Homomorphic encryption and secure comparison," *International Journal of Applied Cryptography*, vol. 1, no. 1, pp. 22–31, 2008.

[80] J. Daugman, "How iris recognition works," *IEEE Trans. on Circuits and Systems for Video Technology (TCSVT)*, vol. 14, no. 1, pp. 21–30, 2004.

[81] J. Daugman, "Probing the uniqueness and randomness of iriscodes: Results from 200 billion iris pair comparisons," *Proceedings of the IEEE*, vol. 94, no. 11, pp. 1927–1935, 2006.

[82] A. De Caro, "Java lattice based cryptography library," :http://gas.dia.unisa.it/projects/jlbc/index.html.

[83] S. K. Debnath, N. Kundu, and T. Choudhury, "Efficient post-quantum private set-intersection protocol," *International Journal of Information and Computer Security*, vol. 17, no. 3-4, pp. 405–423, 2022.

[84] P. Delgado-Santos, R. Tolosana, R. Guest, R. Vera-Rodriguez, F. Deravi, and A. Morales, "Gaitprivacyon: Privacy-preserving mobile gait biomet-

rics using unsupervised learning," *Pattern Recognition Letters*, vol. 161, pp. 30–37, 2022.

[85] D. Demmler, T. Schneider, and M. Zohner, "ABY - a framework for efficient mixed-protocol secure two-party computation." in *Network and Distributed System Security Symposium*, 2015.

[86] J. Deng, J. Guo, and S. Zafeiriou, "ArcFace: Additive angular margin loss for deep face recognition," in *Conf. on Computer Vision and Pattern Recognition (CVPR)*, June 2019.

[87] W. Diffie and M. E. Hellman, "New directions in cryptography," *IEEE Transactions on Information Theory*, vol. 22, no. 6, pp. 644–654, November 1976.

[88] T. V. T. Doan, M.-L. Messai, G. Gavin, and J. Darmont, "A survey on implementations of homomorphic encryption schemes," *The Journal of Supercomputing*, pp. 1–42, 2023.

[89] Y. Dodis, L. Reyzin, and A. Smith, "Fuzzy extractors: How to generate strong keys from biometrics and other noisy data," in *International Conference on the Theory and Applications of Cryptographic Techniques*. Springer, 2004, pp. 523–540.

[90] X. Dong, S. Kim, Z. Jin, J. Y. Hwang, S. Cho, and A. B. J. Teoh, "Open-set face identification with index-of-max hashing by learning," *Pattern Recognition*, vol. 103, p. 107277, 2020.

[91] P. Drozdowski, N. Buchmann, C. Rathgeb, M. Margraf, and C. Busch, "On the application of homomorphic encryption to face identification," in *Intl. Conf. of the Biometrics Special Interest Group (BIOSIG)*. Gesellschaft für Informatik e.V., September 2019, pp. 173–180.

[92] P. Drozdowski, C. Rathgeb, and C. Busch, "Computational workload in biometric identification systems: An overview," *IET Biometrics*, vol. 8, no. 6, pp. 351–368, November 2019.

[93] P. Drozdowski, C. Rathgeb, and C. Busch, "Turning a vulnerability into an asset: Accelerating facial identification with morphing," in *Intl. Conf. on Acoustics, Speech, and Signal Processing (ICASSP)*. IEEE, May 2019, pp. 2582–2586.

[94] P. Drozdowski, C. Rathgeb, B.-A. Mokroß, and C. Busch, "Multi-biometric identification with cascading database filtering," *IEEE Trans.on Biometrics, Behavior, and Identity Science*, vol. 2, no. 3, pp. 210–222, July 2020.

[95] P. Drozdowski, F. Stockhardt, C. Rathgeb, D. Osorio-Roig, and C. Busch, "Feature fusion methods for indexing and retrieval of biometric data: Application to face recognition with privacy protection," *IEEE Access*, vol. 9, pp. 139 361–139 378, October 2021.

[96] P. Drozdowski, F. Struck, C. Rathgeb, and C. Busch, "Benchmarking binarisation schemes for deep face templates," in *Intl. Conf. on Image Processing (ICIP)*. IEEE, October 2018, pp. 191–195.

[97] L. Ducas, V. Lyubashevsky, and T. Prest, "Efficient identity-based encryption over NTRU lattices," in *Intl. Conf. on the Theory and Application of Cryptology and Information Security*. Springer, 2014, pp. 22–41.

[98] L. Ducas and D. Micciancio, "FHEW: Bootstrapping homomorphic encryption in less than a second," in *Proc. Ann. Intl. Conf. on the Theory and Applications of Cryptographic Techniques*. Springer, 2015, pp. 617–640.

[99] R. O. Duda, P. E. Hart, and D. Stork, "Pattern classification," *John Wiely and Sons*, 2001.

[100] P.-A. Dupont, J. Hesse, D. Pointcheval, L. Reyzin, and S. Yakoubov, "Fuzzy password-authenticated key exchange," in *International Conference on the Theory and Applications of Cryptographic Techniques*. Springer, 2018, pp. 393–424.

[101] J.-P. D'Anvers, A. Karmakar, S. Sinha Roy, and F. Vercauteren, "Saber: Module-LWR based key exchange, CPA-secure encryption and CCA-secure kem," in *Proc. Intl. Conf. on Cryptology in Africa*. Springer, 2018, pp. 282–305.

[102] T. Elgamal, "A public key cryptosystem and a signature scheme based on discrete logarithms," *IEEE Trans. on Information Theory*, vol. 31, no. 4, pp. 469–472, 1985.

[103] J. J. Engelsma, A. K. Jain, and V. N. Boddeti, "HERS: Homomorphically encrypted representation search," *IEEE Trans. on Biometrics, Behavior, and Identity Science (T-BIOM)*, 2022.

[104] J. J. Engelsma, K. Cao, and A. K. Jain, "Learning a fixed-length fingerprint representation," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 43, no. 6, pp. 1981–1997, 2019.

[105] J. Ernst and A. Mitrokotsa, "A framework for uc secure privacy preserving biometric authentication using efficient functional encryption," in *International Conference on Applied Cryptography and Network Security*. Springer, 2023, pp. 167–196.

[106] A. Erwig, J. Hesse, M. Orlt, and S. Riahi, "Fuzzy asymmetric password-authenticated key exchange," in *International Conference on the Theory and Application of Cryptology and Information Security*. Springer, 2020, pp. 761–784.

[107] EU Parliament, *EU Quantum Manifesto: A New Era of Technology*, 2016.

[108] European Council, "Regulation of the european parliament and of the council on electronic identification and trust services for electronic transactions in the internal market (eidas regulation)," July 2014.

[109] European Council, "Regulation of the European Parliament and of the Council on the protection of individuals with regard to the processing of personal data and on the free movement of such data (General Data Protection Regulation)," April 2016.

[110] European Data Protection Supervisor, *Report on logging to the SIS II at national level*, 2018.

[111] A. Everspaugh, R. Chaterjee, S. Scott, A. Juels, and T. Ristenpart, "The pythia PRF service," in *USENIX Security Symposium*, 2015, pp. 547–562.

[112] J. Fan and F. Vercauteren, "Somewhat practical fully homomorphic encryption," *Cryptology ePrint Archive*, 2012.

[113] A. Fiat and A. Shamir, "How to prove yourself: Practical solutions to identification and signature problems," in *Conference on the Theory and Application of Cryptographic Techniques*. Springer, 1986, pp. 186–194.

[114] W. Ford and B. S. Kaliski, "Server-assisted generation of a strong secret from a password," in *Proceedings IEEE International Workshops on Enabling Technologies: Infrastructure for Collaborative Enterprises (WET ICE)*. IEEE, 2000, pp. 176–180.

[115] S. Foundation, "Technical information - specifications and software libraries for developers," 2022, https://signal.org/docs/.

[116] M. J. Freedman, Y. Ishai, B. Pinkas, and O. Reingold, "Keyword search and oblivious pseudorandom functions," in *Theory of Cryptography Conference*. Springer, 2005, pp. 303–324.

[117] FRONTEX, "Best practice technical guidelines for automated border control ABC systems," 2015.

[118] E. Fujisaki, T. Okamoto, D. Pointcheval, and J. Stern, "RSA-OAEP is secure under the RSA assumption," in *Annual International Cryptology Conference*. Springer, 2001, pp. 260–274.

[119] J. Galbally, A. Ross, M. Gomez-Barrero, J. Fierrez, and J. Ortega-Garcia, "Iris image reconstruction from binary templates: An efficient probabilistic approach based on genetic algorithms," *Computer Vision and Image Understanding*, vol. 117, no. 10, pp. 1512–1525, 2013.

[120] S. Garg, A. Goel, and M. Wang, "How to prove statements obliviously?" *Cryptology ePrint Archive*, 2023.

[121] C. Gentry, "Fully homomorphic encryption using ideal lattices," in *Proc. Ann. ACM Symposium on Theory of Computing*, 2009, pp. 169–178.

[122] C. Gentry and S. Halevi, "Implementing Gentry's fully-homomorphic encryption scheme," in *Ann. Intl. Conf. on the Theory and Applications of Cryptographic Techniques*. Springer, 2011, pp. 129–148.

[123] R. Gilad-Bachrach, N. Dowlin, K. Laine, K. Lauter, M. Naehrig, and J. Wernsing, "Cryptonets: Applying neural networks to encrypted data with high throughput and accuracy," in *International Conference on Machine Learning*. PMLR, 2016, pp. 201–210.

[124] B. P. Gilkalaye and R. Derakhshani, "Secure authentication using a garbled circuit variant for arithmetic circuits," in *IEEE Intl. Symposium on Technologies for Homeland Security*. IEEE, 2021, pp. 1–7.

[125] M. Gomez-Barrero, J. Galbally, A. Morales, and J. Fierrez, "Privacy-preserving comparison of variable-length data with application to biometric template protection," *IEEE Access*, vol. 5, no. 1, pp. 8606–8619, December 2017.

[126] M. Gomez-Barrero, E. Maiorana, J. Galbally, P. Campisi, and J. Fierrez, "Multi-biometric template protection based on Homomorphic Encryption," *Pattern Recognition*, vol. 67, pp. 149–163, July 2017.

[127] M. Gomez-Barrero, J. Galbally, C. Rathgeb, and C. Busch, "General framework to evaluate unlinkability in biometric template protection systems," *IEEE Transactions on Information Forensics and Security*, vol. 13, no. 6, pp. 1406–1420, 2017.

[128] M. Gomez-Barrero, C. Rathgeb, G. Li, R. Ramachandra, J. Galbally, and C. Busch, "Multi-biometric template protection based on bloom filters," *Information Fusion*, vol. 42, pp. 37–50, 2018.

[129] D. González-Jiménez, F. Pérez-González, P. Comesana-Alfaro, L. Pérez-Freire, and J. L. Alba-Castro, "Modeling gabor coefficients via generalized gaussian distributions for face recognition," in *Proc. IEEE International Conference on Image Processing*, vol. 4. IEEE, 2007, pp. IV–485.

[130] V. D. Goppa, "A new class of linear correcting codes," *Problemy Peredachi Informatsii*, vol. 6, no. 3, pp. 24–30, 1970.

[131] M. Grimmer, R. Ramachandra, and C. Busch, "Deep face age progression: A survey," *IEEE Access*, vol. 9, pp. 83 376–83 393, 2021.

[132] L. K. Grover, "A fast quantum mechanical algorithm for database search," in *Proc. ACM Symposium on Theory of Computing*, 1996, pp. 212–219.

[133] V. Guruswami and M. Sudan, "Improved decoding of Reed-Solomon and algebraic-geometric codes," in *Proceedings Annual Symposium on Foundations of Computer Science*. IEEE, 1998, pp. 28–37.

[134] V. K. Hahn and S. Marcel, "Biometric template protection for neural-network-based face recognition systems: A survey of methods and evaluation techniques," *IEEE Transactions on Information Forensics and Security*, vol. 18, pp. 639–666, 2022.

[135] S. Haider, A. Abbas, and A. K. Zaidi, "A multi-technique approach for user identification through keystroke dynamics," in *Proceedings of the IEEE Intl. Conference on Systems, Man and Cybernetics (SMC)*, vol. 2. IEEE, 2000, pp. 1336–1341.

[136] S. Halevi and V. Shoup, "HElib - an implementation of homomorphic encryption," 2014, homenc.github.io/HElib.

[137] J. Hämmerle-Uhl, G. Penn, G. Pötzelsberger, and A. Uhl, "Size-reduction strategies for iris codes," *Intl. Journal of Computer and Information Engineering*, vol. 9, no. 1, pp. 290–293, 2015.

[138] S. C. C. Han, D. H. Han, and H. Kim, "Web-based keystroke dynamics identity verification using neural network," *Journal of Organizational Computing and Electronic Commerce (JOCEC)*, vol. 10, no. 4, pp. 295–307, 2000.

[139] Y. Han, C. Xu, S. Li, C. Jiang, and K. Chen, "ttPAKE: Typo tolerance password-authenticated key exchange," *Journal of Information Security and Applications*, vol. 79, p. 103658, 2023.

[140] J. Hoffstein, J. Pipher, and J. H. Silverman, "NTRU: a ring-based public key cryptosystem," in *Intl. Algorithmic Number Theory Symposium*. Springer, 1998, pp. 267–288.

[141] C.-Y. Hsu, C.-S. Lu, and S.-C. Pei, "Image feature extraction in encrypted domain with privacy-preserving sift," *IEEE Transactions on Image Processing*, vol. 21, no. 11, pp. 4593–4607, 2012.

[142] H. Huang and L. Wang, "Efficient privacy-preserving face verification scheme," *Journal of Information Security and Applications*, vol. 63, p. 103055, 2021.

[143] H. Huang and L. Wang, "Efficient privacy-preserving face identification protocol," *IEEE Transactions on Services Computing*, 2023.

[144] Z. Huang, W.-j. Lu, C. Hong, and J. Ding, "Cheetah: Lean and fast secure two-party deep neural network inference," in *31st USENIX Security Symposium (USENIX Security 22)*, 2022, pp. 809–826.

[145] A. Ibarrondo, H. Chabanne, V. Despiegel, and M. Önen, "Grote: Group testing for privacy-preserving face identification," in *Proc. ACM Conference on Data and Application Security and Privacy*, 2023, pp. 117–128.

[146] A. Ibarrondo and A. Viand, "Pyfhel: Python for homomorphic encryption libraries," in *Proc. Workshop on Encrypted Computing & Applied Homomorphic Cryptography*, 2021, pp. 11–16.

[147] I. Iliashenko and V. Zucca, "Faster homomorphic comparison operations for BGV and BFV," *Proceedings of Privacy Enhancing Technologies Symposium*, vol. 2021, no. 3, pp. 246–264, 2021.

[148] ISO/IEC JTC 1/SC 37 Biometrics, *ISO/IEC 30136:2018. Information technology — Performance testing of biometric template protection schemes*, International Organization for Standardization, 2018.

[149] ISO/IEC JTC1 SC27 Security Techniques, *ISO/IEC 24745:2022. Information Technology - Security Techniques - Biometric Information Protection*, International Organization for Standardization, 2022.

[150] ISO/IEC JTC1 SC37 Biometrics, *ISO/IEC 19794-1:2011 Information Technology - Biometric Data Interchange Formats - Part 1: Framework*, International Organization for Standardization, June 2011.

[151] ISO/IEC JTC1 SC37 Biometrics, *ISO/IEC 19794-2:2011 Information Technology - Biometric Data Interchange Formats - Part 2: Finger Minutiae Data*, International Organization for Standardization, June 2011.

[152] ISO/IEC JTC1 SC37 Biometrics, *ISO/IEC 19795-1:2021. Information Technology – Biometric Performance Testing and Reporting – Part 1: Principles and Framework*, Intl. Org. for Standardization, June 2021.

[153] A. Jain, S. Prabhakar, L. Hong, and S. Pankanti, "Filterbank-based fingerprint matching," *IEEE Trans. on Image Processing*, vol. 9, no. 5, pp. 846–859, 2000.

[154] A. K. Jain, D. Deb, and J. J. Engelsma, "Biometrics: Trust, but verify," *IEEE Transactions on Biometrics, Behavior, and Identity Science*, vol. 4, no. 3, pp. 303–323, 2021.

[155] S. Jarecki, H. Krawczyk, and J. Xu, "OPAQUE: an asymmetric pake protocol secure against pre-computation attacks," in *International Conference on the Theory and Applications of Cryptographic Techniques*. Springer, 2018, pp. 456–486.

[156] X. Jiang and W.-Y. Yau, "Fingerprint minutiae matching based on the local and global structures," in *Proceedings International Conference on Pattern Recognition (ICPR)*, vol. 2. IEEE, 2000, pp. 1038–1041.

[157] A. K. Jindal, I. Shaik, V. Vasudha, S. R. Chalamala, R. Ma, and S. Lodha, "Secure and privacy preserving method for biometric template protection using fully homomorphic encryption," in *IEEE Intl. Conf. on Trust, sScurity and Privacy in Computing and Communications (TrustCom)*. IEEE, 2020, pp. 1127–1134.

[158] I. Joshi, M. Grimmer, C. Rathgeb, C. Busch, F. Bremond, and A. Dantcheva, "Synthetic data in human analysis: A survey," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024.

[159] R. Joyce and G. Gupta, "Identity authentication based on keystroke latencies," *Communications of the ACM*, vol. 33, no. 2, pp. 168–176, 1990.

[160] A. Juels and M. Sudan, "A fuzzy vault scheme," *Designs, Codes and Cryptography*, vol. 38, no. 2, pp. 237–257, 2006.

[161] A. Juels and M. Wattenberg, "A fuzzy commitment scheme," in *Proc. ACM Conference on Computer and Communications Security*, 1999, pp. 28–36.

[162] P. Kang, S.-s. Hwang, and S. Cho, "Continual retraining of keystroke dynamics based authenticator," in *Proceedings of the 2nd International Conference on Biometrics (ICB)*. Springer, 2007, pp. 1203–1211.

[163] T. Karras, M. Aittala, S. Laine, E. Härkönen, J. Hellsten, J. Lehtinen, and T. Aila, "Alias-free generative adversarial networks," *Advances in Neural Information Processing Systems*, vol. 34, pp. 852–863, 2021.

[164] T. Karras, S. Laine, and T. Aila, "A style-based generator architecture for generative adversarial networks," in *Proc. IEEE/CVF Conf. on Computer Vision and Pattern Recognition*, 2019, pp. 4401–4410.

[165] A. Kerckhoff, "La cryptographie militaire," *Journal of Military Science*, 1883.

[166] R. Kessler, O. Henninger, and C. Busch, "Fingerprints, forever young?" in *Intl. Conf. on Pattern Recognition (ICPR)*, 2021, pp. 8647–8654.

[167] K. Killourhy and R. A. Maxion, "Comparing anomaly-detection algorithms for keystroke dynamics," in *IEEE/IFIP International Conference on Dependable Systems & Networks*, 2009.

[168] S. Kim, K. Lewi, A. Mandal, H. Montgomery, A. Roy, and D. J. Wu, "Function-hiding inner product encryption is practical," in *International Conference on Security and Cryptography for Networks*. Springer, 2018, pp. 544–562.

[169] T. Kim, Y. Oh, and H. Kim, "Efficient privacy-preserving fingerprint-based authentication system using fully homomorphic encryption," *Security and Communication Networks*, vol. 2020, pp. 1–11, 2020.

[170] J. Kolberg, P. Bauspieß, M. Gomez-Barrero, C. Rathgeb, M. Dürmuth, and C. Busch, "Template protection based on homomorphic encryption: Computationally efficient application to iris-biometric verification and identification," in *IEEE Workshop on Information Forensics and Security (WIFS)*, 2019, pp. 1–6.

[171] J. Kolberg, P. Drozdowski, M. Gomez-Barrero, C. Rathgeb, and C. Busch, "Efficiency analysis of post-quantum-secure face template protection schemes based on homomorphic encryption," in *Intl. Conf. of the Biometrics Special Interest Group (BIOSIG)*. Gesellschaft für Informatik e.V., September 2020, pp. 175–182.

[172] J. Kornblum, "Identifying almost identical files using context triggered piecewise hashing," *Digital Investigation*, vol. 3, pp. 91–97, 2006.

[173] B. Krebs, "Facebook Stored Hundreds of Millions of User Passwords in Plain Text for Years," 2019. [Online]. Available: https://krebsonsecurity.com/2019/03/facebook-stored-hundreds-of-millions-of-user-passwords-in-plain-text-for-years/

[174] A. Langlois and D. Stehlé, "Worst-case to average-case reductions for module lattices," *Designs, Codes and Cryptography*, vol. 75, no. 3, pp. 565–599, 2015.

[175] J.-W. Lee, H. Kang, Y. Lee, W. Choi, J. Eom, M. Deryabin, E. Lee, J. Lee, D. Yoo, Y.-S. Kim *et al.*, "Privacy-preserving machine learning with fully homomorphic encryption for deep neural network," *IEEE Access*, vol. 10, pp. 30 039–30 054, 2022.

[176] K. Lewi, P. Mohassel, and A. Roy, "Single-message credential-hiding login," *Cryptology ePrint Archive*, 2020.

[177] Q. Li, Y. Sutcu, and N. Memon, "Secure sketch for biometric templates," in *International Conference on the Theory and Application of Cryptology and Information Security*. Springer, 2006, pp. 99–113.

[178] L. Lin, B. Tian, Y. Zhao, and Y. Niu, "A privacy-preserving gait recognition scheme under homomorphic encryption," in *Proc. Intl. Conf. on Networking and Network Applications (NaNA)*. IEEE, 2022, pp. 406–410.

[179] Y. Lindell, "Secure multiparty computation (MPC)," *Cryptology ePrint Archive*, 2020.

[180] E. Liu and Q. Zhao, "Encrypted domain matching of fingerprint minutia cylinder-code (MCC) with l1 minimization," *Neurocomputing*, vol. 259, pp. 3–13, 2017.

[181] Q. Lou, B. Feng, G. Charles Fox, and L. Jiang, "Glyph: Fast and accurately training deep neural networks on encrypted data," *Advances in Neural Information Processing Systems*, vol. 33, pp. 9193–9202, 2020.

[182] J. Loya and T. Bana, "Privacy-preserving keystroke analysis using fully homomorphic encryption & differential privacy," in *International Conference on Cyberworlds (CW)*. IEEE, 2021, pp. 291–294.

[183] V. Lyubashevsky, C. Peikert, and O. Regev, "On ideal lattices and learning with errors over rings," in *Ann. Intl. Conf. on the Theory and Applications of Cryptographic Techniques*. Springer, 2010, pp. 1–23.

[184] J. MacQueen, "Some methods for classification and analysis of multivariate observations," in *Proceedings of the Fifth Berkeley symposium on Mathematical Statistics and Probability*, vol. 1, no. 14, 1967, pp. 281–297.

[185] G. Mai, K. Cao, P. Yuen, and A. Jain, "On the reconstruction of face images from deep face templates," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2018.

[186] P. Manvi, A. Desai, K. Srinathan, and A. Namboodiri, "Sian: Secure iris authentication using noise," in *Proc. Intl. Joint Conf. on Biometrics (IJCB)*. IEEE, 2022, pp. 1–9.

[187] P. Manvi, A. M. Desai, K. Srinathan, and A. Namboodiri, "S-ban: Secure biometric authentication using noise," in *Proc. Indian Conference on Computer Vision, Graphics and Image Processing*, 2023, pp. 1–9.

[188] C. Marcolla, V. Sucasas, M. Manzano, R. Bassoli, F. H. Fitzek, and N. Aaraj, "Survey on fully homomorphic encryption, theory, and applications," *Proceedings of the IEEE*, vol. 110, no. 10, pp. 1572–1609, 2022.

[189] R. J. McEliece, "A public-key system based on algebraic coding theory," *DNS Report 44, Jet Propulsion Laboratory, California Institute of Technology*, 1978.

[190] B. Meden, P. Rot, P. Terhörst, N. Damer, A. Kuijper, W. J. Scheirer, A. Ross, P. Peer, and V. Štruc, "Privacy–enhancing face biometrics: A comprehensive survey," *IEEE Transactions on Information Forensics and Security*, vol. 16, pp. 4147–4183, 2021.

[191] Q. Meng, S. Zhao, Z. Huang, and F. Zhou, "Magface: A universal representation for face recognition and quality assessment," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2021, pp. 14 225–14 234.

[192] M. K. Morampudi, M. V. Prasad, and U. Raju, "Privacy-preserving iris authentication using fully homomorphic encryption," *Multimedia Tools and Applications*, vol. 79, pp. 19 215–19 237, 2020.

[193] M. K. Morampudi, M. V. Prasad, M. Verma, and U. Raju, "Secure and verifiable iris authentication system using fully homomorphic encryption," *Computers & Electrical Engineering*, vol. 89, p. 106924, 2021.

[194] C. V. Mouchet, J.-P. Bossuat, J. R. Troncoso-Pastoriza, and J.-P. Hubaux, "Lattigo: A multiparty homomorphic encryption library in go," in *Proc. Workshop on Encrypted Computing and Applied Homomorphic Cryptography*, no. CONF, 2020, pp. 64–70.

[195] T. Murakami, R. Fujita, T. Ohki, Y. Kaga, M. Fujio, and K. Takahashi, "Cancelable permutation-based indexing for secure and efficient biometric identification," *IEEE Access*, vol. 7, pp. 45 563–45 582, 2019.

[196] H. Nejatollahi, N. Dutt, S. Ray, F. Regazzoni, I. Banerjee, and R. Cammarota, "Post-quantum lattice-based cryptography implementations: A survey," *ACM Computing Surveys*, vol. 51, no. 6, pp. 1–41, 2019.

[197] T. Okamoto, "Authenticated key exchange and key encapsulation in the standard model," in *International Conference on the Theory and Application of Cryptology and Information Security*. Springer, 2007, pp. 474–484.

[198] M. Olsen, V. Šmida, and C. Busch, "Finger image quality assessment features - definitions and evaluation," *IET Biometrics*, vol. 5, no. 2, pp. 47–64, June 2016.

[199] J. Ortega-Garcia, J. Fierrez-Aguilar, D. Simon, J. Gonzalez, M. Faundez-Zanuy *et al.*, "MCYT baseline corpus: a bimodal biometric database," *Proc. Intl. Conf. Vision, Image and Signal Processing*, vol. 150, no. 6, pp. 395–401, December 2003.

[200] D. Osorio-Roig, C. Rathgeb, P. Drozdowski, and C. Busch, "Stable hash generation for efficient privacy-preserving face identification," *Trans. on Biometrics, Behavior, and Identity Science (TBIOM)*, vol. 4, no. 3, pp. 333–348, July 2021.

[201] H. Otroshi-Shahreza, C. Rathgeb, D. Osorio-Roig, V. Krivokuća, S. Marcel, and C. Busch, "Hybrid protection of biometric templates by combining homomorphic encryption and cancelable biometrics," in *Proc. of the 2022 Intl. Joint Conf. on Biometrics (IJCB)*. IEEE, October 2022.

[202] H. Otroshi-Shahreza, P. Melzi, D. Osorio-Roig, C. Rathgeb, C. Busch, S. Marcel, R. Tolosana, and R. Vera-Rodriguez, "Benchmarking of cancelable biometrics for deep templates," *arXiv:2302.13286*, 2023.

[203] P. Paillier, "Public-key cryptosystems based on composite degree residuosity classes," in *Intl. Conf. on the Theory and Applications of Cryptographic Techniques*. Springer, 1999, pp. 223–238.

[204] C. Peikert, "A decade of lattice cryptography," *Foundations and Trends® in Theoretical Computer Science*, vol. 10, no. 4, pp. 283–424, 2016.

[205] A. Pflug, C. Rathgeb, U. Scherhag, and C. Busch, "Binarization of spectral histogram models: An application to efficient biometric identification," in *2015 IEEE 2nd International Conference on Cybernetics (CYBCONF)*. IEEE, 2015, pp. 501–506.

[206] J. Phillips, P. Flynn, T. Scruggs, K. Bowyer, J. Chang *et al.*, "Overview of the Face Recognition Grand Challenge," in *Conf. on Computer Vision and Pattern Recognition (CVPR)*, vol. 1. IEEE, June 2005, pp. 947–954.

[207] J. Phillips, H. Moon, S. Rizvi, and P. Rauss, "The FERET evaluation methodology for face-recognition algorithms," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 22, no. 10, pp. 1090–1104, October 2000.

[208] K. Pietrzak, "Cryptography from learning parity with noise," in *International Conference on Current Trends in Theory and Practice of Computer Science*. Springer, 2012, pp. 99–114.

[209] N. Popescu-Bodorin and V. E. Balas, "AI challenges in iris recognition: Processing tools for Bath iris image database," in *Proc. Intl. Conf. on Automation and Information*, 2010, pp. 116–121.

[210] G. Pradel and C. Mitchell, "Privacy-preserving biometric matching using homomorphic encryption," in *Intl. Conf. on Trust, Security and Privacy in Computing and Communications (TrustCom)*. IEEE, 2021, pp. 494–505.

[211] H. Proença, "Unconstrained iris recognition in visible wavelengths," in *Handbook of Iris Recognition*. Springer, 2016, pp. 321–358.

[212] M. Qi, J. Chen, and Y. Chen, "A secure biometrics-based authentication key exchange protocol for multi-server TMIS using ECC," *Computer Methods and Programs in Biomedicine*, vol. 164, pp. 101–109, 2018.

[213] M. M. Rahman, T. I. Mishu, and M. A. A. Bhuiyan, "Performance analysis of a parameterized minutiae-based approach for securing fingerprint templates in biometric authentication systems," *Journal of Information Security and Applications*, vol. 67, p. 103209, 2022.

[214] N. K. Ratha, S. Chikkerur, J. H. Connell, and R. M. Bolle, "Generating cancelable fingerprint templates," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 29, no. 4, pp. 561–572, 2007.

[215] N. K. Ratha, J. H. Connell, and R. M. Bolle, "An analysis of minutiae matching strength," in *International Conference on Audio-and Video-Based Biometric Person Authentication*. Springer, 2001, pp. 223–228.

[216] C. Rathgeb and A. Uhl, "A survey on biometric cryptosystems and cancelable biometrics," *EURASIP Journal on Information Security*, vol. 3, March 2011.

[217] C. Rathgeb, J. Kolberg, A. Uhl, and C. Busch, "Deep learning in the field of biometric template protection: An overview," *arXiv preprint arXiv:2303.02715*, 2023.

[218] C. Rathgeb, J. Merkle, J. Scholz, B. Tams, and V. Nesterowicz, "Deep face fuzzy vault: Implementation and performance," *Computers & Security*, vol. 113, p. 102539, 2022.

[219] C. Rathgeb, B. Tams, J. Merkle, V. Nesterowicz, U. Korte, and M. Neu, "Multi-biometric fuzzy vault based on face and fingerprints," in *Proc. Intl. Joint Conference on Biometrics (IJCB)*. IEEE, 2023.

[220] C. Rathgeb, B. Tams, J. Wagner, and C. Busch, "Unlinkable improved multi-biometric iris fuzzy vault," *EURASIP Journal on Information Security*, vol. 2016, pp. 1–16, 2016.

[221] I. S. Reed and G. Solomon, "Polynomial codes over certain finite fields," *Journal of the Society for Industrial and Applied Mathematics*, vol. 8, no. 2, pp. 300–304, 1960.

[222] O. Regev, "On lattices, learning with errors, random linear codes, and cryptography," *Journal of the ACM (JACM)*, vol. 56, no. 6, pp. 1–40, 2009.

[223] R. L. Rivest, L. Adleman, M. L. Dertouzos *et al.*, "On data banks and privacy homomorphisms," *Foundations of Secure Computation*, vol. 4, no. 11, pp. 169–180, 1978.

[224] R. L. Rivest, A. Shamir, and L. Adleman, "A Method for Obtaining Digital Signatures and Public-Key Cryptosystems," *Communications of the ACM*, vol. 21, no. 2, pp. 120–126, 1978.

[225] T. Rohwedder, D. Osorio-Roig, C. Rathgeb, and C. Busch, "Benchmarking fixed-length fingerprint representations across different embedding sizes and sensor types," in *Proc. Intl. Conf. of the Biometrics Special Interest Group (BIOSIG)*. Gesellschaft für Informatik e.V., September 2023.

[226] R. Román, R. Arjona, P. López-González, and I. Baturone, "A quantum-resistant face template protection scheme using kyber and saber public key encryption algorithms," in *Proc. Intl. Conf. of the Biometrics Special Interest Group (BIOSIG)*. IEEE, 2022, pp. 1–5.

[227] K. Ruhloff, D. Cousins, and Y. Polyakov, *The PALISADE Lattice Cryptography Library*, 2017. [Online]. Available: https://git.njit.edu/palisade/PALISADE

[228] A. R. Sadeghi, T. Schneider, and I. Wehrenberg, "Efficient privacy-preserving face recognition," in *Int. Conf. on Information Security and Cryptology*. Springer, 2009, pp. 229–244.

[229] A. Sardar, S. Umer, C. Pero, and M. Nappi, "A novel cancelable face-hashing technique based on non-invertible transformation with encryption and decryption template," *IEEE Access*, vol. 8, pp. 105 263–105 277, 2020.

[230] A. Sarkar and B. K. Singh, "A novel session key generation and secure communication establishment protocol using fingerprint biometrics," in *Handbook of Computer Networks and Cyber Security*. Springer, 2020, pp. 777–805.

[231] W. J. Scheirer and T. E. Boult, "Cracking fuzzy vaults and biometric encryption," in *2007 Biometrics Symposium*. IEEE, 2007, pp. 1–6.

[232] C.-P. Schnorr, "Efficient identification and signatures for smart cards," in *Advances in Cryptology*. Springer, 1989, pp. 239–252.

[233] B. Schoenmakers and P. Tuyls, "Efficient binary conversion for paillier encrypted values," in *Advances in Cryptology-EUROCRYPT 2006: 24th Annual International Conference on the Theory and Applications of Cryptographic Techniques, St. Petersburg, Russia, May 28-June 1, 2006. Proceedings 25*. Springer, 2006, pp. 522–537.

[234] "Microsoft SEAL (release 4.1)," https://github.com/Microsoft/SEAL, Jan. 2023, microsoft Research, Redmond, WA.

[235] J. Šeděnka, K. S. Balagani, V. Phoha, and P. Gasti, "Privacy-preserving population-enhanced biometric key generation from free-text keystroke dynamics," in *IEEE International Joint Conference on Biometrics*. IEEE, 2014, pp. 1–8.

[236] C. E. Shannon, "A mathematical theory of communication," *The Bell System Technical Journal*, vol. 27, no. 3, pp. 379–423, 1948.

[237] C. Shen, T. Yu, H. Xu, G. Yang, and X. Guan, "User practice in password security: An empirical study of real-life passwords in the wild," *Computers & Security*, vol. 61, pp. 130–141, 2016.

[238] P. W. Shor, "Algorithms for quantum computation: discrete logarithms and factoring," in *Proc. Ann. Symposium on Foundations of Computer Science*. IEEE, 1994, pp. 124–134.

[239] V. Shoup, "NTL: A library for doing number theory," 2001.

[240] T. Silde and M. Strand, "Anonymous tokens with public metadata and applications to private contact tracing," in *Financial Cryptography and Data Security*, I. Eyal and J. Garay, Eds. Cham: Springer International Publishing, 2022, pp. 179–199.

[241] K. Simoens, J. Bringer, H. Chabanne, and S. Seys, "A framework for analyzing template security and privacy in biometric authentication systems," *IEEE Transactions on Information forensics and security*, vol. 7, no. 2, pp. 833–841, 2012.

[242] K. Singh, R. Sirdey, and S. Carpov, "Practical personalized genomics in the encrypted domain," in *2018 Third International Conference on Fog and Mobile Edge Computing (FMEC)*. IEEE, 2018, pp. 139–146.

[243] L. Sperling, N. Ratha, A. Ross, and V. N. Boddeti, "HEFT: Homomorphically encrypted fusion of biometric templates," in *Proc. Intl. Joint Conference on Biometrics (IJCB)*. IEEE, 2022.

[244] D. Stebila and M. Mosca, "Post-quantum key exchange for the internet and the open quantum safe project," in *International Conference on Selected Areas in Cryptography*. Springer, 2016, pp. 14–37.

[245] K. Sundararajan and D. L. Woodard, "Deep learning for biometrics: A survey," *ACM Computing Surveys*, vol. 51, no. 3, pp. 1–34, 2018.

[246] J. Surbiryala, R. Raghavendra, and C. Busch, "Finger vein indexing based on binary features," in *2015 Colour and Visual Computing Symposium (CVCS)*. IEEE, 2015, pp. 1–6.

[247] Y. Sutcu, H. T. Sencar, and N. Memon, "A secure biometric authentication scheme based on robust hashing," in *Proc. Workshop on Multimedia and Security*, 2005, pp. 111–116.

[248] E. Tabassi, M. Olsen, O. Bausinger, C. Busch, A. Figlarz, G. Fiumara, O. Henniger, J. Merkle, T. Ruhland, C. Schiel, and M. Schwaiger, "NIST interagency report 8382," National Institute of Standards and Technology, NIST Interagency Report 8382, July 2021.

[249] H. Tamiya, T. Isshiki, K. Mori, S. Obana, and T. Ohki, "Improved post-quantum-secure face template protection system based on packed homomorphic encryption," in *Proc. Intl. Conf. of the Biometrics Special Interest Group (BIOSIG)*. IEEE, 2021, pp. 1–5.

[250] B. Tams, "Decodability attack against the fuzzy commitment scheme with public feature transforms," *arXiv preprint arXiv:1406.1154*, 2014.

[251] B. Tams, "Unlinkable minutiae-based fuzzy vault for multiple fingerprints," *IET Biometrics*, vol. 5, no. 3, pp. 170–180, 2016.

[252] The Biometric Systems Lab (University of Bologna), the Pattern Recognition and Image Processing Laboratory (Michigan State University) and the Biometric Test Center (San Jose State University), "Fingerprint verification competition 2004," March 2004.

[253] W. A. Torres, N. Bhattacharjee, and B. Srinivasan, "Privacy-preserving biometrics authentication systems using fully homomorphic encryption," *International Journal of Pervasive Computing and Communications*, vol. 11, no. 2, pp. 151–168, 2015.

[254] J. R. Troncoso-Pastoriza, D. González-Jiménez, and F. Pérez-González, "Fully private noninteractive face verification," *IEEE Transactions on Information Forensics and Security*, vol. 8, no. 7, pp. 1101–1114, 2013.

[255] N. Tyagi, S. Celi, T. Ristenpart, N. Sullivan, S. Tessaro, and C. A. Wood, "A fast and simple partially oblivious prf, with applications," in *Annual*

International Conference on the Theory and Applications of Cryptographic Techniques. Springer, 2022, pp. 674–705.

[256] Unique Identification Authority of India, "Aadhaar Dashboard," https://www.uidai.gov.in/aadhaar_dashboard/, 2024.

[257] E. Uzun, S. P. Chung, V. Kolesnikov, A. Boldyreva, and W. Lee, "Fuzzy labeled private set intersection with applications to private real-time biometric search," in *30th USENIX Security Symposium (USENIX Security 21)*, 2021, pp. 911–928.

[258] D. K. Vallabhadas and M. Sandhya, "Securing multimodal biometric template using local random projection and homomorphic encryption," *Journal of Information Security and Applications*, vol. 70, 2022.

[259] R. Važan, "Sourceafis fingerprint matcher v3.13.0," 2021, https://sourceafis.machinezoo.com/, accessed 2022-01-18.

[260] M. Wang, K. He, J. Chen, Z. Li, W. Zhao, and R. Du, "Biometrics-authenticated key exchange for secure messaging," in *Proceedings of the ACM SIGSAC Conference on Computer and Communications Security*, 2021, pp. 2618–2631.

[261] Y. Wang, J. Wan, J. Guo, Y.-M. Cheung, and P. C. Yuen, "Inference-based similarity search in randomized montgomery domains for privacy-preserving biometric identification," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 40, no. 7, pp. 1611–1624, 2017.

[262] V. Weger, N. Gassner, and J. Rosenthal, "A survey on code-based cryptography," *arXiv preprint arXiv:2201.07119*, 2024.

[263] T. Wingarz, M. Gomez-Barrero, C. Busch, and M. Fischer, "Privacy-preserving convolutional neural networks using homomorphic encryption," in *2022 International Workshop on Biometrics and Forensics (IWBF)*. IEEE, 2022, pp. 1–6.

[264] T. Wu, "The secure remote password protocol," in *Proc. 1998 Internet Society Symposium on Network and Distributed Systems Security*. Citeseer, 1998, pp. 97–111.

[265] S. Xu, Y. Cao, X. Chen, S.-M. Yiu, and Y. Zhao, "Post-quantum public-key authenticated searchable encryption with forward security: General construction, implementation, and applications," *Cryptology ePrint Archive*, 2023.

[266] W. Yang, S. Wang, K. Yu, J. J. Kang, and M. N. Johnstone, "Secure fingerprint authentication with homomorphic encryption," in *Digital Image Computing: Techniques and Applications*. IEEE, 2020, pp. 1–6.

[267] W. Yang, S. Wang, H. Cui, Z. Tang, and Y. Li, "A review of homomorphic encryption for privacy-preserving biometrics," *Sensors*, vol. 23, no. 7, p. 3566, 2023.

[268] Y. Yang, Q. Zhang, W. Gao, C. Fan, Q. Shu, and H. Yun, "Design on face recognition system with privacy preservation based on homomorphic encryption," *Wireless Personal Communications*, vol. 123, no. 4, pp. 3737–3754, 2022.

[269] A. C. Yao, "How to generate and exchange secrets," in *Symposium on Foundations of Computer Science (SFCS)*. IEEE, 1986, pp. 162–167.

[270] M. Yasuda, T. Shimoyama, J. Kogure, K. Yokoyama, and T. Koshiba, "Packed homomorphic encryption based on ideal lattices and its application to biometrics," in *Intl. Conf. on Availability, Reliability, and Security*. Springer, 2013, pp. 55–74.

[271] M. Yasuda, T. Shimoyama, J. Kogure, K. Yokoyama, and T. Koshiba, "New packing method in somewhat homomorphic encryption and its applications," *Security and Communication Networks*, vol. 8, no. 13, pp. 2194–2213, 2015.

[272] M. Yasuda, "Secure hamming distance computation for biometrics using ideal-lattice and ring-LWE homomorphic encryption," *Information Security Journal: A Global Perspective*, vol. 26, no. 2, pp. 85–103, 2017.

[273] E. Yu and S. Cho, "GA-SVM wrapper approach for feature subset selection in keystroke dynamics identity verification," in *Proceedings of the International Joint Conference on Neural Networks (IJCNN)*, vol. 3. IEEE, 2003, pp. 2253–2257.

[274] E. Zhang, J. Chang, and Y. Li, "Efficient threshold private set intersection," *IEEE Access*, vol. 9, pp. 6560–6570, 2021.

[275] S. Zhang, Z. Yan, W. Liang, K.-C. Li, and C. Dobre, "BAKA: Biometric authentication and key agreement scheme based on fuzzy extractor for wireless body area networks," *IEEE Internet of Things Journal*, 2023.

[276] X. Zhang, C. Huang, D. Gu, J. Zhang, and H. Wang, "BIB-MKS: Post-quantum secure biometric identity-based multi-keyword search over encrypted data in cloud storage systems," *IEEE Transactions on Services Computing*, 2021.

[277] Y. Zhang and F. Koushanfar, "Robust privacy-preserving fingerprint authentication," in *IEEE Intl. Symposium on Hardware Oriented Security and Trust*. IEEE, 2016, pp. 1–6.

[278] Y. Zhang, J. Qin, and L. Du, "A secure biometric authentication based on PEKS," *Concurrency and Computation: Practice and Experience*, vol. 28, no. 4, pp. 1111–1123, 2016.

Paper A

On the Feasibility of Fully Homomorphic Encryption of Minutiae-Based Fingerprint Representations

Pia Bauspieß, Lasse Vad, Håvard Myrekrok, Anamaria Costache, Jascha Kolberg, Christian Rathgeb, and Christoph Busch

Published at International Conference on Information Systems Security and Privacy (ICISSP), 2023

Abstract

Protecting minutiae-based fingerprint templates with fully homomorphic encryption has recently been recognised as a hard problem. In this work, we evaluate state-of-the-art fingerprint recognition based on minutiae templates using post-quantum secure fully homomorphic encryption that operates directly on floating point numbers, such that no simplification or quantisation of the comparison algorithm is necessary. In a practical evaluation on a publicly available dataset, we run a benchmark and provide directions for future work.

A.1 Introduction

Fingerprint patterns allow for an irrevocable and accurate identification of individuals over several decades [166]. Images and templates representing such patterns have therefore, along with other biometric data, been recognised as sensitive personal data by the European Union's General Data Protection Regulation and the ISO/IEC 24745 [149] standard.

In its most recent version from 2022, the standard places particular emphasis on Biometric Information Protection (BIP) in the presence of quantum computers. In their Quantum Manifesto [107], the European Union expects quantum computers to pose a realistic threat within the next 15 years. Comparing this time frame to the the retention period for biometric systems ranging from 5 [110] up to 12 years [166], it becomes evident that long-term protection of biometric data needs to be addressed today.

More concretely, access to a quantum computer would allow an attacker to break the unlinkability, irreversibility, and renewability assurances of classically protected BIP systems, leaving the enrolment data vulnerable for malicious exploitation. These three requirements are defined in ISO/IEC 24745 [149] as *i) unlinkability*, two protected templates stored in different applications cannot be linked to the same subject, *ii) renewability*, new templates can be created from the same biometric instance without the need to re-enrol, and *iii) irreversibility*, it is impossible to retrieve original templates given only protected templates. Considering the quantum challenge, this work proposes a BIP system that achieves long-term protection according to the standard's requirements through the use of post-quantum cryptography.

However, the lift to post-quantum security does not come without challenges. In particular, the combination of accurate minutiae-based fingerprint recognition and BIP through Fully Homomorphic Encryption (FHE) has recently been recognised as a notorious hard problem by leading researchers in biometrics [104].

So far, solutions have only been proposed for fixed-length fingerprint representations [169], or using classically secure cryptography [125]. The novelty and objective of this work is therefore to evaluate minutiae-based fingerprint comparison with FHE on floating point numbers, an encryption scheme which enjoys increasing interest since its proposal in 2017 [65]. As a lattice-based FHE scheme, its post-quantum security is provided by the Ring-Learning with Errors (R-LWE) [183] hardness assumption.

This work presents post-quantum secure minutiae-based fingerprint comparison algorithm [58] using FHE, where the comparison algorithm has not been simplified or quantised in order to be compatible with the encryption scheme. Furthermore, we highlight challenges inherent to the application of FHE to minutiae-based fingerprint comparison and provide an experimental benchmark from which we draw conclusions for future work.

The rest of this paper is structured as follows: Section A.2 contextualises our contribution, before we present our proposed system in Section A.3. We give an experimental evaluation in Section A.4 and draw our conclusions in Section A.5.

A.2 Related Work

Fingerprint recognition has historically been based on minutiae, which are defined as ridge endings and bifurcations of fingerprint ridges. While comparison algorithms with high accuracy have been developed [58, 259], they reflect the complexity inherent to comparing two sets of minutiae such as rotation, non-linear transformation, and absence of an inherent ordering. In their development, they have not necessarily considered the application of encryption schemes, which offer only a limited number of operations that can be computed with feasible computational effort [147]. Therefore, two research directions have emerged that approach the challenge of combining fingerprint recognition with encryption: one is to develop fingerprint representations with simple distance functions as comparison metrics that maintain high recognition accuracy, while the other is to apply and adapt compatible encryption schemes to complex minutiae-based comparators.

Indeed, FHE for fixed-length representations has been proposed for different biometric modalities such as face [31, 44, 171] and iris [170] with high accuracy and real-time efficiency. For fingerprint specifically, the most prominent representation is Jain et al.'s *FingerCode* [153]. Notable works on encrypting this representation include [26, 126, 266]. However, the encryption schemes used are based on classical assumptions and do not hold in the quantum age. A recent

work using FHE with post-quantum security on FingerCode templates is [169]. The FHE scheme [69] applied here only tolerates binary values, which is compatible with FingerCodes, but not with minutiae templates.

Minutiae-based comparators share the difficulty of finding close pairs within the sets of k reference minutiae and l probe minutiae, the mapping between which can be neither injective nor surjective due to potential missing or spurious minutiae. In addition, samples might be rotated, translated or distorted, requiring either prealignment or a rotation-invariant approach. In theory, FHE allows for the evaluation of arbitrary circuits on encrypted input data [121]. In practice however, both alignment and set comparison are functions that can only be described using conditional statements, the number of which in prevalent approaches is high [259, 277]. Their combination with FHE is therefore not straightforward, and more importantly too costly for practical applications [147]. In contrast to that, the comparison of alignment-free fixed-length representations can be performed by computing a simple distance function on the encrypted templates, the result of which is typically decrypted to evaluate the comparison against the decision threshold.

Classically secure homomorphic encryption, which is only partially or somewhat homomorphic [121], has been applied to minutiae-based comparison [125]. However, these schemes lack post-quantum security. This is also true for a approaches based on cancelable biometric templates constructed based on randomized feature transformation, most recently represented by [213], which do not adhere to formal security proofs and are vulnerable to unlinkability attacks. In particular, the indistinguishability under chosen plaintext attacks provided by (F)HE schemes, which gives formal security in terms of ISO/IEC 24745 [149] is not given in the latter. Other works [180, 277, 124] have utilised secure multiparty computation (MPC), which is generally speaking more flexible than FHE. As a drawback, it introduces a communication overhead, and practical post-quantum secure MPC has only been explored recently [55].

Table A.1 gives a qualitative overview of the most relevant related works discussed in this Section and provides a comparison against our proposed approach.

A.3 Proposed System

We study a combination of the minutiae-based fingerprint comparison algorithm Minutia Cylinder-Code (MCC) [58] and the state-of-the-art FHE encryption scheme Cheon-Kim-Kim-Song (CKKS) [65] to illustrate the challenges that arise in the process.

Reference	Template protection category	Crypto-graphic scheme	Variable-length feature representation	Post-quantum security
Barni et al., 2010 [26]	HE	ElGamal [102] Pailler [203]	✗	✗
Gomez-Barrero et al., 2017 [126] Yang et al., 2020 [266]	HE	Pailler [203]	✗	✗
Zhang and Koushanfar, 2016 [277] Gilkalaye and Derakhshani, 2021 [124]	MPC	Garbled Circuits [269]	✓	✗
Gomez-Barrero et al., 2017 [125]	HE	Pailler [203]	✓	✗
Kim et al., 2020 [169]	FHE	TFHE [69]	✗	✓
Ours	FHE	CKKS [65]	✓	✓

Table A.1: Qualitative comparison of related work on cryptographic fingerprint template protection.

A.3.1 Background

Before we describe our proposed system, we introduce the necessary background in this Section. Subsequently, we introduce the baseline verification scheme without encryption, and finally, our proposed protected system.

Throughout this work, we consider a biometric system operating in verification mode. In a setup phase, subjects are enrolled to the system with their fingerprint features. During a verification transaction, a fresh probe sample is captured a biometric claim, i.e., the claimed identity of the data subject, is transferred to the database along with the probe feature set. Then, a comparison between the probe features and the reference template corresponding to the claim is computed, resulting in a comparison score in the range $[0, 1]$, where 1 indicates highest similarity. Finally, this score is compared against a predetermined decision threshold and the comparison trial is accepted or rejected accordingly. In the following section, we describe this comparison algorithm in more detail.

A.3.1.1 Minutia Cylinder Code

Minutia Cylinder-Code (MCC) [58] is a fingerprint comparison algorithm that takes as input two minutiae-based fingerprint templates as standardized in ISO/IEC 19794-2 [151] and outputs a similarity score that can further be used for an automated comparison. Minutiae are significant points in the pattern of fingerprint ridges: ridge endings and bifurcations, where one ridge line splits into two. We remind the reader of the following definition of an ISO/IEC 19794-2 [151] fingerprint template in the notation of [58], Section 3.

Definition A.1 (Fingerprint Template)**.** A fingerprint template is an unordered set $T = \{m_i\}_{i=1}^{N}$ of minutiae m_i, where N is the number of minutiae found in a given fingerprint image. Each minutia is given as a tuple $m = (x_m, y_m, \theta_m)$ of its location in terms of x- and y-coordinate (x_m, y_m) given in pixels from the left upper corner of the sample together with its tangential angle with respect to the x-axis θ_m.

Note that the number of minutiae N varies between captures, not only between different subjects, but also within repeated captures of the same instance. This is due to noise during the capture process: depending on the image quality and capture conditions, minutiae can either be missed during feature extraction, or spurious minutiae can be added, resulting in different length representations of the same fingerprint. In addition, the location of the minutiae are subject to fuzziness, as their location and angle can be distorted through rotation, translation and non-linear transformations. Therefore, minutiae-based fingerprint comparison comprises of the complex problem of accurately comparing two unordered, variable-sized sets of noisy points, a number of which can be spurious.

To address the aforementioned challenges, MCC introduces a local structure associated with each minutiae referred to as a *minutia cylinder*. This structure incorporates information about the neighbourhood of each minutiae, i.e., further minutiae found in close proximity and their spatial and directional relationship with the center minutiae [58]. This approach ensures system interoperability as the cylinder representation is still based on ISO/IEC 19794-2 [151] fingerprint templates. In particular, the variable-length representation is maintained, as the number of minutia cylinders corresponds to the number of minutiae in a fingerprint template. We restate the following definitions from [58], Section 3.

Definition A.2 (Minutia Cylinder)**.** A minutia cylinder is given by a fixed radius R and height 2π centered around the location (x_m, y_m) of a minutia m. It is discretized into small cuboids, called *cells*, which are orientated in the direction of the tangential angle θ_m of the center minutiae. It can be represented as a vector $\mathbf{c}_m \in [0, 1]^n$, where n denotes the total number of cells in a cylinder.

As a minutia cylinder only contains relative information concerning the relationship between the minutiae, such as distance and directional difference, but no global information, it can be considered translation and rotation invariant [58]. The same properties also make it robust against minor non-linear transforms during capture such as different levels of pressure applied on the fingerprint sensor. Most importantly, the fixed-radius neighbourhood is a key component in the handling of missing and spurious minutiae [58].

Definition A.3 (Contribution Score)**.** Each cell inside a minutia cylinder is as-

signed a numerical value C_m, called *contribution score*, which details the likelihood of finding another minutia in a small neighbourhood with a compatible directional difference.

For a more technical definition along with insightful figures, the reader is referred to [58], Section 3.

Definition A.4 (Cylinder Set)**.** Given a fingerprint template T, its corresponding cylinder set is defined as the set of valid cylinders $\mathbf{c}_m$ for $m \in T$. A cylinder is considered valid if it contains a sufficient number of contribution scores, i.e., exceeding a predefined threshold of a minimal number of contribution scores and a minimal number of contributing neighbour minutiae.

Finally, a reference fingerprint template can be compared against a probe feature set based on their cylinder set representations. Therefore, we restate the comparison process given in [58].

Definition A.5 (Conditional Contribution)**.** Let $\mathbf{c}_a$ and $\mathbf{c}_b$ be cylinders corresponding to minutia a in a reference template and minutia b in a probe template. Then, $\mathbf{c}_{a|b} = \mathbf{c}_a$ where $\mathbf{c}_b \neq 0$. In other words, $\mathbf{c}_{a|b}$ contains all contributions from $\mathbf{c}_a$ where $\mathbf{c}_b$ has contribution from corresponding cells.

Definition A.6 (Candidate Pair)**.** Two cylinders represented by $\mathbf{c}_a$ and $\mathbf{c}_b$ are considered a *candidate pair* if and only if they satisfy the following requirements:

1. The directional difference between the two minutiae a and b is not greater than $\frac{\pi}{2}$.
2. At least 60% of corresponding elements in the two vectors $\mathbf{c}_a$ and $\mathbf{c}_b$ are non-zero.
3. $\|\mathbf{c}_{a|b}\| + \|\mathbf{c}_{b|a}\| \neq 0$.

Intuitively, it can be seen that these conditions enable to filter out the most relevant pairings of cylinders. Firstly, the orientation of the minutiae should be reasonably close in order to be considered as a mated comparison trial. Secondly, there is a significant overlap in the contribution scores associated with each minutia cylinder, and thirdly, the contributions in said overlap should occur at similar indices, indicating that the spacial relationships to neighbour minutiae are similar. Based on valid pairings of cylinders according to these criteria, the overall similarity between two cylinders is given by the following definition.

Definition A.7 (Cylinder similarity)**.** The cylinder similarity between two cylinders represented by their vectors $\mathbf{c}_a$ and $\mathbf{c}_b$ is given as

$$\gamma(a,b) = \begin{cases} 1 - \frac{\|\mathbf{c}_{a|b} - \mathbf{c}_{b|a}\|}{\|\mathbf{c}_{a|b}\| + \|\mathbf{c}_{b|a}\|}, & \text{if } \mathbf{c}_a \text{ and } \mathbf{c}_b \text{ are } \textit{candidate pairs}. \\ 0, & \text{otherwise.} \end{cases} \tag{A.1}$$

The cylinder similarity allows to calculate local similarity scores for each minutia pair. From those local scores, a global similarity score indicating the similarity between two fingerprints can be calculated. The authors of [58] propose four different strategies for global score consolidation. In our work, *Local Similarity Sort* (LSS) is applied, where the top k cylinder similarity scores are averaged to produce the global similarity score.

A.3.1.2 Fully Homomorphic Encryption

Fully Homomorphic Encryption (FHE) schemes allow for additions and multiplications of ciphertexts that correspond directly to operations on the corresponding plaintexts [224]. More formally, a cryptographic scheme is *homomorphic* if

$$Enc_{pk}(a * b) = Enc_{pk}(a) * Enc_{pk}(b) \tag{A.2}$$

for an operation $*$. In partially homomorphic encryption schemes, this property is limited to either addition or multiplication. In comparison, FHE schemes allow for a combination of additions and multiplications, making them applicable to a wider variety of use cases.

The public-key encryption scheme used in this work is the Cheon-Kim-Kim-Song (CKKS) [65]. Historically, FHE schemes have first been proposed for integer or binary input data. Only more recently, [65] have proposed a scheme that operated on floating point numbers directly, eliminating the need for input quantisation or significant rounding. While the scheme does come with an approximation error, its order of magnitude is not significant to the application in our work.

Similarly to other FHE schemes, the security of CKKS based on the hardness of the Ring-Learning with Errors (R-LWE) problem [183]. Encryption within such schemes is a probabilistic operation, meaning that every encryption uses fresh randomness. In addition, encryption in CKKS is indistinguishable under chosen-plaintext attacks (IND-CPA), such that an attacker cannot distinguish between an encryption of 0 and an encryption of 1. In particular, an attacker cannot distinguish between two encryptions of the same input, e.g., the biometric template of a specific data subject, and an encryption of a different input,

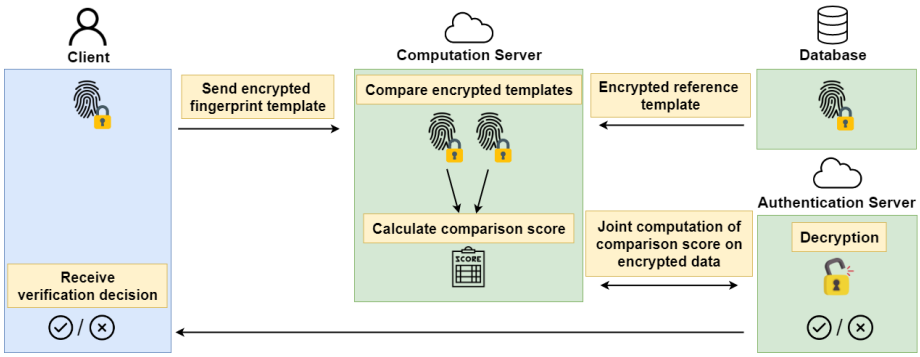

Figure A.1: Simplified flowchart of the proposed solution.

e.g. the biometric template of a different subject. For more details, we refer the reader to the original scheme [65].

A.3.2 Baseline System

The baseline system operates in verification mode on unprotected data without encryption. During enrolment, the reference subjects' fingerprint samples are captured and features are extracted as ISO/IEC 19794-2 [151] fingerprint templates. From the templates, the MCC cylinder sets are constructed as described above, and stored in the reference database. For a verification transaction, a probe subject's features are extracted in the same manner and represented as a cylinder set. Then, the probe cylinder set is compared against the reference cylinder set corresponding to the claimed identity of the probe subject. The comparison outcome is the global similarity score of the two cylinder sets, which is compared against the predefined decision threshold to yield the comparison trial outcome.

A.3.3 Protected System

The protected system builds on the baseline system, but with the addition of FHE. The reference templates are stored in ciphertext form, and the probe features are encrypted before comparison. Through the homomorphic properties of the FHE scheme, the comparison algorithm can be computed on the encrypted data, ensuring privacy protection of the underlying data.

MCC operation	**Enc**	**EvalAdd**	**EvalSub**	**EvalMult**	**EvalAtIndex(· , 1)**	**Dec**
Cylinder encryption	3	—	—	—	—	—
Directional difference	—	—	1	—	—	1
Common validity	—	$n-1$	1	2	$n-1$	1
Denominator	—	$2(n-1)$	2	2	$2(n-1)$	2
Numerator	—	$n-1$	2	1	$n-1$	1
Total	3	$4(n-1)$	6	5	$4(n-1)$	5

Table A.2: Homomorphic operations for the encrypted comparison of two minutia cylinders.

Operation on encrypted data	**Add**	**Subtract**	**Rotate**	**Decrypt**	**Multiply**	**Encrypt**
Relative cost	1	5	24	33	46	52

Table A.3: Relative cost of CKKS [65] operations implemented in PALISADE [227].

We work in an established client-server architecture with a computation server (CS) controlling the database of encrypted reference templates and an authentication server (AS) controlling the secret key for decryption in a semi-honest adversary model [270]. Figure A.1 shows the workflow of the protected system.

In the first step, the client captures a fingerprint sample and generates a cylinder set from it. For each minutia point m, it constructs the encrypted cylinder as a tuple of three CKKS ciphertexts $[Enc_{pk}(\theta_m), Enc_{pk}(\mathbf{c}_m), Enc_{pk}(\mathbf{c}_m^{\text{val}})]$ using coefficient packing. The first ciphertext is the encrypted cylinder angle θ_m, which inherits the minutia angle. The second ciphertext is an encryption of the contribution vector $\mathbf{c}_m$, while the third ciphertext stores the vector $\mathbf{c}_m^{\text{val}}$, which represents the validity of each cell related to minutia m. Even though cylinders are encrypted individually, they cannot be utilised for hill-climbing attacks due to the chosen plaintext security of the encryption scheme. In other words, the separate encryption of multiple cylinders does not lower the privacy protection compared to an encryption of the entire set of cylinders.

For CS to execute the comparison between all probe and reference cylinders, it first determines pairs of cylinders that can be considered *candidate pairs*. Following Definition A.6, the first condition requires the directional difference between two cylinders to be lower than $\frac{\pi}{2}$. This is evaluated in the encrypted domain by subtracting the two encrypted minutia angles $Enc_{pk}(\theta_a) - Enc_{pk}(\theta_b) = Enc_{pk}(\theta_a - \theta_b)$. The resulting difference is decrypted at AS and compared against $\frac{\pi}{2}$ by CS. The comparison is computed in plaintext, as evaluating encrypted con-

ditional statements is complex [147]. However, the difference between two angles does not reveal the orientation of the original minutiae, and therefore, does not leak critical information.

For the second condition, CS verifies that over 60% of the corresponding elements in $\mathbf{c}_a$ and $\mathbf{c}_b$ are non-zero by calculating a common validity vector as the a homomorphic multiplication of two encrypted validity vectors $Enc_{pk}(\mathbf{c}_a^{\text{val}})$ and $Enc_{pk}(\mathbf{c}_b^{\text{val}})$. The number of elements in the resulting packed vector can be obtained by applying the rotation technique first introduced in [44]. The resulting value is decrypted in order to evaluate the condition. If the amount of non-zero elements in the two vectors is below 60% of the total amount of elements, the cylinders are not *candidate pairs* and are not considered further.

The third step is calculating the vectors $Enc_{pk}(\mathbf{c}_{a|b})$ and $Enc_{pk}(\mathbf{c}_{b|a})$ and their norms. For this step, CS multiplies $Enc_{pk}(\mathbf{c}_a)$ and $Enc_{pk}(\mathbf{c}_b)$ with the common validity vector homomorphically, which filters out contributions of cells that should not be taken into account for the cylinder similarity score. The Euclidean norm of the resulting vectors $Enc_{pk}(\mathbf{c}_{a|b})$ and $Enc_{pk}(\mathbf{c}_{b|a})$ can then again be evaluated as above. Then, AS decrypts $Enc_{pk}(||\mathbf{c}_{a|b}||)$ and $Enc_{pk}(||\mathbf{c}_{b|a}||)$ and CS checks that $\|\mathbf{c}_{a|b}\| + \|\mathbf{c}_{b|a}\| \neq 0$.

For the cylinder pairings that can be considered *candidate pairs*, the final cylinder similarity score is given in Definition A.7. The denominator has already been calculated in the previous step, while the numerator is calculated by performing one homomorphic subtraction of $Enc_{pk}(\mathbf{c}_{a|b}) - Enc_{pk}(\mathbf{c}_{b|a}) = Enc_{pk}(\mathbf{c}_{a|b} - \mathbf{c}_{b|a})$, and evaluating the Euclidean norm $||Enc_{pk}(\mathbf{c}_{a|b} - \mathbf{c}_{b|a})||$ of the result as before. The remaining parts of the cylinder similarity $\gamma(a, b)$ are calculated in plaintext, and the method is repeated $m_1 \cdot m_2$ times for m_1 cylinders in the probe and m_2 cylinders in the reference template. The global comparison score is consolidated using local similarity sort [58] and is compared against a threshold that determines whether to accept or reject the verification attempt.

An overview of the workload of homomorphic operations is summarized in Table A.2. Note that the computation of one Euclidean norm requires one homomorphic subtraction and multiplication as well as $n - 1$ additions and rotations by one position [44], where $n = 1536$ is the fixed number of cells in each cylinder. We account for the encryption of the reference template during enrolment, such that only the encryption of the probe template remains. To complement Table A.2, Table A.3 gives the relative cost of the FHE operations.

A.4 Experimental Evaluation

In this Section, we give an experimental evaluation of our proposed system as well as a security analysis according to ISO/IEC 24745 [149]. Further, we compare the performance of our system against the state of the art.

A.4.1 Performance

The experiments have been conducted on an Ubuntu server with version 1.13.0 ubuntu1.1 with 4GHz CPU and 128GB RAM. The proposed system has been evaluated on the publicly available MCYT database [199] containing fingerprint images of 330 subjects with 12 samples of each finger per subject. For feature extraction of the ISO/IEC 19794-2 [151] minutiae templates, the SourceAFIS [259] implementation was used. The MCC [58] algorithm was implemented in C++ based on the original paper without any further optimisations or simplifications. For the implementation of the FHE scheme, the PALISADE library [227] providing the CKKS [65] encryption scheme was used.

The recognition accuracy of our implementation for the baseline system and the protected system is shown in Figure A.2. The biometric performance of the protected system is not impacted through the application of FHE, as all computations are carried out in the same manner as in the baseline system, with the difference being the computation of ciphertexts on contrast to the unencrypted data in the baseline system. As the FHE scheme is able to operate on floating point numbers, no simplification or quantisation was need for our approach. This stands in contrast to other schemes [125, 169], where accuracy loss has to be accepted in order to accommodate the chosen encryption scheme.

Note that the contribution of our work is independent of the biometric performance of the baseline system, which could vary depending on the database used. Instead, the contribution of our proposed system lies in the unimpaired accuracy after the application of BIP, as CKKS is currently the only FHE scheme known to operate on floating point numbers directly [65].

Transaction times for the proposed system are presented in Table A.4. Note that transaction times for the baseline system can be considered negligent in comparison, as they are lower than 50ms throughout all system components on comparable hardware [58]. For the computational performance of the protected system, the relevant metric is the number of cylinders that need to be compared, which corresponds to the number of minutiae in the probe and reference template. In the evaluated database, the median number of minutiae per template was 35, with the lowest number of 6 and highest of 100 minutiae, both of which can be traced back to poor sample quality. The average num-

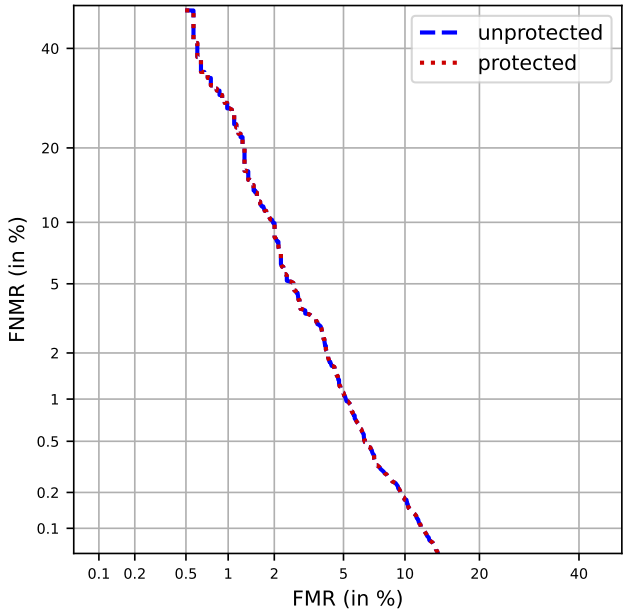

Figure A.2: Detection error trade-off curve for the evaluated MCYT [199] database.

Step	**Cylinder**	**Template**
Key generation	—	0.08
Enrolment	—	0.53
Verification		
Probe encryption	—	0.53
Direct. diff.	0.004	4.13
Common validity	0.017	25.12
Nom. + Denom.	3.80	11410.38
Total	—	11525.03

Table A.4: Transaction times for the proposed system in seconds.

ber of cylinder comparisons for one verification can therefore be extrapolated as $35 \cdot 35 = 1225$.

Evidently, the obtained execution times show that the system is not practical in real-life applications, with a verification transaction taking approximately 192 minutes. The main bottleneck is the computation of the Euclidean norms. This has already been recognised as a challenge in biometric systems [31]. Within the calculation of the norms, the most costly operation is the rotation of ciphertexts, as can be derived from Table A.3.

A.4.2 Security Analysis

We evaluate the protected system with respect to the requirements defined in ISO/IEC 24745 [149]. Firstly, unlinkability in the protected system is given through the chosen-plaintext security of the applied CKKS scheme. By the fresh random component generated for every encryption operation, even two ciphertexts computed from the exact same template look indistinguishable from a random input to an attacker. Therefore, it is not possible for an attacker to link ciphertexts corresponding to a certain data subjects to any other ciphertext within our proposed system, or any other BIP system the subject is enrolled in.

Similarly, the CKKS scheme yields renewability, as a template from the same instance can be re-encrypted and still be used securely in the system. In case the template is no longer available in plaintext form, or decryption is not possible for security reasons, an encryption of 0 can be homomorphically added to the previously stored reference to ensure a newly randomized representation of the ciphertext [28].

Finally, irreversibility of the protected templates is guaranteed through the hardness of the Ring-LWE problem, which the security of the CKKS scheme builds upon. Notably, this assumption only holds true for correct parameter choices [11], which are enforced within the PALISADE library [227].

A.5 Conclusion

Recent standards have placed emphasis on the long-term protection of biometric data. Therefore, this work has evaluated the application of post-quantum secure FHE on minutiae-based fingerprint comparison. The challenge of minutiae-based comparison lies in the variable length of the templates, absence of an inherent order, and thereby more complex comparison which requires conditional statements before a global comparison score can be obtained. In a case study and experimental evaluation, it has been shown that it is not yet practical to

evaluate such algorithms using FHE. The computational overhead of FHE is expected to decrease with further research in cryptography, while at the same time more efficient representations of biometric data need to be found that do not impair the recognition accuracy. In this regard, recent works based on deep neural networks have reported significant improvements for fixed-length fingerprint representation [104]. Until efficient post-quantum protection for high-accuracy fingerprint representations has been developed, classically secure HE or post-quantum secure MPC should be considered.

Paper B

HEBI: Homomorphically Encrypted Biometric Identification

Pia Bauspieß, Marcel Grimmer, Cecilie Fougner, Damien Le Vasseur, Thomas Thaulow Stöcklin, Christian Rathgeb, Jascha Kolberg, Anamaria Costache, and Christoph Busch

Published at International Joint Conference on Biometrics (IJCB), 2023

Abstract

Biometric data stored in automated recognition systems are at risk of attacks. This is particularly true for large-scale biometric identification systems, where the reference database is often accessed remotely. A popular approach for the protection of the stored templates is homomorphic encryption, which grants privacy protection while maintaining the biometric performance of the unprotected system. However, it introduces a significant computational overhead that can render identification transactions infeasible. To reduce this workload, biometric indexing in the encrypted domain has become a recent research interest. In this work, we show that in such schemes, auxiliary indexing data can leak additional privacy-sensitive information that violate standardized requirements for biometric template protection. In response to this leakage, we propose a novel framework HEBI that protects biometric indexing approaches at a post-quantum security level while requiring a computational effort of only 0.12 milliseconds per cluster.

B.1 Introduction

Biometric data allow for an irrevocable identification of individuals over several decades [166]. Therefore, biometric data need to be considered sensitive data requiring long-term protection, even more so than passwords or authentication tokens that can be exchanged upon a security breach. To ensure this protection, the ISO/IEC 24745 standard on biometric information protection [149] defines the following requirements: *i) unlinkability*, two protected templates stored in different applications cannot be linked to the same subject, *ii) renewability*, new templates can be created from the same source if the previously stored reference was leaked without the need to re-enrol, and *iii) irreversibility*, it is impossible to reconstruct original samples given only protected templates. Furthermore, both the computational and biometric performance (i.e., accuracy) of the unprotected system should be preserved.

In biometric identification, where a 1:N search against a large database is performed, biometric templates are at particular risk as reference databases are maintained for long time spans. For example, this is true for criminal databases held by law enforcement agencies or for national citizen registration [256]. In addition, these databases are static targets of attack, as their large storage requirements do not allow for agile changes to their physical security.

Recently, biometric identification protected through *Fully Homomorphic Encryption* (FHE) has been explored to mitigate these security risks [31, 91, 103]. While this approach grants cryptographically sound protection of the biometric tem-

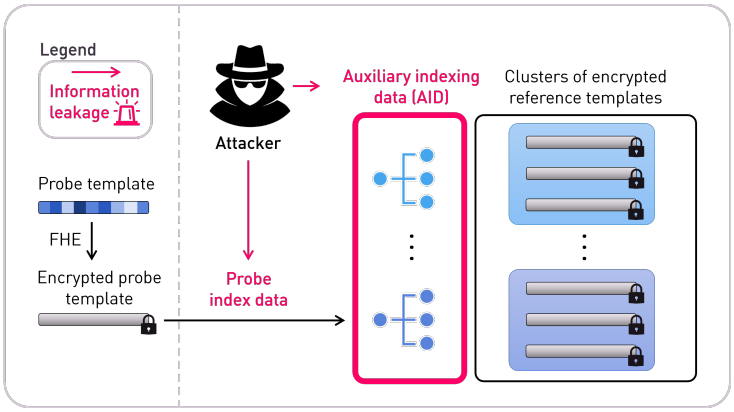

Figure B.1: Biometric information leakage in indexing schemes on encrypted reference databases: an attacker can observe privacy-sensitive information from indexing data, e.g., soft-biometric attributes such as gender of the probe and reference subjects.

plates, it comes with a significant overhead in computational workload. For large-scale databases, workload reduction strategies need to be applied to achieve practical biometric identification systems. Workload reduction strategies have been categorized into two main classes [92]: feature transformation and preselection. Preselection approaches offer a significant speed-up through selecting a smaller subset of the enrolment database that contains the reference identifier with high probability. Using an index string i common to a subset of enrolled references $\mathcal{C}_i$, preselection can be achieved in $\mathcal{O}(1)$ and is therefore efficient.

However, a key challenge with these approaches is the continuous protection of data subject privacy under preselection, i.e., ensuring that the preselection procedure and its outcome do not reveal any information about the underlying subject, or infringe on the unlinkability of the system. This vulnerability is depicted in Figure B.1. It is important to note that the encryption of the feature vectors alone is not sufficient to fulfil this requirement, as the preselection algorithm can reveal additional information about the enrolled subjects, e.g., soft-biometric characteristics such as the gender of ethnicity of the probe and reference subjects.

The risk of information leakage shown in Figure B.1 is particularly high when biometric indexing is based on similarity measures between the enrolled subjects, e.g., in feature-based clustering approaches. These similarity measures

contained in the *Auxiliary Indexing Data* (AID) can potentially reveal sensitive information about the preselected subset such as their shared soft-biometric characteristics. For sound privacy protection in the sense of ISO/IEC 24745 [149], this information needs to be obscured in addition to the protection of the feature vectors.

To mitigate the privacy leakage in biometric indexing, we therefore propose a novel protocol HEBI that can be applied to indexing approaches in the encrypted domain. The key contributions of our work are as follows:

- **Privacy analysis.** To illustrate the significant risks that come with the use of unprotected AID, we give a privacy analysis of existing approaches. We show that we were able to reconstruct the gender and ethnicity of enrolled subjects based only on AID, which must be considered a severe security risk.
- **Formalization of information leakage.** Further, we give a formalization of information leakage in biometric indexing that indicates that such a leakage exists in arbitrary biometric indexing schemes. We use this formalization as further motivation for our work, in addition to the experimental analysis.
- **The novel HEBI protocol.** As our main contribution, we present the HEBI protocol for secure biometric indexing in the encrypted domain. Through the use of lattice-based cryptography [39, 65], our protocol provides post-quantum security in storage, preselection and comparison. We give an experimental evaluation that shows that HEBI can be applied in real-world operational systems at a cost of only 0.12 additional milliseconds for the the post-quantum secure retrieval compared to unprotected preselection systems. At the same time, the biometric performance of the underlying indexing approach is not impacted by the applied cryptographic protection mechanisms.
- **Security analysis.** We provide a comprehensive security analysis of our protocol and show how it mitigates the flaws of unprotected approaches, thus giving full post-quantum security to biometric data under preselection.

The remainder of this article is structured as follows: Section B.2 discusses works that are closely related to ours, before we analyse of the privacy leakage in a previously proposed privacy-preserving biometric indexing scheme in Section B.3. In Section B.4, we introduce more technical cryptographic background information. From this, we present our novel HEBI protocol in Section B.5 that alleviates the presented privacy risks. Section B.6 gives experimental results and a security analysis. Finally, we draw our conclusions in Section B.7.

B.2 Related Work

Workload reduction in homomorphically encrypted biometric identification systems has recently been achieved with post-quantum security [31, 103]. However, both of these works were only based on feature transformation, such that an exhaustive search requiring a linearly increasing costs remains. It is important to note that our HEBI protocol can integrate such feature transformation approaches seamlessly and therefore allows for further improvements in large-scale biometric identification systems.

The cryptographic concept of homomorphic search has previously been applied to biometric identification in [278]. In their work, the authors use the search scheme as a replacement for FHE rather than an additional protection layer for the preselection step. In order to realize homomorphic search on the feature vectors directly, strong statistical assumptions about the feature representation are required, which do not generalize over different modalities. Another recent work [276] applied homomorphic search for biometric authentication instead of identification. Most recently, [33] applied homomorphic search for preselection on an encrypted reference database. However, our HEBI protocol differs non-trivially from the proposal in [33] in several aspects. Firstly, the work by [33] can only be considered as proof-of-concept, as a handcrafted preselection approach is utilized in their work, which underlies the unrealistic assumption of perfect ground truth. In comparison, HEBI is designed for real-world indexing approaches that allow for a meaningful analysis of the overall biometric performance. Secondly, [33] apply a binning approach that does not trivially generalize to other application scenarios apart from their own, while HEBI enables efficient and secure cluster generation independent of the indexing algorithm. Finally, our work offers an extensive analysis of the risk of preselection independent of the concrete indexing approach and shows how to mitigate these risks in a universal approach.

The application of unprotected biometric indexing to biometric identification [90, 137, 195, 200, 205, 229, 246, 261] will be discussed at length in the following Section. These are the schemes our HEBI protocol improves upon through an additional layer of protection during the preselection step. Notably, the choice of protection mechanism for the reference database is independent of the HEBI preselection protocol, though we adhere to FHE-based protection in our work. In addition, HEBI does not impair the originally given biometric performance of the above works.

B.3 Privacy Analysis of Biometric Indexing

Biometric indexing as depicted in Figure B.1 has been applied in a number of recent research works, among others [90, 137, 195, 200, 205, 229, 246, 261]. In this Section, we give further intuition to the privacy implications of such approaches through probability theory.

B.3.1 Formal Model

In this analysis, we investigate the relation between the enrolled reference feature vectors $\{r_j\}_{j=0}^{N-1}$ for a number of reference subjects N and the auxiliary indexing data (AID) represented by index strings $\{i_k\}_{k=0}^{K}$, where K denotes the number of clusters or index strings in the given scheme. We define that every reference feature vector r_j is assigned one and only one index string i_k, while one index string i_k clusters several references (i.e., $K < N$). Upon an identification transaction, a probe feature vector p is extracted from a presented probe sample, and the corresponding index i_k is determined. Then, only the reference features vectors associated with i_k are compared to p in the encrypted domain.

For the formalization of privacy leakage in such indexing schemes, we utilize the information-theoretic concept of *mutual information* $I(X;Y)$, which is defined as

$$I(X;Y) = D_{KL}(P_{(X,Y)}||P_X \otimes P_Y), \tag{B.1}$$

where X and Y are random variables and D_{KL} denotes the Kullback–Leibler divergence [76]. The mutual information can further be expressed in terms of entropy [236]:

$$I(X;Y) = H(X) - H(X|Y), \tag{B.2}$$

where $H(X)$ is the marginal entropy of X and $H(X|Y)$ is the conditional entropy of X given Y. Let $\{X\}_j$ be the variable family that represents the reference feature vectors and $\{Y\}_k$ be the variable family that represents the index strings. In a meaningful indexing scheme, it holds that

$$I(X_k, Y_{i_k}) > I(X_k, Y_{i_m}), \tag{B.3}$$

i.e., the mutual information between the reference feature vector r_j associated with index string i_k should be greater than the mutual information between the same reference feature vector r_j and a different cluster associated with an index string i_m. Otherwise, r_j would be associated with i_m instead. From Equation B.3, it follows that $H(X_k|Y_{i_m}) > H(X_k|Y_{i_m})$. As $H(X_k|Y_{i_m})$ cannot be smaller than 0, it follows that $H(X_k|Y_{i_k}) > 0$. At the same time, the similarity of index strings

does not correspond to the full feature vectors, which would yield no advantage over an exhaustive identification search. Therefore,

$$H(X_k) > H(X_k|Y_{i_k}) > 0, \tag{B.4}$$

and consequently,

$$I(X_k; Y_{i_k}) = H(X_k) - H(X_k|Y_{i_k}) > 0, \tag{B.5}$$

meaning that there is mutual information contained between the feature vectors and index strings. This mutual information defines the leakage of biometric information, which allows for attacks on the probe and reference subjects that can violate their privacy. Indeed, it has been shown that auxiliary data in biometric systems can lead to privacy risks in other applications, e.g., biometric cryptosystems [231]. However, we emphasize that our formal model is not intended to be used as a concrete metric, as mutual information is hard to calculate precisely. Instead, it serves as a logical argument for the existence of privacy leakage in AID.

More empirically, index strings are commonly constructed such that they allow for a clustering of the reference feature vectors based on a more general measure of similarity than the exact comparison between feature vectors. In some approaches [90, 200, 229], the index strings are even derived from the feature vectors directly, representing a down-sampled representation of one or more feature vectors. In the following, we show how to concretely extract privacy-sensitive information from such representations.

B.3.2 Case Study

To illustrate the risks of soft-biometric leakage in biometric indexing in a case study, we analyze the recent work of [200], which is one of the works relying on unprotected index strings discussed above.

In their work, the authors generate a look-up table of short binary strings, or *stable hashes*, which represent distinct clusters of reference templates. They present different methods of obtaining these stable hashes, all of which are based on the feature representations of the enrolled references. In our evaluation, we focus in the first of their proposed approaches, which is the established k-means clustering technique [184]. During the enrolment phase, the clustering algorithm is trained on the enrolment database, which is subsequently encrypted using FHE. The protected templates are stored in the database alongside the look-up table of stable hashes, which in the case of k-means clustering are a binary representation of the cluster centers, or centroids. Upon an identification transaction,

the distance of the probe feature vector to all centroids is calculated, and the closest centroid is determined to be the probe stable hash. Then, the reference subjects with the same stable hash are extracted from the enrolment database, a homomorphic comparison of the encrypted probe feature vector against the encrypted references is computed, and the decision is revealed to the client that initiated the transaction [200].

The advantage of this indexing approach is the error-correcting capability of the clustering approach, which allows for an exact comparison of the stable hashes and is therefore very efficient. The retrieval cost of the look-up operation is constant at $\mathcal{O}(1)$ and can be considered negligible compared to the cost of the homomorphic operations. Furthermore, the low preselection error even on challenging datasets makes the approach in [200] attractive.

However, the vulnerability of the approach with regard to the reference subjects' privacy lies in the stable hash look-up table, which is stored alongside the enrolment database. As argued above, it can be expected that the stable hashes encode information about the probe and reference subjects to some degree, which could be privacy-sensitive information. For example, soft-biometric similarities to the subjects in one cluster could be revealed, which would constitute a violation of ISO/IEC 24745 [149]. Disclosure of soft-biometric data related to the ethnic origin is a breach of the European Union's General Data Protection Regulation [109].

To confirm our hypothesis, we conducted an experimental evaluation of the privacy leakage in the system presented in [200]. For this evaluation, we selected 3,165 samples of 533 subjects from the Face Recognition Grand Challenge (FRGC) database [206] that are compliant with the International Civil Aviation Organization's face image quality requirements for machine-readable travel documents. The code for the stable hash generation from k-means clustering [184] has been provided by the authors to facilitate the reproducibility of their results. In terms of parameters, we followed the original work with $P = 1$ subspaces and $K = 64$.

Figures B.2 and B.3 show the distribution of ethnicity and gender of the 64 clusters. For this analysis, ground truth labels for the image samples were hand-annotated, such that a high accuracy in the labelling can be assumed compared to soft-biometric feature extractors [95]. From the visualization of the distributions, it becomes evident that there exists pooling of soft-biometric characteristics within both dimensions of ethnicity and gender. This can be for example observed in clusters clusters 36, 38 and 39, which exclusively contain female subjects, while clusters 3, 9 and 11 only contain male subjects. Similarly, clusters 10 and 11 exclusively contain Asian subjects, while clusters 21, 22, 23, 42,

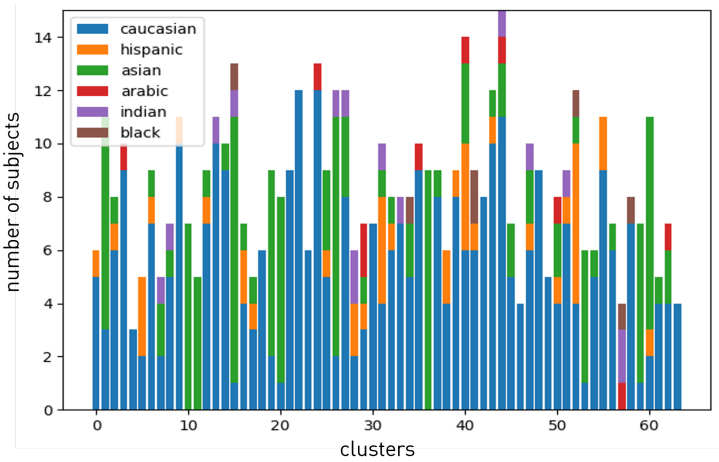

Figure B.2: Distribution of ethnicities over the clusters derived from an ICAO-compliant subset of the FRGC database [206].

46 and 63 only contain Caucasian subjects. While these characteristics are not perfectly separated over all clusters, it is particularly concerning that the clustering effectively exposes underrepresented subgroups. A prominent example is cluster 10, which contains only female Asian subjects. An attacker observing the stable hash corresponding to cluster 10 can therefore with high probability deduct the gender and ethnicity of the probe subject and the reference subjects stored alongside that stable hash.

To extend our analysis, we further evaluated a synthetic face image generation from the centroids to approximate the average features of the subjects in the clusters and their similarity to the synthetic approximation for the respective cluster. We leveraged the StyleGAN3 generator [163] pre-trained on the FFHQ database [164] that includes more than $70,000$ face images with diverse ethnicities, gender labels, and other facial characteristics. To reconstruct latent representations and subsequently derived representative face images from each stable hash ($\mathbf{s}$), we trained a fully connected neural network ($\mathcal{M}$) that maps each stable hash into the semantic manifold of the StyleGAN3 intermediate latent space. We froze the generator ($\mathcal{G}$) weights during training to preserve its capability to generate photo-realistic face images. Further, we applied a simple *mean squared error* loss function to minimize the difference between the reconstructed face images $\hat{x} = \mathcal{G}(\mathcal{M}(\mathbf{s}))$ to the randomly drawn face images x of their corresponding cluster.

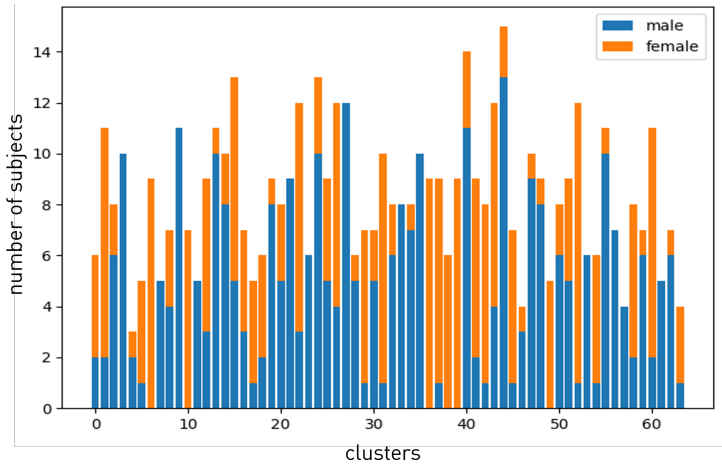

Figure B.3: Distribution of genders over the clusters derived from an ICAO-compliant subset of the FRGC database [206].

For this experiment, the FRGCv2 training subset has been reduced such that each stable hash is assigned with only one face image per identity. This setting prevents the mapping network from oscillating due to the high intra-subject variance. For the optimization of $\mathcal{M}$, the StyleGAN3 truncation factor was set to 0.75, enabling the generation of face images with stable quality. We adopted the Adam optimizer settings from [164] and increased the learning rate to 0.01 to accelerate the training process. The results of this evaluation are shown in Figure B.4.

In Figure B.4, the reconstructed latent representations corresponding to cluster 10 are depicted alongside a selection of bona fide sample images from that cluster, which contains only female Asian subjects. The reconstructed images are based on incrementally scarce training data to show that our GAN-based approach generalizes even in an open-set scenario. The closest approximation has been trained on cluster 10 alone, and cannot be considered a realistic attack. Both the closed-set and the open-set training scenario excluding cluster 10 continue however to show significant similarities to the original identities. Most importantly, the soft-biometric characteristics of gender and ethnicity are preserved. A breach of the latter in particular constitutes a GDPR [109] violation and must be prevented.

To conclude this analysis, significant privacy leakage has been found in the indexing approach by [200]. However, the overall indexing scheme is of high rele-

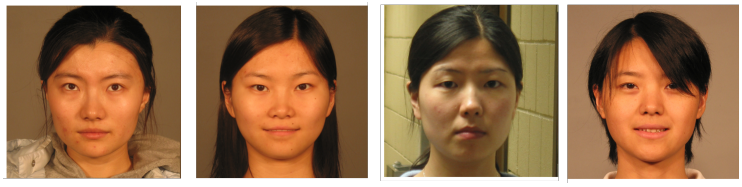

(a) Bona fide samples of subjects from cluster 10.

(b) Left to right: Reconstructions based on untrained, trained on cluster 10 only, trained on all clusters, trained on 70% of all clusters (excluding cluster 10) StyleGAN approximations.

Figure B.4: Comparison of bona fide FRGCv2 samples of cluster 10 and Stlye-GAN presentation attack approximations of cluster 10.

vance to the problem of workload reduction for large-scale biometric identification, as it benefits from a high biometric performance and is therefore desirable to apply.

Looking towards the cryptographic protection of indexing approaches such as [90, 200, 229], the component of the index string that allows for the privacy leakage is their deterministic nature, i.e., in the case of [200], similar feature vectors will always be mapped to the same stable hash. In the remainder of this paper, we are therefore proposing a transformation of this deterministic preselection approach to a non-deterministic preselection, where similar feature vectors are mapped to randomized outputs that look indistinguishable to an attacker. At the same time, they allow for the correct retrieval of the corresponding reference subjects, such that the biometric performance of the indexing approach is not impacted.

B.4 Preliminaries

B.4.1 Fully Homomorphic Encryption

Homomorphic encryption describes a cryptographic technique that allows for the evaluation of functions on encrypted data. More precisely, we call a public-

key encryption scheme homomorphic if

$$\mathrm{Enc}(pk_H, x \odot y) = \mathrm{Enc}(pk_H, x) \odot \mathrm{Enc}(pk_H, y). \tag{B.6}$$

More recently, *Fully Homomorphic Encryption* (FHE) has become practical for application in certain use cases. Following the groundbreaking work by Gentry [121], different schemes have established themselves with respect to their different properties. One of these is the CKKS [65] scheme, which provides the useful advantage of computing on high-precision approximations of floating point numbers directly, where other schemes require integer quantisation [52, 112] or binarisation [69]. In terms of the encrypted comparison of biometric feature vectors, this means that the underlying data does not need to be altered, and no information from the biometric comparison is lost. Therefore, the computations on encrypted templates correspond directly to computations on the unprotected templates, and the biometric performance remains unimpaired.

The security of many FHE schemes, including CKKS, is based on the Ring-Learning with Errors (R-LWE) problem, which is assumed to be secure against attacks implemented on a quantum computer [183]. These cryptosystems therefore provide a high level of protection to the biometric data, and in particular, long-term protection over several decades according to the current basis of knowledge and expectations in the field of cryptography [11].

B.4.2 Public-Key Encryption with Keyword Search

In addition to the protection of the feature vectors, the privacy analysis in Section B.3 has shown that the indexing and retrieval during the preselection process requires additional protection. A recent work on face identification [33] has proposed the use of *Public-Key Encryption with Keyword Search* (PEKS) for the protection of semantic soft-biometric keywords. In this work, we apply this technique to generic biometric indexing approaches.

The cryptographic basis of PEKS lies in *Identity-Based Encryption* (IBE), which was first introduced by Boneh and Franklin in 2001 [46]. Building on this idea, PEKS was proposed as a means of creating ciphertexts for specific semantic keywords instead of identities [45]. In the typical application scenario, a PEKS scheme is used to create an encryption of a keyword together with a corresponding trapdoor. This pair of cryptographic objects can be subjected to a publicly available test function which reveals no information except for the binary decision outcome of the similarity of the underlying keyword of the ciphertext and trapdoor.

A PEKS scheme [39] is defined as a tuple of four algorithms PEKS = (KeyGen, PEKS, Trapdoor, Test):

- $(pk_S, sk_S) \leftarrow \text{KeyGen}(1^k)$: On the input of the security parameter k, this algorithm outputs the public and secret key pair (pk_S, sk_S).
- $s_w \leftarrow \text{PEKS}(pk_S, w)$: On the input of the public key pk_S and a keyword $w \in \{0,1\}^*$, this algorithm outputs a searchable ciphertext s_w.
- $t_w \leftarrow \text{Trapdoor}(sk_S, w)$: On the input of a secret key sk_S and a keyword $w \in \{0,1\}^*$, this algorithm outputs a trapdoor t_w.
- $b \leftarrow \text{Test}(t_w, s_w)$: On the input of a trapdoor $t_w = \text{Trapdoor}(sk_S, w')$ and a searchable ciphertext $s_w = \text{PEKS}(pk_S, w)$, this algorithm outputs a bit $b = 1$ if $w = w'$, and $b = 0$ otherwise.

More recently, PEKS has been implemented based on lattice-based IBE [97] to create lattice-based PEKS [39]. Compared to the original construction, lattice-based PEKS has high computational efficiency and provides post-quantum security through R-LWE [183]. As an important property to the application in this work, PEKS ciphertexts are constructed using a random component, yielding non-deterministic encryption. In the following Section, we will detail how this property ensures privacy protection when applied to biometric indexing.

B.5 The HEBI Protocol

In this Section, we present our novel HEBI protocol for biometric indexing in the encrypted domain. The protocol can be applied to any existing biometric indexing approach that clusters enrolment biometric references to prevent the leakage of sensitive information about the data subjects.

B.5.1 Setting

The HEBI protocol is executed between three parties: A *client* device, a *Database Server* (DS) and a *Trusted Third Party* service (TTP). All three parties are considered in the semi-honest security model, where they may aim to gain information about the data they are exchanging, but are not assumed to deviate from the given protocol. This is an established security assumption in remote biometric authentication [126, 170, 270].

B.5.2 Enrolment

During the enrolment phase, two separate setup operations are performed: initialisation of the encrypted indexing algorithm and the homomorphic encryption of the enrolment database.

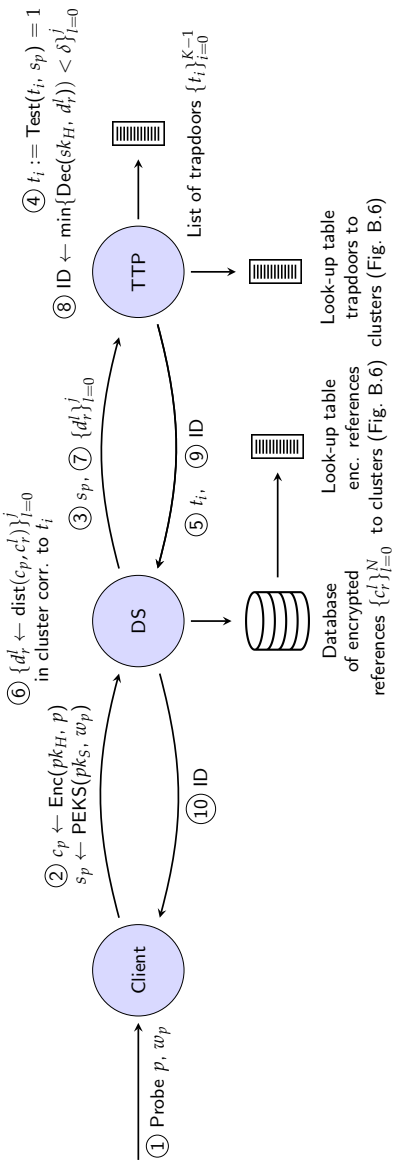

Figure B.5: Identification transaction for encrypted preselection in the HEBI protocol.

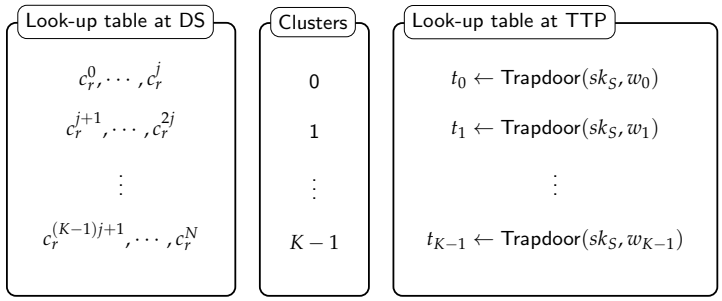

Figure B.6: HEBI look-up tables generated at enrolment.

The indexing algorithm is assumed to require some precomputation on an unencrypted enrolment database [200]. In our protocol, we allow for this precomputation to be conducted during an offline phase prior to the deployment of the system, where the unprotected templates are not exposed to potential attacks. As a result of the indexing algorithm, each biometric reference r will be assigned an index, or cluster, i which can be of arbitrary data representation. If the clustering algorithm does not produce balanced clusters, i.e., the number of subjects per cluster is not consistent, the clusters are padded with random feature vectors to be of equal size.

Once the clusters have been established, the PEKS framework can be initialised. First, TTP generates a number of random PEKS keywords $\{w_i \mid 0 \leq i \leq K-1\}$, where K is the total number of clusters, and fixes a mapping M between the clusters and PEKS keywords , which is made publicly available. Note that the mapping M of clusters to PEKS keywords must be indicated by the clusters' (arbitrarily assigned) order instead of the semantic index string i that could potentially reveal privacy-sensitive information. By making the mapping publicly available, the PEKS keywords do not act as additional secret keys in the system.

From the PEKS keywords, two look-up tables are generated. At TTP, a trapdoor $t_p \leftarrow \mathsf{Trapdoor}(sk_S, w_i)$ is computed and stored for every cluster using the PEKS secret key sk_S. At DS, a mapping of encrypted references to clusters is stored, again based on any order of the clusters without using the index i as the identifier. An overview of the look-up tables is given in Figure B.6.

For the setup of the encrypted enrolment database, TTP generates and stores a key pair of the homomorphic encryption scheme (sk_H, pk_H) and makes pk_H available to the client and DS. For a reference feature vector r, the client can enrol a data subject by computing $c_r \leftarrow \mathrm{Enc}(pk_H, r)$ and sending c_r encrypted

biometric reference to DS. Since the assignment of subjects to clusters is initially fixed, coefficient packing can be applied to facilitate further workload reduction [31].

B.5.3 Identification

During an identification transaction in HEBI, the client captures a probe sample and obtains its feature representation p. The client determines the index i_p of the probe with respect to the applied indexing algorithm and uses the public mapping M to determine the corresponding PEKS keyword w_p. Using the public key pk_H of the HE scheme, the client encrypts the probe feature vector by computing $c_p \leftarrow \text{Enc}(pk_H, p)$. It further computes the encrypted probe index $s_p \leftarrow \text{PEKS}(pk_S, w_p)$, and sends c_p and s_p to DS, which forwards s_p to TTP.

Upon receiving s_P, TTP determines the corresponding trapdoor t_i for which $\text{Test}(t_i, s_p) = 1$ holds true. Using the look-up table mapping trapdoors to clusters (see Figure B.6), TTP sends the cluster identifier to DS, where the homomorphic comparisons are computed between the encrypted probe c_p and the encrypted references $\{c_r^l\}_{l=0}^{j}$ in the cluster corresponding to t_i. The encrypted comparison scores are sent to TTP, which decrypts them and determines the identification outcome, which is forwarded to the client. Note that throughout this transaction, DS and TTP do not have access to unprotected feature vectors or the index strings i that could reveal sensitive information. An overview of an identification transaction is given in Figure B.5.

Our HEBI protocol can be seen as an independent layer of protection to arbitrary indexing schemes. Furthermore, it can also be combined with interchangeable template protection approaches for the feature vectors themselves, e.g., different FHE schemes or irreversible feature transformations. It is therefore versatile in its application and can be considered for applications beyond face recognition.

B.6 Experimental Evaluation

To show the practicality of our HEBI protocol, we give an experimental evaluation for the application to stable hashes [200]. By applying the additional layer of security, the privacy concerns outlined in Section B.3 will be mitigated.

The experiments were conducted on the same subset of 533 subjects of the FRGCv2 [206] database with 3,165 ICAO-compliant samples. In addition, 529 subjects with 1413 samples from the FERET [207] database of ICAO-compliant

quality were used for the evaluation. From the samples, features are extracted with the open-source feature extraction model ArcFace [86] which produces face templates of 512-dimensional floating point vectors with documented good performance on the used dataset [31]. For the stable hash generation using k-means clustering, the parameters $P = 1$ subspace and $K = 64$ clusters are chosen in accordance with the size of the database. The experiments where implemented in Python and C++ on macOS Monterey 12.4 with an M2 processor at 3.50 GHz CPU clock frequency.

For the homomorphic operations, the CKKS [65] FHE scheme was applied, as it does not impair the biometric performance. The implementation of the state-of-the-art FHE library OpenFHE [22] was applied, where CKKS parameters corresponding to 128 bits of security were chosen [11]. For further workload reduction, coefficient packing for a quadratic speed-up as previously proposed by [31] was applied, showing the compatibility of HEBI with such approaches. The squared Euclidean distance was applied as the comparison metric. For the lattice-based PEKS scheme, the implementation by [39] was used.

B.6.1 Results

The results of the experimental evaluation are presented in Tables B.1 and B.2. In terms of execution times (Table B.1), it can be seen that the majority of the workload is absorbed by the FHE comparisons on the encrypted feature vectors, an observation which is consistent with related work [93, 95, 103]. It is important to note that this workload can differ for different FHE schemes and has generally been found to be lower for integer-quantised and binary encryption, which introduces a trade-off with the biometric performance [171]. The baseline and preselection accuracy can be seen in Table B.2, were a closed-set identification scenario was evaluated. Aside from this concrete instantiation however, we stress that HEBI is independent of the concrete preselection procedure and inherits and maintains the accuracy of the underlying indexing algorithm in question.

The main focus of this evaluation is the overhead of a secure indexing using HEBI over unprotected preselection. From Table B.1, it can be derived that the protected preselection using PEKS takes 7.69 milliseconds for 64 clusters or 0.12 milliseconds per cluster. As the cost for the preselection scales linearly with the number of clusters rather than the size of the enrolment database, this cost is expected to grow significantly slower than the cost for an exhaustive identification search. For larger databases, the original work on stable hashing [200] proposes a number of $K = 1024$ clusters, the cost of which can be approximated at 123.04 milliseconds, which can be considered real-time. Depending on the in-

System function	Time (ms)
Probe stable hash generation	0.28
Probe encryption	2.27
PEKS search	7.69
FHE comparisons	9,996.00
Total	10,006.24
Baseline (exh. search)	334,891.00

Table B.1: HEBI execution times for 533 subjects and 64 clusters.

Database	Enroll Samples	Search Samples	False Negative	True Positive	Preselection Accuracy	Baseline Accuracy
FERET [207]	529	884	19	865	0.9785	1.0000
FRGCv2 [206]	533	2,632	207	2,425	0.9214	0.9971

Table B.2: Accuracy of the stable hash clustering [200] for the FERET [207] and FRGCv2 [206] databases and $K = 64$ clusters.

dexing algorithm used, there exists a trade-off between the number of clusters, the preselection error, and the number of biometric references per cluster. Overall, it becomes evident however that the lattice-based PEKS scheme adds only a negligible overhead to the identification system at less than 8% of the total cost, while providing post-quantum protection under preselection. Compared to the baseline system, the workload is reduced down to 3%. The communication cost for HEBI consists of 2.66MB for a CKKS public key, 267.4KB for a CKKS ciphertexts, 27.2KB for a PEKS public key, 52KB for a PEKS ciphertext, and 27KB for a PEKS trapdoor.

B.6.2 Security Analysis

The security of both the FHE and the PEKS scheme are based on the R-LWE [183] problem, which is assumed to be post-quantum secure. The HEBI protocol maintains the post-quantum security through all steps of the identification transaction, including preselection. Contrary to unprotected indexing approaches such as [90, 200, 229], the PEKS ciphertexts are generated in a non-deterministic manner, which makes them indistinguishable over the given clusters. A privacy attack as discussed in Section B.3 is thereby prevented.

With regards to the requirements formulated in ISO/IEC 24745 [149], irreversibility is given through the security assumption of R-LWE [183]. Unlinkability and

renewability can be derived directly from the IND-CPA security of both the CKKS [65] and PEKS [39] schemes, i.e., the indistinguishability under chosen plaintext attacks. Through this property, an attacker cannot distinguish between an encryption of 0 and and encryption of 1. In biometric identification, this extends to the indistinguishability of encrypted templates: even if an attacker gains access to two encryptions of the same template, they cannot be distinguished from arbitrary inputs in a feasible manner. The same property holds for the encryption of index strings through the PEKS scheme. Therefore, it is not possible for an attacker to link data subjects to other subjects enrolled under the HEBI protocol or another system.

Finally, the performance preservation of HEBI is given through the application of CKKS [65] and PEKS [39], as neither scheme impairs the biometric performance. The operations in the encrypted domain correspond directly to the operations in an unprotected biometric system. In terms of computational performance of HEBI, our experimental evaluation has shown that the overhead of the PEKS scheme is small, while a trade-off between the preselection error and homomorphic workload persists. Further limitations of HEBI include the assumption of the semi-honest adversary model. Although this is an established assumption in biometric template protection, it does not fully reflect the capabilities of real-world adversaries. In addition, we have only evaluated the efficiency of HEBI for fixed-length feature representations, which can be considered a limitation.

B.7 Conclusion

This work firstly revealed that indexing schemes can leak privacy-sensitive biometric information. Motivated by this, we introduced the HEBI protocol for biometric indexing in the encrypted domain. Index strings in biometric identification systems allow for the reconstruction of privacy-sensitive information about the data subjects, which stands in violation to ISO/IEC 24745 as well as the GDPR. As a solution to this problem, HEBI gives post-quantum secure protection to the feature vectors alongside their auxiliary indexing data in storage, preselection, and comparison. HEBI is independent of the indexing algorithm and protection of the enrolment database and adds only negligible computational overhead per indexing cluster.

Paper C

MT-PRO: Multibiometric Template Protection Based On Homomorphic Transciphering

Pia Bauspieß, Chiara-Marie Zok, Anamaria Costache, Christian Rathgeb, Jascha Kolberg, and Christoph Busch

Published at IEEE International Workshop on Information Forensics and Security (WIFS), 2023

Abstract

Reliable authentication of individuals is the foundation of trusted digital interaction. Biometrics lend themselves ideally to this goal. However, biometric data must be protected under computation according to European laws and international standards. Over the past ten years, fully homomorphic encryption has become a popular tool for biometric template protection. However, it comes with the security risk of cryptographic key material, which requires careful management and could be leaked, leaving the stored templates vulnerable to attacks. To meet this challenge, we present the novel MT-PRO protocol utilising homomorphic transciphering to improve the security of such systems against offline decryption attacks. Our protocol does not impair the biometric performance and allows for multibiometric comparisons of fixed-length feature representations. Furthermore, we evaluated our protocol on public datasets with open-source implementation available at `https://github.com/dasec/MT-PRO` and discuss its real-world application potential.

C.1 Introduction

Trustworthy digital communication requires reliable authentication mechanisms, i.e., the ability to tie a human user to their digital identity. The need for reliable authentication is present in many applications, ranging from online banking and legal transactions to telemedicine. Biometric characteristics are uniquely suited to provide such authentication mechanisms, as they allow for a persistent identification of individuals [166].

However, there exist a number of concerns regarding biometric authentication, which can be classified into two main categories: concerns about the reliability (or security) of biometric authentication, and concerns about the protection of biometric feature vectors stored and used in the system (i.e., privacy). For the privacy protection of biometric reference templates, the ISO/IEC 24745 standard on biometric information protection [149] defines clear requirements: *i) unlinkability*, two protected templates stored in different applications cannot be linked to the same subject, *ii) renewability*, new templates can be created from the same source if the previously stored reference was leaked without the need to re-enrol a subject, and *iii) irreversibility*, it is impossible to reconstruct original samples given only protected templates. Furthermore, both the computational and biometric performance (i.e., accuracy) of the unprotected system should be preserved.

Recent solutions to biometric template protection apply *Fully Homomorphic Encryption* (FHE) for encrypted storage and comparison of biometric feature vec-

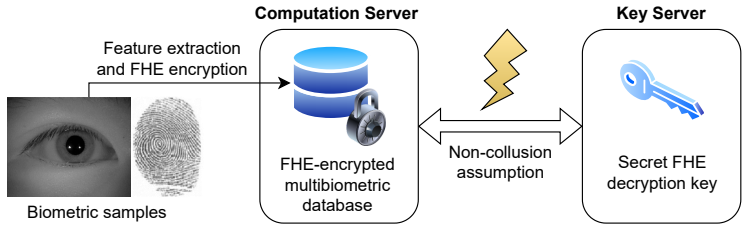

Figure C.1: Security risk in FHE-based biometric template protection: if the non-collusion assumption is violated, the encrypted reference database can be decrypted by an attacker, leaving the enrolled templates vulnerable to attacks.

tors [31, 103, 126, 270]. The established architecture for these works includes a two-server setup, where a trusted key server manages the cryptographic key material, while a computation server has access to the encrypted reference database (Fig. C.1). This scenario is typically considered in a semi-honest adversary model (aside from [28]), where the two servers must not collaborate. If they do, or an attacker gains access to the cryptographic key material in another way, the database could be decrypted and the enrolled subjects would be vulnerable to impersonation attacks.

This non-collusion assumption can be considered the weakest point in FHE-based template protection systems, as it does not reflect real-world adversary capabilities. This has lead to a decreased trust in outsourced biometric authentication compared to on-device biometric authentication, e.g., FaceID [15]. Security against attackers who have obtained secret components of a biometric system has previously only been achieved using cancelable biometrics [201], which can decrease the accuracy of the system. The accuracy of biometric comparisons however determines the reliability of biometric authentication, and thereby, its security. As biometric feature representations are noisy due to intra-class variance, they introduce the risk of false-accepts, which can lower the security of the biometric system. In response, *multibiometric* systems have received increased interest in recent years [94, 126, 219, 243]. Through the combination of multiple biometric modalities (e.g., iris and fingerprint), the false-accept rate can be lowered significantly [219], increasing the overall security level.

A secure and reliable biometric authentication system would therefore address both of the aforementioned research challenges: security and privacy. In this work, we present such a system with our novel MT-PRO protocol that utilises the cryptographic concept of *Homomorphic Transciphering* (HT) [71]. Using HT, the protected database receives an additional layer of encryption, such that the

leakage of the FHE secret does not enable a viable attack on the database. We describe our contribution as follows.

- We present the novel MT-PRO protocol for secure and privacy-preserving multibiometric verification with HT. To the best of our knowledge, this is the first application of HT to biometric template protection.
- Our MT-PRO protocol is secure against an attacker who has obtained both the protected multibiometric database and the corresponding FHE secret key. Compared to related work considering this attack model, our protocol does not impair the biometric performance of the system. We give a vulnerability analysis of established FHE-based BTP approaches with regard to these *offline attacks* and compare our work to the state-of-the-art in the field.
- We present a reproducible experimental evaluation of MT-PRO and give a comprehensive security analysis, showing how the shortcomings of current FHE-based BTP approaches have been addressed.

The remainder of this article is structured as follows: Section C.2 discusses related work and gives context to our contribution, before we define the cryptographic backbones of our work in Section C.3. As our main contribution, Section C.4 presents our proposed MT-PRO protocol for HT-based multibiometric template protection secure against offline attacks, including a vulnerability analysis of previous work. The experimental evaluation of MT-PRO is presented in Section C.5 with a security analysis, before we offer conclusions in Section C.6.

C.2 Related Work

The concept of HT has previously received interest from various research fields, including cloud computing [57] and privacy-preserving genomic comparisons [242]. However, these previous works have only used FHE schemes based on integer plaintexts, which in the context of real-valued biometric feature representations lead to accuracy loss through quantization. In comparison, our MT-PRO protocol utilises an encryption scheme that operates directly on floating point data [65], such that no accuracy is lost in the encrypted domain.

More recently, the problem of FHE-based template protection schemes secure against offline decryption attacks has been investigated in biometric research, with [201] proposing a combination of *Cancelable Biometrics* (CB) and FHE to mitigate the leakage of secret key material. However, the application of CB yields

Reference	BTP approach	Preserve accuracy	Prevent offline attacks	Post-quantum security
Canteaut et al. 2017 [57]	FHE + HT*	✗	✓	✓
Singh et al. 2018 [242]	FHE + HT*	✗	✓	✓
Boddeti 2018 [44]	FHE	(✓)	✗	✓
Otroshi et al. 2022 [201]	CB + FHE	✗	✓	✓
Sperling et al. 2022 [243]	FHE	✓	✗	✓
Ours	FHE + HT	✓	✓	✓

*not applied to biometric data

Table C.1: Qualitative comparison of related work.

an accuracy loss [216] in addition to requiring quantisation to accommodate for integer-based FHE.

Regarding the aspect of an additional layer of encryption in MT-PRO, a notable recent work is [219], who utilise the concept of password-hardening for fuzzy vaults. While [219] also add a password-derived symmetric key to their scheme, the symmetric decryption is performed on the client side. Thereby, the client gains access to the original protected database entry, i.e., the locked fuzzy vault, and can potentially perform offline attacks. In MT-PRO on the other hand, the symmetric decryption is performed inside the FHE circuit on the server side, such that the client does not learn the protected reference template, while the server does not learn the symmetric key. Table C.1 gives an overview of related works.

C.3 Background

C.3.1 Homomorphic Encryption (HE)

HE is a cryptographic technique that allows for computation on encrypted data that translate directly to computation on the underlying plaintext. HE schemes are classified by the arithmetic operations they allow for, where FHE allows for the evaluation of arbitrary arithmetic circuits [121]. For the scope of our work, we give a simplified definition of the following FHE functionalities [65]:

- $(sk, pk) \leftarrow \mathit{HomKeyGen}(1^\lambda)$: on input of the security parameter λ, generates a secret key sk and public key pk, where pk includes the homomorphic evaluation keys.

- $c_m \leftarrow HomEnc(pk, m)$: on input of the public key pk and a message m, outputs a ciphertext c_m.
- $c_{f(m_1,m_2)} \leftarrow HomEval(pk, c_{m_1}, c_{m_2})$: on input of the public key pk and two ciphertexts c_{m_1} and c_{m_2}, outputs an encryption $c_{f(m_1,m_2)}$ of the evaluation of a function f on the underlying plaintext messages m_1 and m_2.
- $m' \leftarrow HomDec(sk, c_m)$: on input of the secret key sk and ciphertext c_m, outputs a message m'. It holds that $m = m'$ with overwhelming probability.

C.3.2 Homomorphic Transciphering (HT)

HT [71] combines FHE and symmetric encryption. We first define a symmetric cipher with the following functions:

- $k \leftarrow SymKeyGen(1^\lambda)$: on input of the security parameter λ, this function generates a key k .
- $c_m \leftarrow SymEnc(k, m)$: on input of the key k and a message m, this function outputs a ciphertext c_m.
- $m \leftarrow SymDec(k, c_m)$: on input of key k and ciphertext c_m, this function outputs the message m.

Let (sk, pk) be a FHE key pair as defined above. Then, HT allows for the transformation of a symmetric encryption $SymEnc(k, m)$ of a message m to a homomorphic encryption of the same message m, i.e., $HomEnc(pk, m)$, using a homomorphic encryption of the symmetric key, i.e., $HomEnc(pk, k)$. An illustration of the HT functionality can be seen in Fig. C.2.

The transciphering functionality performs a homomorphic evaluation of the decryption circuit of the symmetric cipher. Thereby, the party computing the transciphering does not gain access to the symmetric key k or the message m. Typically, a client device will compute the symmetric encryption of m which requires less computational workload and bandwidth, while a server will compute the transciphering operation and retrieve the homomorphic encryption of m. It is important to note that not all symmetric ciphers are considered *FHE-friendly*, i.e., only symmetric ciphers specifically developed for an application to transciphering can be used [71].

C.4 Protocol

We will now describe our MT-PRO protocol in detail. We begin with a description of the unprotected and protected baseline system using FHE, including a

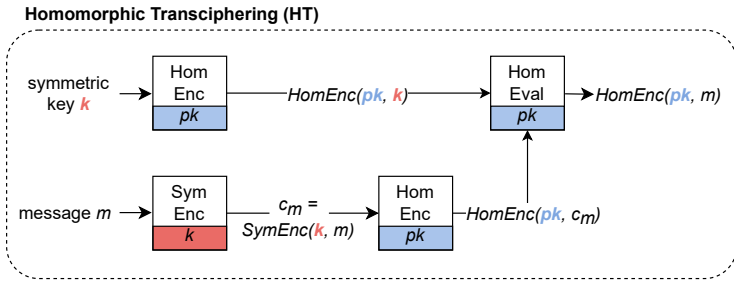

Figure C.2: FHE encryption using Homomorphic Transciphering (HT) [71].

vulnerability analysis under offline attacks. Then, we will describe the integration of HT and discuss its benefits and drawbacks.

C.4.1 Pre-processing

Biometric characteristics can be captured by various sensors depending on the biometric modality. In our protocol, we consider combinations of multiple biometric modalities, known as *multibiometrics*. We consider only feature vectors that can be expressed as fixed-length, ordered vectors. However, our protocol is unconstrained in terms of the length of single vectors, number of vectors, and data type (i.e., binary, integer, or floating point values). In particular, a combination of different feature representations and comparison functions can be used in MT-PRO. After capturing and feature extraction, we consider the reference template or probe feature vector as a concatenation of individually extracted vectors. The cryptographic solution for deriving the combined comparison score will be explained in further detail later in this Section.

C.4.2 Two-Server Architecture

Our MT-PRO protocol builds on the established architecture [31, 103, 126, 270] consisting of a computation server and key server as described above. A client capturing and extracting the reference and probe feature vectors interacts with the computation server in order to initiate an enrolment or verification transaction. In prior works, both servers are considered to act as semi-honest adversaries, i.e., such that they do not deviate from the given protocol, and do not collude in sharing any data they receive or store. We will continue our description of the baseline system under this model before considering the risk of offline attacks and its impact on this security assumption.

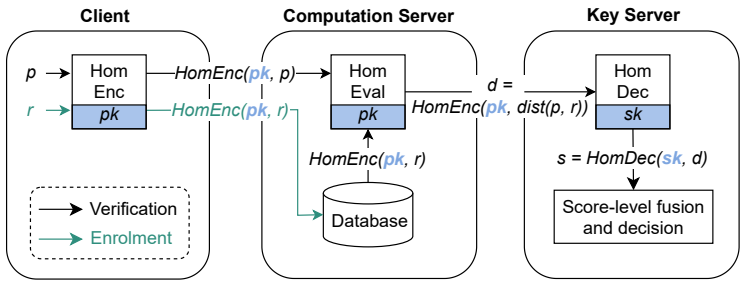

Figure C.3: FHE-protected verification baseline system as used in [31, 103, 126, 270].

C.4.3 Unprotected Baseline System

The unprotected baseline system performs enrolment and verification transactions on plaintext data. During enrolment, the unencrypted reference template is stored in the database. Then, for a verification transaction, a fresh probe feature set is sent to the computation server, who computes the comparison score and determines the verification outcome.

C.4.4 Protected Baseline System

The protected baseline system shown in Fig. C.3 performs the same transactions as the unprotected system, however, while operating on encrypted instead of plaintext data. During enrolment, the client encrypts the reference template to a ciphertext $HomEnc(pk, r)$, which is stored in the database. During verification, the client encrypts the probe feature vector to $HomEnc(pk, p)$, which is sent to the computation server. Through the properties of FHE, the distance score can be computed based on the encrypted reference and probe templates, yielding an encrypted comparison score $d = HomEnc(pk, dist(p, r))$. The key server, using the FHE secret key sk, can decrypt the score to $HomDec(sk, d)$ and determine the verification outcome after threshold comparison.

C.4.5 Vulnerability Analysis

Considering real-world adversaries, the FHE-protected baseline system described in Section C.4.4 established over the past ten years [31, 103, 126, 270] can be vulnerable to the following attack. Having gained access to the protected reference database consisting of ciphertexts $HomEnc(pk, r)$, and the FHE secret key sk, an attacker can easily decrypt and obtain the reference templates, from which sam-

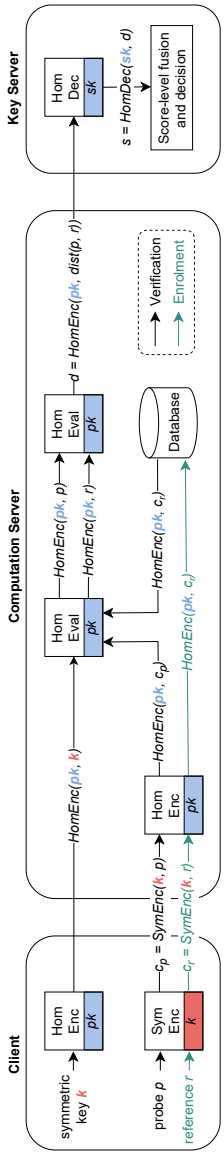

Figure C.4: Proposed MT-PRO protocol based on HT [71] and FHE [65] ensuring protection of the encrypted database under offline decryption attacks. If an attacker gains access to the database and the FHE secret key, it cannot decrypt the encrypted references due to the additional layer of symmetric encryption.

ples can be reconstructed with high confidence [59, 119]. If a template is compromised, its biometric instance (e.g., a finger or eye) can no longer be used for trustworthy authentication due to the risk of impersonations attacks, which could be viable for several decades [166]. We call this attack scenario an *offline decryption attack* or *offline attack*, as the attack can be executed without active access to the system and thereby in an unobtrusive manner.

It is important to note that the addition of zero-knowledge proofs in previous works such as [28] does not withstand such offline attacks, and can therefore not be considered a complete solution to the security challenge. While zero-knowledge proofs guarantee that the computations have been calculated correctly, and can therefore aid in the detection of an attacker deviating from the protocol, they do not protect the encrypted database from decryption once an attacker has gained access to the FHE secret key. We will therefore now present our MT-PRO protocol secure against offline attacks.

C.4.6 Enrolment in MT-PRO

In the MT-PRO enrolment phase (Fig. C.4), the client computes a symmetric encryption of the reference template $c_r = SymEnc(k, r)$ instead of a homomorphic encryption as in the baseline system, using the symmetric key k. Then, the client sends c_r to the computation server, who computes an additional layer of encryption around the symmetric ciphertext through encrypting it homomorphically to $HomEnc(pk, c_r)$. This ciphertext is then stored in the reference database.

C.4.7 Verification in MT-PRO

During MT-PRO verification, also shown in Fig. C.4, the client repeats the symmetric encryption for the freshly extracted probe features and computes $SymEnc(k, p)$. In addition, it computes a homomorphic encryption of its symmetric secret key k, yielding $HomEnc(pk, k)$. Then, both ciphertexts are sent to the computation server, who executes the HT. Upon receiving $SymEnc(k, p)$, the computation server computes a homomorphic encryption $HomEnc(pk, c_p)$. Then, $HomEnc(pk, c_p)$, $HomEnc(pk, c_r)$ and $HomEnc(pk, k)$ are inputs to the HT circuit as described in [71]. Using $HomEnc(pk, k)$, the homomorphic evaluation of the symmetric decryption function is computed. As outputs, the computation server obtains FHE ciphertexts $HomEnc(pk, p)$ and $HomEnc(pk, r)$, and the comparison score is computed, which will be described in the following. The key server decrypts the comparison score and determines the verification outcome.

C.4.8 Multibiometric Comparisons in MT-PRO

In the MT-PRO protocol, combinations of multiple biometric modalities can be used. For this, we extend the concept of coefficient packing presented by [31], where multiple templates are concatenated and encrypted into the same ciphertext. Two challenges arise with regard to multibiometrics: different template lengths and different comparison functions. Through sharing the template order and length (but no information about the underlying data), the computation server can execute the respective comparison functions for each subcomponent of the multibiometric template. To ensure that no information is overwritten, the masking technique from [91] is applied, where only the final comparison score at the start position within the multibiometric plaintext vector is revealed. The individual scores are then combined though an average score level fusion.

C.4.9 Key Management in MT-PRO

Regarding the management of the additional symmetric secret key k within MT-PRO, several options arise:

C.4.9.1 Device key

The symmetric key k can be embedded into the client device, as is typical in IoT applications. This approach has the advantage that the data subject does not need to manage any key material. As the key is static, the reference database can be encrypted as described above, and the transciphering will be correct upon verification. However, the risk of key leakage is larger as one key is used for potentially many subjects, and all reference database entries corresponding to the device key must be re-encrypted when the key is updated. Additionally, a subject can only be verified from the same device that was used during enrolment.

C.4.9.2 Static User Key

Alternatively, the secret key can be made user-specific. To ease key management on the user side, a password-derived key can be used. As long as this key is static, i.e., derived from the password in a deterministic manner, the protocol can be executed as in the case of a device key. Upon a key update, only the corresponding entry for one subject needs to be re-encrypted in the database. Note that contrary to classical password authentication, no hashed password is stored at the computation server, further improving the protection against offline attacks.

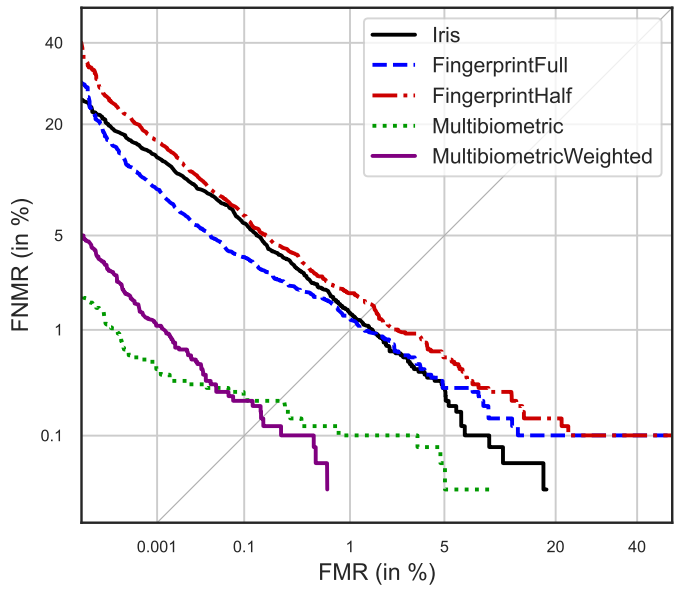

Figure C.5: DET curve showing the multibiometric system performance, where FMR is the false-match rate and FNMR is the false non-match rate [152].

C.4.9.3 Ephemeral User Key

The symmetric key can also be derived from password-authenticated key exchange, corresponding to session keys that are different for each authentication attempt. This approach yields a higher security level for the symmetric key as it is no longer feasible to brute-force. As a significant drawback however, the reference database cannot be encrypted with a symmetric cipher, as the keys used during enrolment and verification will be different. This is useful in classical HT scenarios where large amounts of data are encrypted and HT is mainly used for workload reduction on the client side, but not applicable to prevent offline attacks.

C.5 Experimental Evaluation

We implemented our MT-PRO protocol using the framework by [71], which is based on the Lattigo [194] FHE library. The cryptographic components are the stream cipher HERA [71] and the CKKS [65] FHE scheme. All parameters are chosen at a security level of 128bits. For reproducibility of our results, our implementation is available at `https://github.com/dasec/MT-PRO`. Due to the

high RAM requirements of the HT framework, a Debian GNU 11 server with an AMD EPYC processor at 32x2.8GHz CPU and 128GB RAM was used.

To illustrate the multibiometric verification, we used the newly available deep fingerprint embeddings by [225] of the MCYT [199] database, and deep iris embeddings by [202] of the CASIA Iris Thousand database [70]. For both modalities, 512-dimensional real vectors are extracted, where the fingerprint feature vectors can be split into two 256-dimensional vectors representing the textural and minutiae-derived information, respectively. Cosine distance is used as the comparison function. We evaluated the individual performance (shown as *Iris* and *FingerprintFull* in Fig. C.5), the performance using the first 256 dimensions of the fingerprint features (*FingerprintHalf* in Fig. C.5), the combined performance of both 512-dimensional feature sets (*Multibiometric* in Fig. C.5) as well as the multibiometric performance where only the first 256 dimensions of the fingerprint features are used. For the latter, the individual modalities' scores were scaled according to their dimension (*MultibiometricWeighted* in Fig. C.5). By using different-length feature representations, we show the functionality of MT-PRO described in Section C.4.8 compared to previous approaches considering only feature representations of the same length [31]. MT-PRO can also be instantiated with binary feature representations using the Hamming distance for comparison.

C.5.1 Results

The biometric performance of our MT-PRO protocol can be observed in Fig. C.5, where the weighted multibiometric system is the preferred approach. We note however that our protocol is independent of the multibiometric combinations, and that the individual system performance will depend on the modalities and feature representations used. Due to the use of floating-point based FHE [65], the biometric performance of the unprotected baseline system is maintained.

The computational performance can be viewed in Table C.2. It can be seen that the transciphering operation, i.e., transferring the symmetrically encrypted probe and reference to their homomorphically encrypted representation, is the most expensive operation at 107.64 seconds, followed by the FHE operations at 66.40 seconds. This shows that while the concept of HT is meaningful on a theoretical basis, it is not yet applicable in real-world systems. Further improvements on the cryptographic components are required to improve these transactions times, as further dimensionality reduction of the biometric templates would not yield a significant improvement. Due to larger parameter choices required for HT, the baseline cost of FHE comparisons is also higher than in previous works [31].

MT-PRO Component	Time (s)
Symmetric template encryption	0.42
Homomorphic template encryption	0.21
User key encryption	3.40
Template transciphering	107.64
FHE comparisons	66.40
Comparison score decryption	0.17
Enrolment	0.63
Verification	330.22
Protected Baseline Verification (without HT)	66.78

Table C.2: MT-PRO Execution times for verification and enrolment.

C.5.2 Security Analysis

MT-PRO fulfils the ISO/IEC 24745 [149] requirements of unlinkability and renewability due to security of the FHE scheme against chosen-plaintext attacks [65]. Post-quantum secure irreversibility is provided by the Ring-Learning With Errors [183] hardness assumption of the FHE and HT schemes [71].

C.5.2.1 Security Against Offline Decryption Attacks

We reconsider the adversary from Section C.4.5 that has gained access to the encrypted database and the FHE secret key. In MT-PRO, the adversary only has access to a database with entries $HomEnc(pk, c_r) = HomEnc(pk, SymEnc(k, r))$. Therefore, FHE decryption only yields $SymEnc(k, r)$, which cannot be decrypt without the key k. This security guarantee assumes that the database can be attacked in storage, while the computation server is not corrupted during verification. If the adversary gains access to *HomEnc(pk,k)* during verification, the database could be decrypted. However, an attack on the database is the more realistic attack scenario from a forensic standpoint, as databases are static and outsourced targets.

C.5.2.2 Security Under Full Disclosure Model

The ISO/IEC 30136 [148] standard on performance testing of biometric template protection schemes defines the *full disclosure* attack model for biometric systems, where an adversary has access to all algorithms and all secrets used in the system. The standard adds that this security assumption can be restricted to the adversary knowing a subset of the secret information handled throughout the

system. Thereby, the security of MT-PRO against offline decryption attack can be considered as a partial fulfilment of the full disclosure model, as MT-PRO remains secure if the FHE secret key is leaked to an attacker. Notably, MT-PRO achieves this security without accuracy loss of the biometric comparisons, which is an advantage compared to previous work [201]. However, the symmetric key k as well as its homomorphic encryption $HomEnc(pk, k)$ must be kept secret. As k can be freshly derived from a user-password for each authentication attempt as described in Section C.4.9, it is not easily accessible to an attacker. Additional protection of k could be achieved through the use of multi-party computation, however, at the cost of an additional computational overhead.

C.6 Conclusion

In this work, we presented the MT-PRO protocol for fully homomorphic encryption-based biometric template protection secure against offline decryption attacks even if an attacker gains access to the secret key of the homomorphic encryption scheme. To achieve this, we applied homomorphic transciphering to template protection for the first time, yielding a system with post-quantum security and unimpaired biometric performance. Our experimental evaluation showed that homomorphic transciphering is not yet feasible. Therefore, further improvement of the cryptographic components is required.

Paper D

Type^2: A Secure and Seamless Biometric Two-Factor Authentication Protocol Using Keystroke Dynamics

Pia Bauspieß, Patrick Bours, Christian Rathgeb, and Christoph Busch

Published at Norwegian Information Security Conference (NISK), 2023

Abstract

Password-based user authentication comes with impersonation risks due to poor quality passwords or security breaches of service providers. An additional layer of security can be provided to the authentication through keystroke dynamics, i.e., measuring and comparing users' typing rhythm for their password. While this two-factor authentication is efficient and unobtrusive, the privacy of the biometric characteristics must be ensured. Therefore, we present the Type^2 protocol for secure two-factor authentication based on keystroke dynamics, where the anomaly detection of the latter is executed in the encrypted domain. In an experimental evaluation, we show that our proposed protocol achieves real-time efficiency with an overhead of less than 130 milliseconds compared to password-only authentication.

D.1 Introduction

Reliable user authentication is an important building block in an increasingly digital world [108]. In many authentication scenarios, it is important to ensure that data is disclosed only to the intended receiver, and not to a third party using the receiver's device with their stolen authentication credentials. This applies, e.g., to the disclosure of medical data, but also the agreement of legal contracts or financial transactions.

One of the most common digital authentication methods, passwords, do not inherently provide this security. Trust in password-authenticated communication can be impaired by the fact that many users choose simple passwords that are easy to brute-force [237], or their password may have been compromised by a large-scale attack on a service provider [173].

Biometrics can make such impersonation attacks harder and provide additional confidence in the authentication. In particular, one efficient and unobtrusive way of adding a second trust factor to password-based authentication are keystroke dynamics, i.e., measuring and comparing the users' typing rhythm for their password [167]. Thereby, a second authentication factor can be derived from the already provided password through extracting the timing information from the user's typing rhythm. This motivation is visualized in Figure D.1.

However, such biometric characteristics are classified as sensitive by the European Union's General Data Protection Regulation (GDPR) [109] and must be protected according to the ISO/IEC 24745 international standard on biometric information protection [149]. The latter defines the following three requirements for secure biometric authentication: *i) unlinkability*, two protected biometric tem-

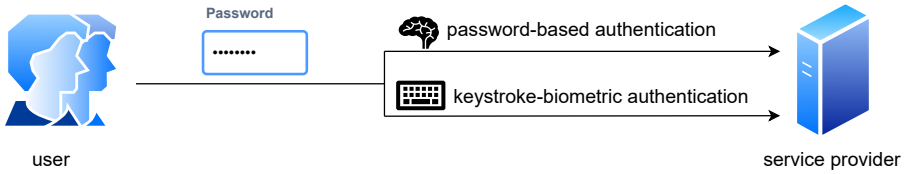

Figure D.1: Seamless integration of biometric authentication using keystroke dynamics.

plates stored in different applications cannot be linked to the same subject, *ii) renewability*, new templates can be created from the same biometric instance without the need to re-enrol, and *iii) irreversibility*, it is impossible to retrieve original templates given only protected templates. In addition, the biometric performance (i.e., accuracy) as well as the computational performance of the unprotected system should be preserved.

In this work, we present the Type^2 protocol for secure two-factor authentication based on keystroke dynamics, where the biometric comparisons are executed in the encrypted domain. To this end, *Fully Homomorphic Encryption* (FHE) [121] is applied to the biometric features both during enrolment and verification. More concretely, we investigate the compatibility of established anomaly detectors for keystroke dynamics [167], and present an analysis of the applicability and feasibility of FHE to these detectors. Further, we give a comprehensive security analysis of Type^2 with regard to adaptations that have to be made in order to apply FHE to different detectors. We evaluated our Type^2 protocol experimentally on publicly available data [167] and libraries [227]. Our proposed protocol can be instantiated with detectors that achieve real-time user authentication at an overhead of less than 130 milliseconds per authentication attempt.

The rest of this article is structured as follows: Section D.2 discusses works that are closely related to ours, before Section D.3 gives more technical background information. Our protocol and main contribution is presented in Section D.4, together with its experimental evaluation given in Section D.5. Finally, we draw our conclusions in Section D.6.

D.2 Related Work

One of the first to discuss the application of homomorphic encryption to keystroke dynamics were Šeděnka et al. [235]. In their work, the authors indicate that their key generation protocol could also be instantiated with FHE, but that they re-

	Encryption Scheme	Accuracy Preservation	Performance Preservation	Post-quantum Security
Šeděnka et al. 2014 [235]	DGK	✓	✗	✗
Acar et al. 2019 [6]	BFV	✗	✓	✓
Loya et al. 2021 [182]	CKKS	✓	✗	✓
Ours	CKKS	✓	✓	✓

Table D.1: Qualitative comparison of related works on keystroke dynamic authentication with (fully) homomorphic encryption.

frained from this choice due to the significant computational overhead of FHE, in particular with respect to the schemes and implementations that were available in 2014. Therefore, they use additively homomorphic encryption only [79], which only allows for additions of ciphertexts, and therefore limits the complexity of detectors. In their evaluation, they use an in-house dataset that does not allow for reproducibility of their research. Nevertheless, we can estimate a comparison of the efficiency, as the authors of [235] achieve execution times in the magnitude of minutes, whereas our Type^2 protocol can be executed in the order of milliseconds.

More recently, Acar et al. [6] presented a privacy-preserving multi-factor authentication system named *PINTA*, where they consider keystroke dynamics as one potential authentication factor. The authors evaluate their protocol on the established and publicly available keystroke dynamics dataset provided by [167], in addition to other modalities such as mouse movements. Their multi-factor authentication protocol uses fuzzy hashing in combination with FHE, which impairs the accuracy of the system. Furthermore, the FHE scheme used by [6] is the BFV [50, 112] encryption scheme, which operates on integers and therefore requires a quantisation of keystroke dynamic features. The computational cost of their authentication decision was evaluated at around 370 milliseconds.

The most closely related work to ours was presented by Loya et al. [182] in 2021. In their work, the authors evaluate a neural network with differential privacy during the training process, while the keystroke dynamic features are protected using the CKKS [65] encryption scheme. This is the same FHE scheme we will use for our experimental evaluation. In addition, the work by [182] utilizes the same established dataset for keystroke dynamic evaluation provided by [167]. However, the execution times of [182] are not applicable for real-time applications, as they are no lower than 14 seconds.

D.3 Background

D.3.1 Password-Authenticated Key Exchange

For the first component of our Type^2 protocol, *Password-Authenticated Key Exchange* (PAKE) [155] is used. Compared to traditional hashing and salting of passwords, PAKE provides additional security against offline attacks and can be considered the state-of-the-art in password authentication. Popular approaches include the SRP protocol [264] used among others in the Apple iCloud, or the more recent the OPAQUE [155] protocol. Similar to biometric authentication, a PAKE protocol is defined through a registration phase, where the user's password information is enrolled into the system in a protected manner, and an authentication phase, where a cryptographic key is exchanged successfully if and only if the correct password is provided again. The PAKE component in our protocol can be easily exchanged and we therefore do not focus on it further for the scope of this work, but refer the reader to the works of [264] and [155] directly.

D.3.2 Fully Homomorphic Encryption

FHE allows for the evaluation of arithmetic circuits on encrypted data [121] and has been determined to fulfil the ISO/IEC 24745 [149] requirements for biometric information protection [31, 44, 126, 270]. For the scope of our work, we define an FHE scheme through the following algorithms:

- $(sk, pk) \leftarrow KeyGen(1^\lambda)$: on input of the security parameter λ, generates a secret key sk and public key pk, where pk includes the homomorphic evaluation keys.
- $c_m \leftarrow HomEnc(pk, m)$: on input of pk and a message m, outputs a ciphertext c_m.
- $c_{f(m_1,m_2)} \leftarrow HomEval(pk, f, c_{m_1}, c_{m_2})$: on input of pk, a public function f, and two ciphertexts c_{m_1} and c_{m_2}, outputs an encryption $c_{f(m_1,m_2)}$ of the evaluation of f on the underlying plaintext messages m_1 and m_2.
- $m' \leftarrow HomDec(sk, c_m)$: on input of sk and ciphertext c_m, outputs a message m'.

These operations can be applied to vectorized data, where all evaluations will be performed element-wise, yielding an improvement in terms of computational overhead [44]. It holds that $Dec(sk, HomEval(pk, f, c_{m_1}, c_{m_1})) = f(m_1, m_2)$ [65].

D.3.3 Keystroke Dynamics

In this work, we focus on keystroke dynamic features that can be extracted from password timings measured using the same keyboard for each authentication attempt. For a given password, this feature set will always be of fixed length n, and the order of typed letters will be the same, easing the task of anomaly detection. Different features that can be measured from password typings are [167]: *(i) keydown-keydown time:* time interval between a key is pressed and the consecutive key is pressed, *(ii) keyup-keydown time:* time interval between a key is released and the consecutive key is pressed, and *(iii) hold time:* time interval between a key is pressed and the same key is released.

Using these timings, the typical typing pattern of a user is established during the enrolment or training phase. In this step, the mean vector over a set of timing vectors is stored, with additional information such as the covariance of the features. For neural network-based approaches, this step corresponds to the training of the weights. For a verification transaction, a fresh probe timing vector is captured from the data subject. The probe features are compared against the stored reference template and a distance score or *anomaly score* [167] is computed. Using a predefined threshold, the anomaly score can be used to grant or deny the subject access to the system. The combined algorithms of enrolment and verification are referred to as an *anomaly detector* in the following.

D.4 Proposed System

In this Section, we describe the Type^2 protocol with FHE protection and necessary modifications and limitations for all of the anomaly detectors described in [167]. An overview of our proposed system is given in Figure D.2.

In the enrolment phase, both the password and biometric reference of a subject are enrolled into the system. For the password w, the PAKE registration is performed according to the chosen approach [264, 155]. Additionally, an FHE key pair (sk, pk) is generated by the key server, and the public key pk is shared with the other parties. We assume that an attacker has access to the public key. The client uses pk to encrypt the keystroke timing features after the reference vector r has been established in the training process. The *Computation Server* (CS) stores $c_r \leftarrow HomEnc(pk, r)$.

In the first step of the verification protocol, the subject provides a password w', which is input to the PAKE protocol. If the PAKE authentication phase is successful, the system proceeds to the keystroke anomaly detection. For an optimized user experience, both processes can also be run in parallel. Using the

timing features t' extracted from w', the client computes the probe ciphertext $c_p \leftarrow HomEnc(pk, t')$ and sends it to the CS. Using the encrypted reference template c_r corresponding to the biometric claim, the CS computes the detector $d(c_r, c_p)$, and sends d to the key server. Here, d can be decrypted and the threshold comparison of the decrypted anomaly score against threshold τ is computed. The system outputs a bit $b = 1$ if the anomaly score is smaller than τ, and $b = 0$ otherwise.

D.4.1 Adversary Model

In our work, we consider all parties to operate in the *semi-honest adversary model*. In this model, the participating parties do not deviate from the given protocol, but may aim to collect information that is available to them. It can be argued that a more realistic model is given through the *malicious adversary model*, where parties are allowed to deviate from the given protocol to gain further information. This model has been discussed in the context of biometric template protection [28], where zero-knowledge proofs are applied for the protection against malicious adversaries. Our proposed Type^2 protocol is compatible with such proofs, however, we do not focus further on malicious adversaries in our work. We assume that the capture process of the timing features takes place in a controlled environment during enrolment, resulting in trusted reference vectors. During verification, the system may be confronted with presentation attacks. However, this work focuses on the application of FHE to keystroke dynamic features in a manner that does not alter the unencrypted accuracy of the system.

D.4.2 Euclidean Detector

The squared Euclidean distance used in the Euclidean detector [99] has been studied for FHE-based template protection for other biometric modalities, mostly for face [31, 44]. As the square-root operation is not supported by FHE, the squared Euclidean distance is preferred to the original Euclidean distance. During the enrolment phase, [167] describe that the mean vector over the set of training vectors is computed and stored as reference vector. As the enrolment is considered an offline process, the mean vector is computed on the unencrypted training vectors. Then, the client enrolls a subject by encrypting the mean reference vector r as $c_r \leftarrow HomEnc(pk, r)$, and sends c_R to CS.

For a verification transaction, the client encrypts the probe feature vector p as $c_p \leftarrow HomEnc(pk, p)$, which is sent to CS. Here, CS computes the squared Eu-

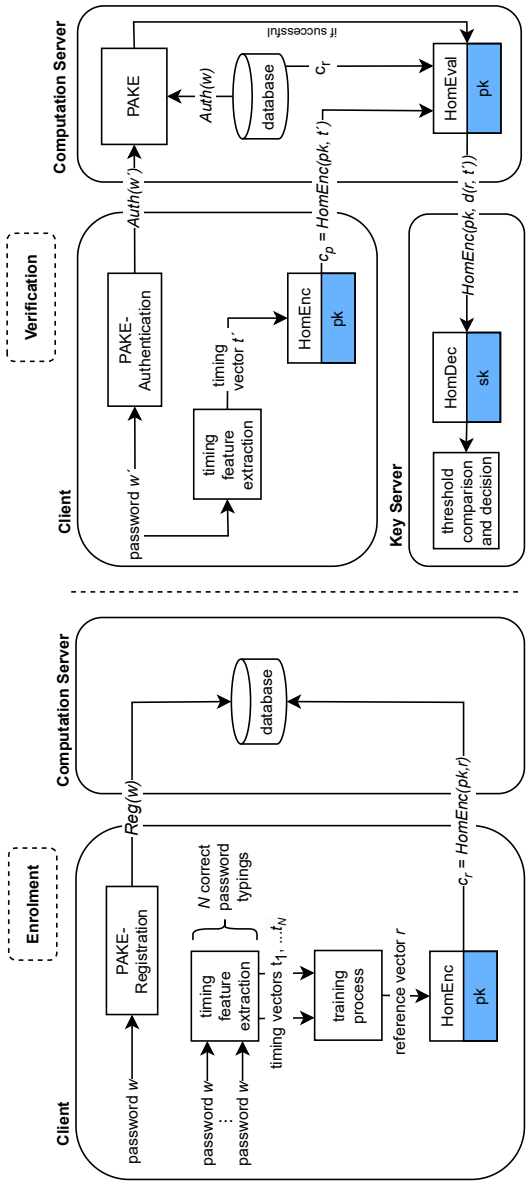

Figure D.2: Enrolment (left) and verification (right) transactions in the Type^2 protocol.

clidean distance

$$d_{Euclid}(r,p) = \sum_{i=0}^{n-1} (r_i - p_i)^2 \tag{D.1}$$

as established in the recent literature [44]: the two ciphertexts c_r and c_p are subtracted, yielding an element-wise subtraction of their elements. The resulting vector is multiplied with itself, corresponding the square of elements in the vector. To facilitate the computation of the sum over the packed vector, the established rotate-and-add technique is applied [44]. The total cost of FHE operations required for the Euclidean detector is summarized in Table D.2.

D.4.3 Normed Euclidean Detector

The normed Euclidean detector expands upon the Euclidean detector through normalizing the final anomaly score, i.e., dividing it by the multiplied norm of the probe and reference feature vectors [42]. During the enrolment phase, the norm of the reference template is computed and encrypted to an additional ciphertext $c_r^{||\cdot||} \leftarrow HomEnc(pk, ||r||)$, which is sent to CS together with the encrypted reference template c_r, and both ciphertexts are stored at CS.

During verification, the client computes $c_p^{||\cdot||} \leftarrow HomEnc(pk, ||p||)$ in addition to c_p. The computation of the Euclidean distance follows the description in Section D.4.2. In addition to the Euclidean distance, one homomorphic multiplication $c_p^{||\cdot||} \cdot c_r^{||\cdot||}$ is performed. Both the encrypted squared Euclidean distance and the encrypted multiplied norms are sent to the key server for decryption, which calculates the final anomaly score. Ideally, the division would also be computed in the encrypted domain. However, division is not directly supported by FHE operations. In the following, we note the described approach as *approach A* and describe a second option yielding a more private computation (*approach B*).

For a fully private computation of the normed Euclidean distance *(approach B)*, the client can compute the inverted Euclidean norms $1/||r||$ and $1/||p||$ from the reference and probe feature vectors during enrolment and verification, respectively. Then, it can produce ciphertexts $c_r^{1/||\cdot||} \leftarrow HomEnc(pk, 1/||r||)$ and $c_p^{1/||\cdot||} \leftarrow HomEnc(pk, 1/||p||)$. The computation on the FHE ciphertexts described above then corresponds to

$$\frac{d(c_p, c_r)}{c_p^{||\cdot||} \cdot c_r^{||\cdot||}} = d(c_p, c_r) \cdot c_p^{1/||\cdot||} \cdot c_r^{1/||\cdot||}, \tag{D.2}$$

where d denotes the squared Euclidean distance described in Section D.4.2. This more private computation comes at a cost of one additional homomorphic multiplication and thereby an increased multiplicative depth of the circuit. The cost for all homomorphic operations is higher for parameter sets that allow this additional circuit depth. Therefore, approach B must be expected to have higher computational workload than approach A.

D.4.4 Manhattan Detector

The Manhattan detector utilizes the Manhattan distance, which is another established distance metric in pattern recognition [99]. It is defined as

$$d_{Manhattan}(r, p) = \sum_{i=0}^{n-1} |r_i - p_i|. \tag{D.3}$$

On unencrypted data, only a bit shift is required for the computation of the absolute difference between the reference and probe feature vector elements. However, a bit shift is not an available computation in FHE. When computing on integer or float values, the computation of the absolute value corresponds to a conditional statement. The evaluation of conditional statements is by design infeasible on encrypted data, as the result of the computation needs to be known in order to evaluate the statement. Recent approaches have explored conditional statements in FHE, however, they cannot be considered feasible for real-world applications [147].

Therefore, the only encrypted computation that can be performed during verification for the Manhattan detector is the difference between the reference and probe feature vectors, i.e., $c_r - c_p$, and the absolute values and the sum are computed on the plaintext data at the key server. It can be argued that the protected computation of the difference yields an additional protection of the features, in particular during storage, but also during the comparison, as it can be challenging for an attacker to determine the original features based on the difference alone. However, the aforementioned privacy limitations apply. Further privacy protection could be given through a random negation of both probe and reference feature vectors. However, this approach would correspond to an additional multiplication of the encrypted probe feature vector during verification, thereby increasing the authentication workload.

For the filtered Manhattan distance, outliers are excluded during the training phase [159]. As the enrolment phase is computed on plaintext data however, this does not impact the application of FHE to the detector in question.

D.4.5 Scaled Manhattan Detector

The scaled Manhattan distance utilizes mean absolute deviation a_i of the i-th feature of the training vector as a scale factor for the final anomaly score [17]. Similarly to the normed Euclidean distance, this additional vector a can be computed on the plaintext reference vectors during enrolment. Due to the lack of a division operation in FHE, we apply the same transform as in Section D.4.3 and let the client encrypt the inverse $1/a$ into a ciphertext $c_r^{1/a} \leftarrow HomEnc(pk, 1/a)$, which is stored at CS alongside the encrypted reference vector c_r. Then, we can express the anomaly score as

$$d^{scaled}_{Manhattan}(r,p) = \sum_{i=0}^{n-1} \frac{|r_i - p_i|}{a_i} = \sum_{i=0}^{n-1} \left| \frac{r_i - p_i}{a_i} \right| = \sum_{i=0}^{n-1} \left| (r_i - p_i) \cdot \frac{1}{a_i} \right| \tag{D.4}$$

and calculate the values $\frac{r_i - p_i}{a_i}$ in the encrypted domain at the following cost for a verification transaction (see Table D.2): first, one encryption of c_p is computed, then one subtraction of $c_r - c_p$. Subsequently, the inverted mean absolute deviation vector $c_r^{1/a}$ is multiplied to the difference, and the result is decrypted. As in Section D.4.4, the absolute values and computation of the sum must be conducted on plaintext data, as the evaluation of conditional statements such as the absolute value are not feasible on FHE-encrypted data.

For the computation of the scaled Manhattan distance, the mean absolute deviation vector a should be stored in encrypted form at CS. It can be assumed that a encodes sensitive information about the biometric reference stored at CS, and can therefore be considered to be of similar sensitivity as the feature vectors themselves. Scaling on the decryption comparison score in plaintext can therefore not be considered a secure approach.

D.4.6 Mahalanobis Detector

The Mahalanobis detector [99] is based on the Mahalanobis distance:

$$d_{Mahalanobis}(r,p) = (r-p)^\top S^{-1}(r-p), \tag{D.5}$$

where S denotes the covariance matrix over the training vectors. Both S and the mean reference vector r are computed in plaintext. Then, the following ciphertexts are computed by the client: an encryption of the mean reference feature vector c_r, and each column of the inverted covariance matrix S^{-1} into a ciphertext c_i^S, where $\{c_i^S \leftarrow HomEnc(pk, S_i^{-1})\}_{i=0}^{n-1}$. During verification, the client obtains and encrypts a probe feature vector and sends the resulting ciphertext c_p to CS. In the first step of the distance computation, CS computes

$(r-p)^\top S^{-1}$ on the corresponding ciphertexts through one subtraction of $c_r - c_p$, and n multiplications of the resulting vector with each of the ciphertexts c_i^S. The vector-matrix multiplication is completed by computing the sum over each $(c_r - c_p) \cdot c_i^S$, which is computed as described in Section D.4.2. The total cost for the Mahalanobis detector is given in Table D.2. The approach to the normed Mahalanobis detector [42] follows the same procedure as the normed Euclidean detector described in Section D.4.2 as approach B. In addition to the computations for the Mahalanobis distance score, the inverted probe and reference feature vector norms $c_r^{1/||\cdot||} \leftarrow HomEnc(pk, 1/||r||)$ and $c_p^{1/||\cdot||} \leftarrow HomEnc(pk, 1/||p||)$ are encrypted. Then, the final comparison score is obtained after a multiplication by both ciphertexts to the original score, i.e., $d(c_p, c_r) \cdot c_p^{1/||\cdot||} \cdot c_r^{1/||\cdot||}$. The additional encryption (of $c_p^{1/||\cdot||}$) and two multiplications can be observed in Table D.2.

D.4.7 Nearest-neighbor Detector

The nearest-neighbor approach [138] expands the Mahalanobis detector described in Section D.4.6 by computing the Mahalanobis distance to every training vector (instead of the mean reference vector), and choosing the lowest out of these comparison scores as the final outcome. Its cost with regard to FHE operations can therefore be determined as the N-fold effort of the Mahalanobis detector, where N is the number of training vectors. As discussed above, conditional statements cannot be evaluated efficiently in FHE. Therefore, all N distance scores need to be decrypted, and the lowest score is determined in the plaintext domain. The nearest-neighbour approach can therefore not be fully realized in FHE, and furthermore has an infeasible overhead in terms of the number of required FHE operations.

D.4.8 Neural-Network Detector

The neural network detector utilizes a simple fully connected neural network with one hidden layer. The enrolment phase corresponds to the training phase of the network, while the comparison score is achieved through inference over one probe sample [99]. This inference can be expressed as two matrix multiplications with the encrypted probe feature vector, and can therefore be computed similarly to the Mahalanobis distance. As the network only has one output node, the second multiplication corresponds to a similar vector multiplication as in Section D.4.6. The total cost with regard to the originally proposed parameter choices [167] can be viewed in Table D.2. The FHE protection for the auto-associative neural-network detector introduced by [138] is similar to the

previously described approach with the difference of n output nodes and an additional distance computation. These additional costs can be viewed in Table D.2.

D.4.9 Fuzzy Logic Detector

The fuzzy logic detector [135] applies a succession of logical statements, i.e., conditional statements, to classify the probe feature set instead of classic distance metric. While the reference and probe features can still be sent and stored encrypted, all computations can only be computed in plaintext due to the challenge of evaluating conditional statements on encrypted data. FHE protection can therefore not be meaningfully applied to this detector.

D.4.10 Outlier-Counting Detector

The outlier-counting detector presented by [135] is derived from the scaled Manhattan distance. However, the final score is a count of element-wise scores above a predefined threshold, rather than the distance scores itself. For every feature in the feature vector, a so-called z-score defined as

$$z_i = \frac{|r_i - p_i|}{\sigma_i}, \tag{D.6}$$

where σ_i is the standard deviation of the i-th feature calculated during the training phase. We therefore apply the same transformation as in Sections D.4.3 and D.4.5, and store a ciphertext $c_{1/\sigma} \leftarrow \mathit{HomEnc}(pk, 1/\sigma)$ at CS during enrolment. Here, the vector $1/\sigma$ contains all inverse standard deviations $1/\sigma_i$ for every feature i. During verification, client and CS proceed as in Section D.4.5 and obtain the encrypted result of the computation $c'_z = (c_r - c_p) \cdot c_{1/\sigma}$. As argued above, neither the absolute value nor the threshold comparisons can be computed in the encrypted domain. Therefore, c'_z is decrypted and the remaining computations are executed over the plaintext vector.

D.4.11 One-Class Support Vector Machine Detector

For the one-class *Support Vector Machine* (SVM) detector [273], the training phase is again conducted on the unencrypted training vectors. After training is completed, the determined hyperplane h used as the separator is encrypted into a ciphertext $c_h \leftarrow \mathit{HomEnc}(pk, h)$ and stored at CS. A verification transaction then corresponds to a projection of the encrypted probe feature set p into the higher-dimensional separator space of the SVM, i.e., a matrix multiplication, the cost of which is presented in Table D.2.

D.4.12 k-Means Detector

The application of the established k-means clustering algorithm [184] has been proposed for keystroke dynamics by [162]. In terms of the application of FHE to this detector, the approach corresponds to the Euclidean detector described in Section D.4.2. For each of the k centroids, the Euclidean distance between the centroid and the probe feature vector is computed, and the closest distance is determined to be the final comparison score. However, as the evaluation of this last conditional statement is not feasible within FHE, all three distances are decrypted, and the minimal distance is determined over the plaintext data. This means that final comparison score was fully computed in the encrypted domain, however, the algorithm reveals additional information in plaintext that may impact the privacy of the enrolled subjects, i.e., the discarded distances to the remaining $k - 1$ centroids. This limitation is also indicated in Table D.2 for better transparency with regard to the different approaches.

D.4.13 Workload and Feasibility Discussion

We have now described all keystroke anomaly detectors from the seminal study by [167] and their challenges and adaptions under FHE encryption. Due to the limitations of FHE computations discussed so far, we can classify these detectors into three categories: *(1)* vector-based distance metrics such as the Euclidean and Manhattan distance, *(2)* detectors requiring matrix-vector or matrix multiplications, which introduce a significantly higher workload in FHE operations than the detectors discussed above. These include the (normed) Mahalanobis detector [99] as well as neural network-based approaches, including SVMs, as evaluated in [182]. And finally, *(3)*, detectors require the evaluation of conditional statements, which cannot be realized efficiently in FHE [147]. These include the nearest-neighbour [138], fuzzy logic and outlier counting [135], and k-means [162] detectors. We give the computational workload of all detectors in Table D.2. Further context to Table D.2 is provided through the relative cost of FHE operations given in Table D.3. With regard to their feasibility however, detectors from categories *(2)* and *(3)* are not evaluated them experimentally. The experimental workload for some detectors of category *(3)* however can be estimated based on the Euclidean and Manhattan distance. E.g, the workload of outlier counting can be estimated as the workload of the scaled Manhattan distance, while the workload of the k-means detector corresponds to the k-fold workload of the Euclidean detector.

Detector	Enc	EvalAdd	EvalSub	EvalMult	EvalAtIndex	Dec
Euclidean	1	$n-1$	1	1	$n-1$	1
Euclidean (normed) (appr. A)*	2	$n-1$	1	2	$n-1$	2
Euclidean (normed) (appr. B)	2	$n-1$	1	3	$n-1$	1
Manhattan**	1	—	1	—	—	1
Manhattan (filtered)**	1	—	1	—	—	1
Manhattan (scaled)**	1	—	1	1	—	1
Mahalanobis	1	$2n(n-1)$	—	n^2	$2n(n-1)$	1
Mahalanobis (normed)	2	$2n(n-1)$	—	n^2+2	$2n(n-1)$	1
Nearest-neighbour*	N	$2Nn(n-1)$	—	N^2n	$Nn(n-1)$	N
Neural-network (standard)	1	$\lceil \frac{2n}{3} \rceil n-1$	—	$\lceil \frac{2n}{3} \rceil^2$	$2n(n-1)$	1
Neural-network (auto-assoc)	1	$2(n^2-n)$	1	n^2+n+1	$n-1(2n+1)$	1
Outlier-counting**	1	—	1	1	—	1
SVM (one-class)	1	$n+m-2$	m	$m \cdot m$	$(n-1)$	1
k-means*	1	$k(n-1)$	k	k	$k(n-1)$	k

Table D.2: FHE operations during verification for keystroke anomaly detectors [167], where n is the feature dimension, N is the number of training vectors, k is the number of centroids in the k-means clustering, and m is the dimension of the SVM projection space. Detectors marked with ** can only be partly computed on encrypted data, while detectors marked with * reveal more information than the final comparison score.

Operation on encrypted data	Add	Subtract	Rotate	Decrypt	Multiply	Encrypt
Relative cost	1	5	24	33	46	52

Table D.3: Relative cost of CKKS [65] operations implemented in PALISADE [31, 227].

D.5 Experimental Evaluation

We implemented our Type^2 protocol using the CKKS [65] scheme implemented in the PALISADE [227] C++ FHE library at a security level of 128bits for all variants of the Euclidean and Manhattan detectors. All execution times were measured on an Intel i7 CPU @ 2.60GHz with 32GB RAM and an Ubuntu 20.04 operating system. As a dataset, we used the established CMU keystroke dynamics dataset provided by [167] and maintain all features and the split into training and testing data. For the 400 timing vectors captured from each of the 51 subjects in the dataset, the first 200 password timings were used for the training of each detector, and samples from the remaining timings for verification.

The execution times for enrolment and verification for the five discussed detectors are given in Table D.4, where N is the number of subjects to be enrolled in the system. As discussed in Section D.4, the Manhattan detectors have the fastest execution times as they use the lowest number of homomorphic operations. However, they cannot be considered fully secure, as the pre-computation step is decrypted before anomaly score can be calculated. The Euclidean detectors grant more privacy, with the plain Euclidean and the normed Euclidean (approach B) being the only fully private detectors with regard to evaluation under FHE. For the latter, the impact of the increased multiplicative depth of 2 instead of 1 can be observed. The encryption of reference or probe data, which consists of two encryption operations for the feature vector and its norm (or inverted norm) for both approach A and B to the normed Euclidean detector, therefore increases to 21 milliseconds instead of 8 milliseconds due the parameter set required to accommodate the increased circuit depth.

In terms of the biometric performance, we refer the reader to the original evaluation conducted in [167], which we give in Table D.5. Through the application of the CKKS [65] with correct parameter choices, the biometric performance if not altered in the encrypted domain. In particular, we chose a scaling factor of 50 bits for the CKKS scheme, such the accuracy of the detectors is not affected by the application of the encryption scheme. Therefore, the accuracy evaluations given by [167] are maintained.

D.5.1 Security Analysis

Our proposed Type^2 protocol fulfils the ISO/IEC 24745 [149] requirements unlinkability, renewability, and irreversibility. Firstly, irreversibility is given through the hardness of the Ring-Learning with Errors (R-LWE) problem [183], which the CKKS [65] FHE scheme builds upon. As R-LWE is believed to be secure against attacks implemented on a quantum computer [10], our Type^2 pro-

Detector	Enrolment (ms)	Verification (ms)
Euclidean	$4N$	117
Euclidean (normed) (appr. A)*	$8N$	125
Euclidean (normed) (appr. B)	$21N$	338
Manhattan**	$4N$	4
Manhattan (filtered)**	$4N$	4
Manhattan (scaled)**	$8N$	8

Table D.4: Experimentally determined execution times in milliseconds for the evaluated detectors. Detectors marked with ** can only be partly computed on encrypted data, while detectors marked with * are computed on encrypted data, but reveal more information than the final comparison score.

Detector	Equal-Error Rate (EER)	Standard Deviation
Euclidean	0.171	0.095
Euclidean (normed) (appr. A)*	0.215	0.119
Euclidean (normed) (appr. B)		
Manhattan**	0.153	0.092
Manhattan (filtered)**	0.136	0.083
Manhattan (scaled)**	0.096	0.069

Table D.5: Biometric performance for the evaluated detectors taken from [167]. Detectors marked with ** can only be partly computed on encrypted data, while detectors marked with * are computed on encrypted data, but reveal more information than the final comparison score.

tocol inherits this post-quantum security. Secondly, unlinkability and renewability are provided through the IND-CPA security of the CKKS scheme, i.e., its indistinguishability under chosen-plaintext attacks. Thereby, an attacker cannot distinguish between two encryptions of the same feature vector and two encryptions of different feature vectors. Finally, our protocol preserves both the biometric and computational performance of the unprotected authentication as shown in Section D.5. The choice of the PAKE, which is an independent component of the protocol next to the FHE protection, determines the security of the authentication as a second factor. However, post-quantum protection may not be necessary for the PAKE component, as the user password does not require long-term protection as sensitive biometric features do. This yields more flexibility with regard to the chosen PAKE approach, where computational efficiency lower than the workload for the biometric authentication should be considered [155, 264].

D.6 Conclusion

In this work, we have presented the Type^2 protocol for secure two-factor authentication based on keystroke dynamics as second trust factor, where the protection of sensitive biometric data is ensured through fully homomorphic encryption. For five established keystroke anomaly detectors, we showed the potential and limitations of their evaluation under fully homomorphic encryption. In an experimental evaluation, we show that our protocol outperforms the state-of-the-art with execution times of under 130 millisecond per authentication attempt. While the assumption of the semi-honest adversary model remains a limitation, the cryptographic principles applied in this work can be used to extend the Type^2 protocol in more realistic adversary models. With advances of the cryptographic components, more complex detectors, e.g., neural networks, could be investigated in future research. Furthermore, it would be interesting to extend the Type^2 protocol to other behavioral features using mobile phones as the capture device.

Paper E

BRAKE: Biometric Resilient Authenticated Key Exchange

Pia Bauspieß, Tjerand Silde, Matej Poljuha, Alexandre Tullot, Anamaria Costache, Christian Rathgeb, Jascha Kolberg, and Christoph Busch

Accepted for publication in IEEE Access, 2024

Abstract

Biometric data are uniquely suited for connecting individuals to their digital identities. Deriving cryptographic key exchange from successful biometric authentication therefore gives an additional layer of trust compared to password-authenticated key exchange. However, biometric data are sensitive personal data that need to be protected on a long-term basis. Furthermore, efficient feature extraction and comparison components resulting in high intra-subject tolerance and inter-subject distinguishability, documented with good biometric performance, need to be applied in order to prevent zero-effort impersonation attacks.

In this work, we present a novel protocol for *Biometric Resilient Authenticated Key Exchange* that fulfils the above requirements of biometric information protection compliant with the international ISO/IEC 24745 standard. In our protocol, we present a novel modification of unlinkable fuzzy vault schemes that allows their connection with oblivious pseudo-random functions to achieve resilient protection against offline attacks crucial for the protection of biometric data. Our protocol is independent of the biometric modality and can be implemented based on the security of discrete logarithms as well as lattices. We provide an open-source implementation of both instantiations of our protocol which achieve real-time efficiency with transaction times of less than one second from the image capture to the completed key exchange.

E.1 Introduction

Biometric characteristics provide accurate and non-repudiable identification of individuals over several decades [166]. This makes them suited for bridging the gap between real and digital identities in a way passwords or other machine-generated identifiers cannot. At the same time however, these properties also make them uniquely vulnerable. In particular, biometric information cannot be revoked or replaced in the same way a password or cryptographic token can. Once a digital representation of a biometric characteristic, further referred to as a biometric template, has been leaked, the underlying source (e.g., a particular finger or eye), can no longer be used securely for authentication. In fact, biometric templates provide no protection of the underlying data, as they can be reversed to samples sufficient for attacks [59, 119, 185].

Due to this risk, biometric data have been recognised as sensitive personal data by the European Union's General Data Protection Regulation (GDPR) [109] and the ISO/IEC 24745 international standard on biometric information protection [149]. The latter defines three security requirements for secure biometric

systems: *i) unlinkability and renewability*, meaning that an attacker cannot connect two protected biometric templates stored in different applications, and new templates from the same source look indistinguishable to a previously stored reference, *ii) irreversibility*, it should be impossible for an attacker to retrieve original samples given only protected templates, and *iii) performance preservation*, the computational performance and the recognition accuracy of the system should not be impacted significantly by adding a layer of protection to the original data.

At first sight, the performance preservation requirement in ISO/IEC 24745 seems to be a question of convenience only. However, it details a second and crucial dimension that determines the security of biometric authentication: the accuracy of the underlying biometric comparison function. Contrary to passwords, which can be compared in an exact manner, captured samples of the same biometric characteristic are never exactly equal, but *fuzzy*. They are subject to noise such as ageing, environmental influence, or image quality. Comparison of two samples is therefore based on some measure of similarity. If this measure is too imprecise, or the feature representation is not discriminative enough, an authentication system is not capable of accurately distinguishing between mated comparisons, where the samples stem from the same subject, and non-mated authentication attempts, where the samples stem from different subjects. Trust in the derived authentication would consequently be low.

Recently, the idea of building authenticated key exchange on the basis of biometrics has gained interest with the proposal of Biometrics-Authenticated Key Exchange (BAKE) [260]. Analogously to Password-Authenticated Key Exchange (PAKE) [155], a client and server negotiate a shared cryptographic key that should be equal if and only if the biometric authentication was successful.

With their protocol, the authors of [260] achieve security in terms of the protection of the biometric data with classical security assumptions. However, their biometric comparator is vulnerable, as we show by reproducing their results experimentally. The reason for this imprecision is a fingerprint comparison algorithm that is specific to their protocol, but has not been evaluated in terms of biometric performance (i.e., accuracy). We provide this evaluation and show that the algorithm is not able to distinguish between mated comparison trials within the same identity and non-mated comparison trials between different identities in a sufficient manner (see Appendix E.6). More generic protocols both on symmetric fuzzy PAKE (fPAKE) [100] and asymmetric fuzzy PAKE (fuzzy aPAKE) [106] have been proposed. However, with regard to biometrics, they have the following shortcomings: fPAKE [100] does not achieve protection of the biometric data, which is shared with the server in plaintext. Fuzzy aPAKE [106] achieves security in both dimensions in theory, but is inefficient in practice as

it is based on generic oblivious transfer which is performed once for each bit in the biometric template. In addition, [100] and [106] only enable comparison of fixed-length biometric representations. The most accurate comparison metric for fingerprints, one of the most popular biometric modalities, is however based on variable-length representations, the similarity of which cannot be expressed as a simple distance function.

E.1.1 Contribution

In this work, we present a protocol for *Biometric Resilient Authenticated Key Exchange* (BRAKE) that addresses the deficiencies of previous works [100, 106, 260]. Our BRAKE protocol achieves effective protection of the biometric data against offline attacks through the application of an *Oblivious Pseudo-Random Function* (OPRF). Our protocol is efficient with execution times of under one second on commodity hardware from the biometric capture to the completed key exchange, including communication cost. To the best of our knowledge, our protocol is the first to achieve secure biometric authenticated key exchange with high biometric and computational performance, thus fulfilling ISO/IEC 24745. More precisely, we contribute:

- Biometric resilient authenticated key exchange secure against offline attacks: through a novel modification of unlinkable fuzzy vault schemes, we build a seamless integration of biometric authentication into oblivious pseudo-random functions to achieve resilient protection against offline attack, which is crucial for the long-term protection of biometric data according to the ISO/IEC 24745 [149] standard.
- Classical and post-quantum security: Our two-round protocol can be instantiated both with a discrete logarithm OPRF [155] and Diffie-Hellman key exchange [87] as well as a lattice-based OPRF [12] and the state-of-the art post-quantum key encapsulation mechanism CRYSTALS Kyber [49], which was recently standardized in NIST IR 8413 [10]. Through our protocol's compatibility with lattice-based primitives, which are assumed to be post-quantum secure, we further achieve long-term protection of the underlying biometric data.
- Interchangeability of biometric modalities: our protocol can be instantiated with different fuzzy vault schemes that have been designed for different biometric modalities and feature representations. In particular, it is compatible with both fixed-length and variable-length representations of biometric characteristics.
- Open-source implementation: an implementation of our protocol based

on discrete logarithms as well as lattices is available at `https://github.com/dasec/DL-BRAKE` and `https://github.com/dasec/PQ-BRAKE`, respectively. We show that our protocol achieves real-time efficiency with transaction times of under one second from the fingerprint image capture at the sensor to the completed key exchange. To support the reproducibility of our results, we provide automated installation scripts with all dependencies alongside our implementation.

E.1.2 Related Work

We briefly discuss the state-of-the-art to motivate two principles for secure biometrics-authenticated key exchange: recognition accuracy and reciprocal interaction.

The main concern with the protocol proposed in [260] is the generation of the biometric secret key constructed from fingerprint representations. The authors use a simplified version of the well-studied nearest-neighbour approach first proposed by [156], which they chose due to its anticipated rotation invariance. However, this algorithm and its flaws have been studied for two decades, specifically, its inability to tolerate missing genuine minutiae [58]. It has therefore been found unusable in practice, and improved rotation-invariant fingerprint recognition algorithms have been proposed that mitigate the known shortcomings [58]. Such improved algorithms require a more complex comparison subsystem however, and are not compatible with the constructor offered in [260]. Notably, the authors of [260] fail to state the recognition accuracy of their iris and fingerprint based protocols, and do not give an experimental evaluation detailing the security with regard to the biometric performance.

Their construction for iris is based on the established fixed-length feature representation IrisCode [80] and can be assumed to achieve adequate accuracy as long as the sample quality is high. It is worth noting that the state-of-the-art in iris recognition is based on samples captured under near-infrared light, and therefore requires designated capture devices, i.e., near-infrared sensors. Such specific sensors are however not part of most personal communications devices such as smartphones. The use of classical iris recognition in the Signal [115] protocol as motivated by [260] is therefore not meaningful. In such a scenario, iris recognition in the visual spectrum would need to be considered, which is a more challenging task and provides, as of today, lower accuracy [211].

Secondly, the public keys derived from the biometric secret keys in [260] are vulnerable to offline attacks: in their construction, any adversary can guess a biometric template and check if it corresponds to the public key in hand, without interacting with another party. In such an attack, the adversary does not

Scheme	Year	Feature representation	Cryptographic primitives	Asymmetric	Efficient	Accurate	Post-quantum security	Compliant with ISO/IEC 24745
fPAKE [100]	2018	binary, fixed-length	GC + ECC	✗	✓	✓	✗	✗
fuzzy aPAKE [106]	2020	binary, fixed-length	ECC + OT	✓	✗	✓	✗	✗
BAKE [260]	2021	integer, variable-length	ECC + LWE	✓	✓	✗	✓	✗
iPAKE [48]	2023	binary, fixed-length	ECC + PAKE	✗	✓	✓	✗	✗
ttPAKE [139]	2023	binary, fixed-length	Secret sharing + OT	✓	✓	✓	✗	✗
BAKA [275]	2023	binary, fixed-length	ECC + Blockchain	✓	✓	✓	✗	✗
BRAKE (*ours*)	2023	integer, variable-length	ECC + PAKE	✓	✓	✓	✓	✓

Table E.1: Comparison of our protocol to related work.

have to guess an exact biometric feature representation, but succeeds as soon as she finds an input that is close enough with regard to the distance metric used. This probability can be expressed as the false-match rate of the biometric system, i.e., the proportion of authentication attempts from non-mated samples falsely accepted as authentication attempts of an enrolled data subject. Again, low biometric accuracy leads to a low effort in an offline search attack.

Even with assumed high biometric accuracy, offline attacks expose biometric data to high risks. Therefore, we construct our protocol such that interaction is required for every adversarial guess, which allows for rate-limiting that can be enforced as long as at least one party remains honest. The concept of enforcing interaction through a third party OPRF service in itself is not new [111]. However, the construction previously presented by [111] is neither trivially compatible with fuzzy secrets such as biometric features, nor with lattice-based primitives as our proposed protocol. In particular, no lattice-based partially OPRF as required for the protocol given in [111] is known as of today, and its construction lies outside of the scope of this work.

An overview of how our proposed scheme compares to related works can be found in Table E.1. An efficient solution to fuzzy PAKE was presented by [100]. However, the solution is constructed as a symmetric protocol, where the server learns the biometric reference template. The approach of [100] does therefore not fulfil the ISO/IEC 24745 [149] requirements. Building on this line of research, [48] recently proposed fuzzy PAKE based on Error-Correcting Codes (ECC). While their protocol is efficient with a small overhead compared to [100] and improves upon the security of [100], the symmetric construction remains an obstacle with regard to ISO/IEC 24745 [149].

A different line of research emerged with the fuzzy asymmetric PAKE construction of [106]. Here, the asymmetric protocol does not allow the server to learn the biometric reference template. However, the expensive computation of bitwise Oblivious Transfer (OT) makes the solution impractical for real-world applications. More recently, [139] proposed their solution ttPAKE to typo-tolerance PAKE, which can be considered related to the challenges posed by biometric authentication with regard to the fuzziness of input data. Their solution builds on the idea of [106], but is based on double-layered secret sharing. While their protocol is asymmetric, the password is shared with the server in the setup phase for the purpose of constructing a secret-shared password table, and is deleted by the semi-honest server afterwards. If this protocol were applied to biometric data, this plaintext disclosure of the authentication secret would violate the ISO/IEC 24745 [149] requirements. Another recent work presents BAKA [275], a protocol for biometric authentication and key agreements based on fuzzy extractors. However, this work applies blockchain to store biomet-

ric data, which is an inherent violation of the ISO/IEC 24745 [149] renewability requirement. Through the immutability of blockchain records, compromised reference templates cannot be renewed. Furthermore, none of the above works apart from [260] have been instantiated using post-quantum secure cryptographic primitives.

Further recent works are concerned with authentication based on fuzzy input data, however, with different aims to our work. Motivated by more private solutions for TLS authentication, [176] proposed single message Credential Hiding Login (CHL). Their one-round protocol allows for efficient user authentication both for static and fuzzy secrets, with biometric authentication as a possible application. Their scheme is based on the security of Learning with Errors (LWE) problems and can be instantiated with post-quantum secure parameters. In contrast to our work however, not session keys are exchanged as a result from the successful login. Another solution to biometric authentication based on functional encryption was recently presented by [105]. While their solution is computationally efficient, no key material is generated from the successful biometric two-factor authentication. Similarly, [265] presented post-quantum secure biometric authentication using searchable encryption, a cryptographic technique related to functional encryption as applied in [105].

Other related works have been directed on extracting uniformly distributed cryptographic keys directly from biometric templates without running an interactive protocol [89]. Similar to [100] and [106], only fixed-length representations are considered that can be compared with some distance metric. From fuzzy extractors, two-factor authentication protocols have been built [212]. More recently, [230] proposed a session key generation protocol specifically for fingerprint based on so-called cancelable biometrics, which are one-way transforms on the biometric data that are not based on well-studied cryptographic problems and can therefore not be assumed to underlie specific hardness assumptions.

E.1.3 Structure of Paper

The rest of this paper is structured as follows: In Section E.2, background information and definitions required for the construction of our protocol are presented. As our main contribution, Section E.3 presents our BRAKE protocol with security definitions and proof sketches, before we give concrete instantiations based on discrete logarithms and lattices in Section E.4. Section E.5 presents the experimental evaluation of the protocol and practical comparison with related work, before we outline our conclusions in Section E.6.

E.2 Preliminaries

The framework for automated and interoperable biometric recognition has been standardised in ISO/IEC 19794-1 [150], and subsequent parts of the standard define biometric data interchange formats for the modalities fingerprint, face, iris, voice, handwritten signatures, and vascular biometrics. For the scope of our work, we look at the three most prevalent modalities fingerprint, face, and iris, for which well-tested fuzzy vault schemes exist.

E.2.1 Biometric Performance Metrics

Biometric performance testing and reporting is standardised in ISO/IEC 19795-1 [152] and subsequent parts. The evaluation of biometric systems is based on two components: error rates and throughput rates. For a verification scenario, the most important error metrics are:

- *False Non-Match Rate (FNMR):* proportion of mated comparisons that resulted in a reject decision.
- *False Match Rate (FMR):* proportion of non-mated comparisons that resulted in an accept decision.

The FMR can be thought of as the security level of the biometric system, detailing how many zero-effort impostors were able to be verified. In most scenarios, systems with a FMR below 1% are considered secure, while high-security applications such as automated border control require a FMR lower than 0.1% [117]. The FNMR on the other hand can be considered as the convenience level of the system, detailing how many mated comparison trials were not able to be verified. A FNMR up to 5% is considered acceptable [117].

Factors impacting the recognition performance of a biometric system are first and foremost the sample quality both during enrolment and verification, and the robustness of the feature representation and comparison algorithm with regard to rotation, translation, and noise of the samples [198, 248]. Furthermore, any feature transformation such as binarisation may impact the accuracy of the system.

E.2.2 Entropy of Biometric Representations

The entropy of biometric data is a topic that is often referred to in works about fuzzy cryptographic primitives [100]. In the literature, the entropy of a face has been determined at 56 bits [7], a minutiae-based fingerprint representation at 82 bits [215], and an iris at 249 bits [81]. However, these numbers can only be

considered as an upper bound of the entropy of a certain biometric instance, as the amount of information in a biometric sample heavily depends on the capture device used and its fidelity (e.g., its resolution) as well as the feature extraction algorithm used. Indeed, [251] argues that it is not in all scenarios appropriate to use the entropy of a single biometric template as a measure for security, which is an overestimate when it comes to comparisons between biometric features. Here, the false-accept security defined as $\log_2(FMR^{-1})$ gives a more accurate measure, as it is sufficient for an attacker to guess a template that is close enough to a reference template.

E.2.3 Fuzzy Vault

The concept of fuzzy vaults was first introduced by [160], who propose a scheme that allows to *lock* a biometric feature secret set t with a secret polynomial f using a biometric feature secret set t using a probabilistic algorithm. The output of this algorithm is a locked fuzzy vault that can be *unlocked* using a second biometric feature set t', if there are enough points the intersection of t and t'. We give a short definition of their original scheme before we move on to the state-of-the-art for different biometric modalities.

Definition E.1 (Fuzzy Vault Scheme [160])**.** Let $\mathcal{C}$ be an error-correcting code, $H : \mathcal{C} \to \{0,1\}^{2\lambda}$, for security parameter λ, be a cryptographic hash function H, and let τ a biometric comparison threshold. Then, a *fuzzy vault scheme* is a set of the following algorithms:

- $(f, H(f), V) \leftarrow \texttt{lock}(t)$: On input of a biometric feature set t, the algorithm samples a random secret $f \in \mathcal{C}$ and outputs a locked fuzzy vault V together with the hash digest $H(f)$.
- $f' \leftarrow \texttt{unlock}(V, H(f), t')$: On input of a locked fuzzy vault V and a biometric feature set t', the algorithm outputs an opening polynomial $f' \in \mathcal{C}$. The unlocking can be verified by comparing $H(f)$ to $H(f')$.

A basic authentication protocol based on the fuzzy vault scheme is given in Figure E.1.

Instantiation for Fingerprint

The original schemes by [160] and a similar scheme by [73] have been proven to be insecure due their construction based on large point clouds to hide the secret f, which are vulnerable to correlation attacks [250]. Therefore, [251] presented an improved scheme to mitigate correlation attacks (see [251], Section

Enrolment

Setup party		**Server**
t reference template		
$(f, H(f), V) \leftarrow \texttt{lock}(t)$		
	$\xrightarrow{H(f), V}$	
		store $(H(f), V)$

Verification

Client		**Server**
t' probe template		$(H(f), V)$
	$\xrightarrow{\text{request } H(f), V}$	
	$\xleftarrow{H(f), V}$	
$f' \leftarrow \texttt{unlock}(V, H(f), t')$		
return $H(f) = H(f')$		

Figure E.1: Fuzzy vault authentication protocol based on [160].

1.2.3), building on the initial proposal by [89]. These improved fuzzy vault schemes fulfil the requirements of ISO/IEC 24745 [149].

The improved fuzzy vault scheme has first been constructed for minutiae-based fingerprint representations [251]. From the pattern of fingerprint ridge lines, significant points known as *minutiae* are extracted as compact and distinguishing features, specifically, ridge endings and bifurcations, namely the location and orientation where one ridge line splits into two. In the scheme by [251], minutiae are encoded into a finite field $\mathbb{F}_{p'}$ using absolute pre-alignment and quantisation to account for a certain degree of noise with regard to the position of the minutiae. The set of minutiae $t \subset \mathbb{F}_{p'}$ is then considered the biometric template. A polynomial $f \in \mathbb{F}_{p'}[x]$ of degree $\tau - 1$ is chosen uniformly at random and locked as

$$\texttt{lock}(t) = (f, f(x) + \prod_{a \in t}(x - a)) =: (f, V). \tag{E.1}$$

To unlock the vault, V is evaluated on the probe minutiae set t' and decoded using a Reed-Solomon decoder, yielding

$$\texttt{unlock}(V, t') = \text{decode}(\{(b, V(b)) \mid b \in t'\}) =: f'. \tag{E.2}$$

Lemma E.1 (Theorem 1 in [251])**.** Let $(f, H(f), V) \leftarrow \texttt{lock}(t)$ be a commitment to a polynomial $f \in \mathbb{F}_{p'}[x]$ with minutiae set t, and $f' \leftarrow \texttt{unlock}(V, H(f), t')$ an unlocking of V using a minutiae set t'. Then, $f = f'$ if and only if $|t \cap t'| \geq \tau$.

Analogue constructions exist for iris [220] and face [218] recognition, which we refer the reader to for full details.

E.2.4 Cryptographic Primitives

Definition E.2 (Pseudo-Random Function, [60])**.** A family of functions $f_k : \{0,1\}^\lambda \times \{0,1\}^n \to \{0,1\}^{n'}$, with key $k \in \{0,1\}^\lambda$, are called Pseudo-Random Functions (PRFs) if the following holds:

- $f_k(x)$ is efficiently computable from k and x.
- It is not efficiently decidable whether one has access to a computation oracle for $f_k(\cdot)$ or to an oracle producing uniformly random bit-strings of length n.

Definition E.3 (Oblivious Pseudo-Random Function, [116])**.** A two-party protocol π between a client and a server is an Oblivious Pseudo-Random Function (OPRF) if there exists some PRF family f_k, such that π privately realizes the following functionality:

- Client has input x; Server has input k.
- Client outputs $f_k(x)$; Server outputs nothing.

Definition E.4 (Hashed Diffie-Hellman OPRF, [114])**.** Let $\mathbb{G}$ be a cyclic group of prime order p, $x \in \{0,1\}^*$ the client input, $k \in \mathbb{Z}_p$ the evaluator's secret key, $H_{\mathbb{G}} : \{0,1\}^* \to \mathbb{G}$ and $H_{\mathbb{Z}_p} : \{0,1\}^* \to \mathbb{Z}_p$ cryptographic hash functions that output values in $\mathbb{G}$ and $\mathbb{Z}_p$, respectively. The protocol `HashDH` consists of the following algorithms:

- $(B, r) \leftarrow \texttt{blind}(x)$: The client samples a random $r \leftarrow\!\!\$\ \mathbb{Z}_p$ and outputs r and $B \leftarrow [r]H_{\mathbb{G}}(x)$.
- $S \leftarrow \texttt{eval}(B, k)$: On input $B \in \mathbb{G}$, the evaluator outputs $S \leftarrow [k]B$.
- $U \leftarrow \texttt{unblind}(S, r)$: On input $S \in \mathbb{G}$ and $r \in \mathbb{Z}_p$, the client outputs $U \leftarrow H_{\mathbb{Z}_p}(x, [r^{-1}]S)$.

As a result of this protocol, the client privately obtains $H_{\mathbb{Z}_p}(x, [k]H_{\mathbb{G}}(x))$ without learning k and without the evaluator learning the input x nor the output U.

Definition E.5 (Key Encapsulation Mechanism, [197])**.** A Key Encapsulation Mechanism (KEM) is a scheme with three algorithms `KeyGen`, `encap` and `decap`, where

- $(\texttt{pk}, \texttt{sk}) \leftarrow \texttt{KeyGen}(1^\lambda)$: takes as input the security parameter λ and outputs a public key `pk` and a secret key `sk`.
- $(\texttt{ctx}, \gamma) \leftarrow \texttt{encap}(\texttt{pk})$: takes as input a public key `pk`, samples a session pre-key γ, and outputs γ and an encapsulation `ctx` of γ under the public key `pk`.
- $\gamma' \leftarrow \texttt{decap}(\texttt{ctx}, \texttt{sk})$: takes as input an encapsulated session pre-key `ctx` and a secret key `sk` and outputs a decapsulated session pre-key γ'.

We require that for all $(\texttt{pk}, \texttt{sk})$ generated from `KeyGen` we have, except with negligible probability, that $\gamma = \texttt{decap}(\texttt{encap}(\gamma, \texttt{pk}), \texttt{sk})$, and that the scheme is IND-CCA secure.

E.2.5 Lattice-Based Cryptography

Lattice-based cryptography builds upon certain lattice problems which are considered hard to solve even for quantum computers, and these can be used as the basis for designing a variety of cryptographic systems [204]. The two most popular lattice problems are the *Learning With Errors* (LWE) decision-problem introduced in [222] and the *Short Integer Solution* (SIS) search-problem introduced in [8]. In this work, we use the module variants of these problems, where we are working over cyclotomic rings $R_q = \mathbb{Z}_q[X]/\langle X^N + 1\rangle$ where N is a power of two and q a prime. The norm of elements in R_q is computed on coefficient vectors of polynomials in $\mathbb{Z}$.

Definition E.6 (Module-LWE)**.** . Let χ be a bounded distribution over R_q^d and let $s \leftarrow \chi$ be a secret vector. Then, sample $A_i \in R_q^{d\times d}$ uniformly at random and $e_i \leftarrow \chi$, and finally set $(A_i, b_i = A_i \cdot s + e_i)$ in $R_q^{d\times d} \times R_q^d$. The M-LWE$_{d,s,\chi}$ decision-problem is to decide with non-negligible advantage whether m independent samples $\{(A_i, b_i)\}_{i=1}^m$ are computed as above or sampled from the uniform distribution over $R_q^{d\times d} \times R_q^d$.

Definition E.7 (Module-SIS)**.** . Given m uniform vectors $a_i \in R_q^d$, the M-SIS$_{d,m,\beta}$ problem is to find polynomials $s_i \in R_q$ such that all $||s_i|| \leq \beta$ and

$$\sum_{i=1}^{m} a_i \cdot s_i = 0 \in R_q.$$

E.3 Biometric Resilient Authenticated Key Exchange

In this Section, we introduce our protocol for Biometric Resilient Authenticated Key Exchange (BRAKE) built from a fuzzy vault scheme, an OPRF, and a KEM.

E.3.1 Setting

For our proposed protocol, we assume that a biometric capture device is linked to a client which performs the preprocessing and feature extraction, and acts as a communicating party in the protocol. Its communication counterparts are a server which controls a database of locked fuzzy vaults and client reference public keys, and an evaluator which is in possession of a secret OPRF key. In practice, the evaluator can be instantiated by a trusted execution environment at the server. For this reason, we do not model direct communication between the

client and the evaluator, but work under the weaker assumption that all communication between client and evaluator is seen by the server. This is a common practice in biometric information protection [270], as it allows for enhanced network security choices that protect the party handling secret key material. Furthermore, we assume that authenticated channels are established between all parties, e.g., through TLS. Thereby, mutual authentication can be established between a client and the server.

E.3.2 Modification of Fuzzy Vault Schemes

In the original improved fuzzy vault schemes, the decoding algorithm with highest performance both in terms of execution times and accuracy is the Guruswami-Sudan decoder [133]. Thereby, unlocking a fuzzy vault with feature vector t' corresponds to a randomised brute-force decoding strategy, where subsets of t' are chosen uniformly at random and evaluated as unlocking sets for the reference fuzzy vault. During this randomised decoding, a candidate polynomial f' is generated for each subset and compared against the stored hash $H(f)$ corresponding to the biometric reference template t. When a candidate polynomial is found for which $H(f) = H(f')$, the decoding attempts are stopped. If no candidate polynomial is found within a certain number of decoding attempts, the underlying comparison of t and t' is classified as a non-mated comparison trial.

In our protocol however, we do not wish to store $H(f)$ at the server as it allows for offline brute-force attacks. Instead, we run the full decoding attempts until the threshold for non-mated comparison trials is reached, even when we expect a mated comparison trial. During decoding, we temporarily store all candidate polynomials and sort them with respect to their frequency. For a mated comparison, we expect the correct candidate polynomial f' for which $H(f') = H(f)$ to appear as the most frequently reconstructed polynomial due to the large overlap of the sets t and t'. A similar strategy is applied in [73] and is supported by our experimental evaluation, showing only a negligible deviation with regard to the biometric performance.

Notably, the FMR and thereby security of the system is not affected by the change to highest-frequency decoding. In both cases, no non-mated comparisons yield matching candidate polynomials within the list decoder threshold. Therefore, the polynomial that occurs with the highest frequency is also not a matching candidate polynomial. Consequently, the FMR is not affected by the change from hash-verified decoding to highest-frequency decoding.

In addition, the frequency pattern found in a mated comparison does not give an attacker an advantage in terms of an offline-brute force attack. Through the

additional roots of the randomly generated secret polynomial f, a number of seemingly correct polynomials of degree $\tau - 1$ could be interpolated by an attacker that is not in possession of a mated feature set. Therefore, a brute-force attack on a locked vault alone, without the confirmation of $H(f)$ or a successful key exchange, corresponds to a non-mated comparison attempt with no clear frequency pattern.

E.3.3 Protocol

In this Section, we give the formal definition of our proposed protocol for biometric resilient authenticated key exchange.

Definition E.8 (Biometric Resilient Authenticated Key Exchange)**.** A three-party protocol BRAKE between a client, a server and an evaluator is a Biometric Resilient Authenticated Key Exchange, if it realizes the following functionalities:

- Enrolment: A trusted setup party inputs a biometric reference template t and corresponding identifier id. The setup party computes a locked vault (f, V) based on t. The evaluator inputs a key k. Then the parties jointly compute a client public key $\mathtt{cpk}_t$ derived from f. The server outputs $(V, \mathtt{cpk}_t = \mathtt{eval}(f, k), \mathtt{id})$ and the other parties outputs nothing. The enrolment protocol is detailed in Figure E.2.
- Verification: The client inputs a biometric probe feature set t' and a biometric claim id, the server inputs $(V, \mathtt{cpk}_t, \mathtt{id})$ and the evaluator inputs k. The client requests the locked vault V for id and interpolates a polynomial f' from t'. The parties jointly compute a key exchange on input f'. The server outputs a session key ρ and the client outputs a session key ρ' and a bit indicating if $H(\rho) = H(\rho')$. The verification is detailed in Figure E.3.

Here, the client will output the bit 1 if and only if $|t \cap t'| \geq \tau$ for τ the biometric verification threshold. For the algorithms defined in Definition E.8, we require the following building blocks:

Definition E.9 (Building blocks)**.** We define the following building blocks for the BAKE protocol:

- $\mathtt{pp} \leftarrow \mathtt{setup}(1^\lambda)$: The setup algorithm defines a universe $\mathcal{P}$, randomness space $\mathcal{R}$, key space $\mathcal{K}$ and a cryptographic hash function $H : \{0,1\}^* \rightarrow \{0,1\}^{2\lambda}$. Further, the setup algorithm defines an error-correcting code $\mathcal{C}$ with correction capacity τ. These are incorporated in the public parameters pp and all following algorithms implicitly inherit pp.

- $(f, V) \leftarrow \texttt{lock}(t)$: The algorithm takes as input a biometric template t, samples a random polynomial $f \in \mathcal{C}$, and outputs f and a locked fuzzy vault V. Note that the fuzzy vault scheme do not include the hash digest $H(f)$.
- $f' \leftarrow \texttt{unlock}(V, t')$: The algorithm takes as input a biometric probe feature vector t' and locked fuzzy vault V, and outputs an opening polynomial f'.
- $(B, r) \leftarrow \texttt{blind}(f)$: The algorithm samples a random element $r \in \mathcal{R}$ and outputs an element $B \in \mathcal{P}$.
- $S \leftarrow \texttt{eval}(B, k)$: On input $B \in \mathcal{P}$ and key $k \in \mathcal{K}$, the server outputs an evaluation $S \in \mathcal{P}$.
- $\texttt{sk} \leftarrow \texttt{unblind}(S, r)$: On input $S \in \mathcal{P}$ and $r \in \mathcal{R}$, the algorithm outputs an evaluation t U that can further be used as (or to generate) a client secret key $\texttt{csk} \in \mathcal{K}$.
- $(\texttt{sk}, \texttt{pk}) \leftarrow \texttt{KeyGen}(1^\lambda)$: The algorithm outputs a secret key $\texttt{sk} \in \mathcal{K}$ and a public key $\texttt{pk} \in \mathcal{P}$.
- $\texttt{pk} \leftarrow \texttt{pkGen}(\texttt{sk})$: The algorithm takes as input a secret key $\texttt{sk} \in \mathcal{K}$ and outputs a public key $\texttt{pk} \in \mathcal{P}$.
- $(\texttt{ctx}, \gamma) \leftarrow \texttt{encap}(\texttt{cpk})$: The algorithm takes as input a client public key cpk, samples a session pre-key γ and outputs γ and an encapsulation ctx of γ under cpk.
- $\gamma' \leftarrow \texttt{decap}(\texttt{ctx}, \texttt{csk})$: The algorithm takes as input an encapsulated session pre-key ctx and a client secret key csk and outputs a decapsulated session pre-key γ'.
- $\rho \leftarrow \texttt{KDF}(\texttt{cpk}, \texttt{spk}, \texttt{cpk}_e, \texttt{spk}_e, \gamma)$: The key derivation function KDF takes as input the client and server static and ephemeral public keys cpk, spk, $\texttt{cpk}_e$, $\texttt{spk}_e$ as well as a pre-key γ and outputs a session key $\rho \in \{0,1\}^{2\lambda}$.

The detailed functioning of the BRAKE protocol can be seen in Figures E.2 and E.3. We also give a short semantic description in the following. During enrolment (Figure E.2), a client public key $\texttt{cpk}_t$ is derived from a biometric reference template t and the OPRF key k, and is stored at the server together with a locked fuzzy vault V of t using a secret random polynomial f. First, the client generates f and locks the vault with template t. Note that now, the fuzzy vault scheme no longer includes the hash digest $H(f)$ of the secret polynomial sampled during locking. Then, the client initiates the OPRF evaluation on input f. The evaluator evaluates the blinded input B using the OPRF key k, and the client is able to

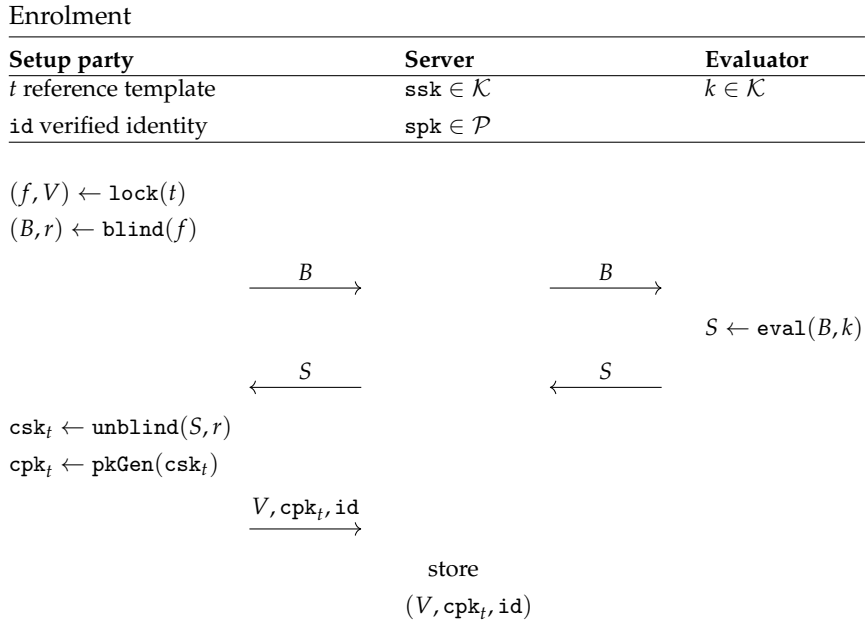

Figure E.2: BRAKE enrolment protocol.

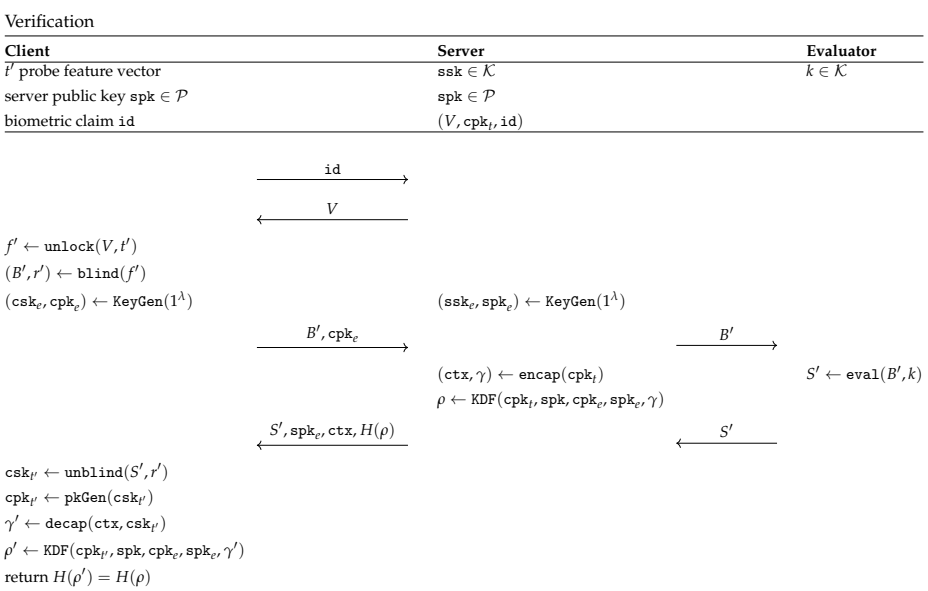

Figure E.3: BRAKE verification protocol.

unblind and obtain its secret key $\mathtt{csk}_t$, from which it computes the corresponding public key $\mathtt{cpk}_t$. To conclude the enrolment step, the client sends the tuple $(V, \mathtt{cpk}_t, \mathtt{id})$ to the server to be stored for future reference.

For verification and key exchange (Figure E.3), the client requests the fuzzy vault V stored at the server for identity id, and, using a biometric probe t', unlocks the vault to a polynomial f'. Then, the OPRF evaluation on f is computed analogously to the enrolment step. At the same time, the client and server generate ephemeral key pairs to prepare the key exchange. Additionally, the server has a static key pair $(\mathtt{ssk}, \mathtt{spk})$ generated during setup that is not derived from any biometric information. For the key exchange, we assume that the client has access to the static server public key spk as discussed above. Once all keys have been generated, the server encapsulates a session pre-key γ using the client's public key $\mathtt{cpk}_t$. The client can decapsulate γ if and only if the secret reconstructed from the fuzzy vault was correct, i.e., in the case where t and t' are closer than threshold τ. Finally, the session key ρ is derived from γ using the client and server static and ephemeral public keys $\mathtt{cpk}$, $\mathtt{spk}$, $\mathtt{cpk}_e$, $\mathtt{spk}_e$ in the key derivation function KDF. We note that the hashed session key ρ allows for the authentication to be explicit.

E.3.4 Security Definitions

Following the definition of the BRAKE protocol in Figures E.2 and E.3, we give formal definitions of the security of the protocol. For simplicity, we implicitly model the use of identifiers within the enrolment database. In theory, an adversary wants to learn a biometric feature vector that is close to any enrolled template. In practice however, it always needs to choose a specific identity to attack or run attacks on multiple specific identities in parallel. The following definitions and proof sketches model security in the case where a template t is enrolled in the database held by the server, and an honest client would use a feature vector t' to authenticate.

Notation. Denote by $f^{-1} = \log_2(FMR^{-1})$ the false-accept security of a biometric feature extractor and comparator, let ℓ be the rate limit enforced by the server and the evaluator, and let $\ell_{\mathcal{A}}$ be the brute-force capacity of the attacker $\mathcal{A}$.

Definition E.10. (Correctness) We say that a BRAKE protocol is correct if a capture subject presenting a biometric probe feature vector t' and identifier id can successfully authenticate to an honest server if and only if $|t \cap t'| \geq \tau$ for a fixed biometric verification threshold τ, except with negligible probability.

Definition E.11. (Client Privacy) We say that a BRAKE protocol has client privacy if an adversary $\mathcal{A}$ controlling the client has the following advantage in obtaining a biometric feature vector t' that is close to an enrolled biometric template t:

$$\Pr\left[\begin{array}{l}\mathtt{dist}(t,t') < \tau \ : \\ \end{array} \begin{array}{c} \mathtt{pp} \leftarrow \mathtt{setup}(1^\lambda) \\ \{V, \mathtt{cpk}_t\} \leftarrow \mathtt{enroll}(\mathtt{pp}, t) \\ \forall i \in [\ell] : \begin{cases} (B', \mathtt{cpk}_e) \leftarrow \mathcal{A}(\mathtt{pp}, V) \\ (\mathtt{ssk}_e, \mathtt{spk}_e) \leftarrow \mathtt{KeyGen}(1^\lambda) \\ S' \leftarrow \mathtt{eval}(B', k) \end{cases} \\ t' \leftarrow \mathcal{A}(S', \mathtt{spk}, \mathtt{spk}_e, \mathtt{ctx}) \end{array}\right] \leq \ell f^{-1} + \mathsf{negl}(\lambda).$$

Definition E.12. (Server Privacy) We say that a BRAKE protocol has server privacy if an adversary $\mathcal{A}$ controlling the computation server has the following advantage in obtaining a biometric feature vector t' that is close to an enrolled biometric template t:

$$\Pr\left[\begin{array}{l}\mathtt{dist}(t,t') < \tau \ : \\ \end{array} \begin{array}{c} \mathtt{pp} \leftarrow \mathtt{setup}(1^\lambda) \\ \{V, \mathtt{cpk}_t\} \leftarrow \mathtt{enroll}(\mathtt{pp}, t) \\ \forall i \in [\ell] : \begin{cases} B' \leftarrow \mathcal{A}(\mathtt{pp}, \{V, \mathtt{cpk}_t\}) \\ S' \leftarrow \mathtt{eval}(B', k) \end{cases} \\ t' \leftarrow \mathcal{A}(S') \end{array}\right] \leq \ell f^{-1} + \mathsf{negl}(\lambda).$$

If client and server run the protocol BRAKE honestly, the evaluator only sees the blinded element, which is information-theoretically secure, and hence, in-

dependent of the biometric template. We therefore do not model evaluator privacy.

The advantage of an adversary controlling both the client and the server effectively reduces to server privacy. In this scenario, the information the adversary needs to guess is the evaluated element S'. However, as discussed above, the evaluator cannot distinguish between evaluation requests for different biometric feature vectors corresponding to mated authentication attempts, or repeated evaluation requests for a single identity aimed at running a brute-force search. Therefore, rate-limiting at the evaluator can be enforced by user-specific OPRF keys. This way, the evaluator will learn the identifier of the user attempting to authenticate, but is not able to gain any more knowledge about her biometric data, while effectively preventing the server from learning it.

The advantage of an adversary controlling both the client and the evaluator initially reduces to the definition of client privacy, as the adversary seeks to learn the reference public key stored during enrolment. However, after running one (unsuccessful) authentication attempt for a specific identity, the adversary will receive the encapsulated key derived from the biometric reference data of the data subject in question. From that point on, it can guess a biometric feature vector, issue an evaluation by use of the evaluation key, and compare the resulting key against the obtained one. Therefore, we realistically model an adversary controlling both the client and the evaluator as being able to run an offline search on the biometric enrolment database. Due to the architecture considerations, this scenario is somewhat unlikely in practice, and a more realistic threat is the server and evaluator colluding.

Definition E.13. (Client-Evaluator Privacy) We say that a BRAKE protocol has client-evaluator privacy if an adversary $\mathcal{A}$ controlling both the client and the authentication server does not have an advantage in obtaining a biometric feature vector t' that is close to any enrolled biometric template t above running a brute-force search on V:

$$\Pr\left[\begin{array}{c}\mathtt{dist}(t,t') < \tau \;:\; \forall i \in [\ell] : \begin{cases} \mathtt{pp} \leftarrow \mathtt{setup}(1^\lambda) \\ \{V, \mathtt{cpk}_t\} \leftarrow \mathtt{enroll}(\mathtt{pp}, t) \\ (B', \mathtt{cpk}_e) \leftarrow \mathcal{A}(\mathtt{pp}, \mathtt{id}, V) \\ (\mathtt{ssk}_e, \mathtt{spk}_e) \leftarrow \mathtt{KeyGen}(1^\lambda) \\ S' \leftarrow \mathcal{A}(B', k) \\ \mathtt{ctx} \leftarrow \mathtt{encap}(\rho, \mathtt{cpk}_t) \\ t' \leftarrow \mathcal{A}(S', \mathtt{spk}, \mathtt{spk}_e, \mathtt{ctx}) \end{cases}\end{array}\right] \leq \ell_{\mathcal{A}} f^{-1} + \mathsf{negl}(\lambda).$$

Definition E.14. (Server-Evaluator Privacy) We say that a BRAKE protocol has server-evaluator privacy if an adversary $\mathcal{A}$ controlling both the server and the evaluator does not have an advantage in obtaining a biometric feature vector

t' that is close to any enrolled biometric template t above running a brute-force search on V:

$$\Pr\left[\texttt{dist}(t,t') < \tau \;:\; \begin{array}{c} \texttt{pp} \leftarrow \texttt{setup}(1^\lambda) \\ \{V, \texttt{cpk}_t\} \leftarrow \texttt{enroll}(\texttt{pp}, t) \\ f' \leftarrow \texttt{unlock}(V, t') \\ B' \leftarrow \texttt{blind}(f') \\ (\texttt{csk}_e, \texttt{cpk}_e) \leftarrow \texttt{KeyGen}(1^\lambda) \\ t' \leftarrow \mathcal{A}(\texttt{pp}, \texttt{id}, V, B', k, \texttt{cpk}_t, \texttt{cpk}_e) \end{array}\right] \leq \ell_{\mathcal{A}} f^{-1} + \texttt{negl}(\lambda).$$

E.4 Concrete Instantiations

We now give two concrete instantiations of BRAKE, where the first is based on the hardness of discrete logarithms, while the second utilises lattice-based cryptography. Thereby, we show that both classical security and post-quantum security can be achieved using BRAKE. For both instantiations, the modified improved fuzzy vault scheme described in Section E.3.2 is used. The detailed description of the instantiations includes their cryptographic building blocks, complete instantiated protocols, and security proofs.

E.4.1 Instantiation Based on Discrete Logarithms

In this Section, we give an instantiation of the protocol defined in Figures E.2 and E.3 using cryptographic primitives that build on the security of discrete logarithms (DL). Concretely, we instantiate the universe $\mathcal{P}$ with a cyclic group $\mathbb{G}$, which can be the group of points on an elliptic curve, and the key space $\mathcal{K}$ and randomness space $\mathcal{R}$ with a scalar field $\mathbb{Z}_p$, where p is the prime order of $\mathbb{G}$. Further, we also define two hash functions $H_{\mathbb{G}} : \{0,1\}^* \to \mathbb{G}$ and $H_{\mathbb{Z}_p} : \{0,1\}^* \to \mathbb{Z}_p$.

Building on these foundations, the respective algorithms of Definition E.9 are instantiated with the Hash-DH OPRF defined in Definition E.4 and ephemeral Diffie-Hellman key exchange with a key-derivation function KDF. The detailed protocols for enrolment and verification are defined in Figures E.4 and E.5, respectively. In the following, we refer to the verification protocol in Figure E.5 as DL-BRAKE. We note that in the setting where the evaluator rate-limits the number of evaluations per user, the protocol can trivially be updated to send the identity of the user (or a fixed pseudonym) together with the blinded value, and the evaluator evaluates a partially oblivious PRF where the identity is a public input to the function together with the secret evaluation key. Implementing the techniques from [240, 255] allows us to perform this slightly different evaluation without (noticeable) increased computation nor communication compared to the protocol we have described.

DL-BRAKE enrolment protocol

Setup party	**Server**	**Evaluator**
t reference template	$\texttt{ssk} \in \mathbb{Z}_p$	$k \in \mathbb{Z}_p$
	$\texttt{spk} \in \mathbb{G}$	

$f \leftarrow\!\!\$\ \mathbb{F}_{p'}[x] : \deg(f) = \tau - 1$

$V(x) = f(x) + \prod_{a \in t}(x - a)$

$r \leftarrow\!\!\$\ \mathbb{Z}_p$

$B = [r]H_{\mathbb{G}}(f)$

B ⟶ B ⟶

$S = [k]B$

S ⟵ S ⟵

$U = [r^{-1}]S = [k]H_{\mathbb{G}}(x)$

$\texttt{csk}_t \leftarrow H_{\mathbb{Z}_p}(U)$

$\texttt{cpk}_t = [\texttt{csk}_t]G$

$V, \texttt{cpk}_t, \texttt{id}$ ⟶

store $(V, \texttt{cpk}_t, \texttt{id})$

Figure E.4: DL-BRAKE enrolment protocol instantiated with discrete-logarithm OPRF and Diffie-Hellman key exchange.

DL-BRAKE verification protocol

Client		**Server**	**Evaluator**
t' probe feature vector		$\mathtt{ssk} \in \mathbb{Z}_p$	$k \in \mathbb{Z}_p$
$\mathtt{spk} \in \mathbb{G}$		$\mathtt{spk} \in \mathbb{G}$	
		$(V, \mathtt{cpk}_t, \mathtt{id})$	
	$\xrightarrow{\mathtt{id}}$		
	$\xleftarrow{V}$		
find $\{(b, V(b)) : b \in t'\}$			
and decode to $f' \in \mathbb{F}_{p'}[x]$			
$r' \leftarrow\!\!\$\ \mathbb{Z}_p$			
$B' = [r']H_{\mathbb{G}}(f')$			
$\mathtt{csk}_e \leftarrow\!\!\$\ \mathbb{Z}_p$		$\mathtt{ssk}_e \leftarrow\!\!\$\ \mathbb{Z}_p$	
$\mathtt{cpk}_e = [\mathtt{csk}_e]G$		$\mathtt{spk}_e = [\mathtt{ssk}_e]G$	
	$\xrightarrow{B', \mathtt{cpk}_e}$	$\xrightarrow{B'}$	
		$\rho \leftarrow \mathtt{KDF}([\mathtt{ssk}_e]\mathtt{cpk}_e,$	$S' = [k]B'$
		$[\mathtt{ssk}]\mathtt{cpk}_e, [\mathtt{ssk}_e]\mathtt{cpk}_t,$	
		$\mathtt{cpk}_e, \mathtt{spk}_e, \mathtt{cpk}_t, \mathtt{spk})$	
	$\xleftarrow{S', \mathtt{spk}_e, H(\rho)}$	$\xleftarrow{S'}$	
$U' = [r'^{-1}]S' = [k]H_{\mathbb{G}}(x')$			
$\mathtt{csk}_{t'} \leftarrow H_{\mathbb{Z}_p}(U')$			
$\mathtt{cpk}_{t'} = [\mathtt{csk}_{t'}]G$			
$\rho' \leftarrow \mathtt{KDF}([\mathtt{csk}_e]\mathtt{spk}_e,$			
$[\mathtt{csk}_e]\mathtt{spk}, [\mathtt{csk}_{t'}]\mathtt{spk}_e,$			
$\mathtt{cpk}_e, \mathtt{spk}_e, \mathtt{cpk}_{t'}, \mathtt{spk})$			
return $H(\rho') = H(\rho)$			

Figure E.5: DL-BRAKE verification protocol instantiated with discrete-logarithm OPRF and Diffie-Hellman key exchange.

E.4.2 DL-BRAKE Security Proofs

In this Section, we provide theorems stating the security of the DL-BRAKE based on the hardness of discrete logarithms, and we sketch the security proofs.

Theorem E.1 (Correctness)**.** Assume that a probe sample t' is within the verification threshold τ compared to a biometric template $t_{\texttt{id}}$ for some registered identity `id`. Then the DL-BRAKE protocol in Figure E.5 is correct.

Proof sketch. This follows directly from the construction. If the comparison result of the probe feature set t' to a biometric template $\mathrm{t}_{\texttt{id}}$ is within the the verification threshold τ for some registered identity `id`, then the client will successfully reconstruct the correct polynomial f' using interpolation. From the correctness of the `OPRF`, the `KEM`, and the `KDF`, we then conclude that the client and the server compute the same values, and the data subject is correctly authorised. If the distance between probe and reference feature set is more than τ points, by correctness of Lagrange interpolation, two different polynomials will be reconstructed, and, but for a collision in the hash function, the key exchange will fail. □

Theorem E.2 (Client Privacy)**.** Let $\mathcal{A}_0$ be an adversary against *client privacy* in the DL-BRAKE protocol in Figure E.5 with advantage ϵ_0. Then there exists an adversary $\mathcal{A}_1$ against the fuzzy vault V with advantage ϵ_1 and an adversary $\mathcal{A}_2$ against the `OPRF` with advantage ϵ_2, such that $\epsilon_0 \leq \epsilon_1 + f^{-1}(1 + \epsilon_2)$. The runtime of $\mathcal{A}_0$ is essentially the same as of $\mathcal{A}_1$ and $\mathcal{A}_2$.

Proof sketch. We consider a single log-in attempt by an adversary $\mathcal{A}_0$ controlling the client. If $\mathcal{A}_0$ guesses a biometric probe, the probability that this probe is close to the reference sample is approximately f^{-1}. Furthermore, if $\mathcal{A}_0$ with probability ϵ_0 can output a valid probe sample t' given access to the fuzzy vault V, we can trivially turn $\mathcal{A}_0$ into an adversary $\mathcal{A}_1$ against V with the same advantage. Moreover, if $\mathcal{A}_0$ with advantage f^{-1} can output a valid probe sample t' when having access to values evaluated with key k, then we can turn $\mathcal{A}_0$ into an adversary $\mathcal{A}_2$ against the `OPRF`. Finally, we observe that the `KEM` are independent of $t_{\texttt{id}}$, and hence, an adversary $\mathcal{A}_0$ cannot learn anything from interacting with this protocol. We conclude that the protocol achieves client privacy. □

Theorem E.3 (Server privacy)**.** Let $\mathcal{A}_0$ be an adversary against server privacy in the DL-BRAKE protocol in Figure E.5 with advantage ϵ_0. Then there exists an adversary $\mathcal{A}_1$ against the fuzzy vault V with advantage ϵ_1 and an adversary $\mathcal{A}_2$ against the `OPRF` with advantage ϵ_2, such that $\epsilon_0 \leq \epsilon_1 + f^{-1}(1 + \epsilon_2)$. The runtime of $\mathcal{A}_0$ is essentially the same as of $\mathcal{A}_1$ and $\mathcal{A}_2$.

We omit the proof of Theorem E.3 since it is similar to Theorem E.2.

Theorem E.4 (Client-Evaluator Privacy)**.** Let $\mathcal{A}_0$ be an adversary against client-evaluator privacy in the DL-BRAKE protocol in Figure E.5 with advantage ϵ_0 controlling both the client and the evaluator. Then $\epsilon_0 \leq f^{-1}$ and $\mathcal{A}_0$ has no advantage in guessing a biometric probe within the threshold of an enrolled template above a brute-force search.

Proof sketch. We consider a colluding malicious client and malicious evaluator. Assume that $\mathcal{A}_0$ runs the verification protocol once on any input probe t' and receives $(S', \mathtt{spk}_e, H(\rho))$ from the server. Then $\mathcal{A}_0$ can guess a biometric probe, interpolate to get a polynomial f' and execute the `OPRF` on input f' using the evaluator's key k. For each guess, $\mathcal{A}_0$ can check if the `KDF` output corresponds to $H(\rho)$. No information about any enrolled template $t_{\mathtt{id}}$ is encoded in the messages from the server. □

Theorem E.5 (Server-Evaluator Privacy)**.** Let $\mathcal{A}_0$ be an adversary against server-evaluator privacy in the DL-BRAKE protocol in Figure E.5 with advantage ϵ_0 controlling both the server and the evaluator. Then $\epsilon_0 \leq f^{-1}$ and $\mathcal{A}_0$ has no advantage in guessing a biometric template within the threshold of an enrolled template above a brute-force search.

Proof sketch. We consider a colluding malicious server and malicious evaluator. Then $\mathcal{A}_0$ can guess a biometric probe, interpolate to get a polynomial f' and execute the `OPRF` on input f' using the evaluator's key k. For each guess, $\mathcal{A}_0$ can check if $[H_{\mathbb{Z}_p}(B')]G = \mathtt{cpk}_r$. No information about any enrolled template $t_{\mathtt{id}}$ is encoded in the messages from the client. □

E.4.3 Instantiation Based on Lattices

Our BRAKE protocol can also be instantiated with lattice-based cryptographic primitives, which are assumed to yield post-quantum security for certain parameter choices [11]. Two components in the protocol need to be instantiated: the OPRF and the KEM.

A construction of a lattice-based OPRF has recently been proposed by [12], which builds on the security of the M-LWE problem defined in Section E.2.5 for $d = 1$ (often referred to as the Ring-Learning With Errors (R-LWE) problem [183]). Additionally, this specific construction has the additional property of being *verifiable* (making it a VOPRF), i.e., the client has a guarantee that the output received from the OPRF evaluation is truly correct and calculated with the server's publicly committed key k [12, 60].

However, the zero-knowledge proof appended to the lattice-based PRF for verifiability are not practical for real-world application due to proof sizes of several gigabytes [12]. The authors of [12] give a rough indication of the amounts in question at approximately 2^{40} bits or around 128 GB of communication data for realistic parameter choices of $\log_2(q) \approx 256$ and ring dimension 16384. Therefore, we only look at the case of passive security against dishonest clients for the lattice instantiation, which can be significantly simplified by replacing the PRF with a hash function. We will give a detailed description of the modifications applied to the lattice-based VOPRF by [12] in the following.

E.4.3.1 Lattice OPRF

An option that is made possible by removing the zero-knowledge proofs is the ability to heavily reduce the computation time and communication cost generated by the PRF. Originally, the PRF is evaluated as

$$F_k(x) := \lfloor a_x \cdot k \rceil_{q'} \in R_{q'}^{d'},$$

where a_x is a lattice PRF [24]. This evaluation can be replaced with the PRF $F'_k(x) := \lfloor a_x \cdot k \rceil_{q'}$ where a_x a pseudorandom ring element output by a hash function evaluated on some secret input x. This truncation shrinks the calculations from a vector of polynomials to just single polynomials in $R_{q'}$.

In practical terms, the input a_x we wish to evaluate the OPRF on, is the random polynomial f generated by the fuzzy vault scheme. Therefore, the element f needs to be mapped to a ring element in a deterministic fashion. The procedure is described in the following steps:

1. Concatenate every coefficient of f into a string cf.
2. Create $h := H(cf)$ using a cryptographic hash function.
3. Produce N coefficients of the polynomial a_x by creating a hash of the form $h_i := H(i \| h)$ for $i = 0, ..., N-1$ using the same hash function as before and converting hashes into integers. Here, $\|$ denotes concatenation.
4. Reduce the coefficients of a_x mod q (if needed).

This procedure results in a polynomial a_x which is an element of the ring $R_q = \mathbb{Z}_q[X]/\langle X^N + 1 \rangle$ and can subsequently be used to compute an M-LWE sample. Using the truncated PRF described above, the lattice-based OPRF construction by Albrecht et al. [12] can be modified as will be described in the following Section. Figure E.6 shows the functioning of the modified OPRF, using the truncated PRF, in more detail. Here, D_σ is a uniform distribution over R_q which produces ternary values, and $D_{\sigma'}$ is a uniform distribution over R_q which

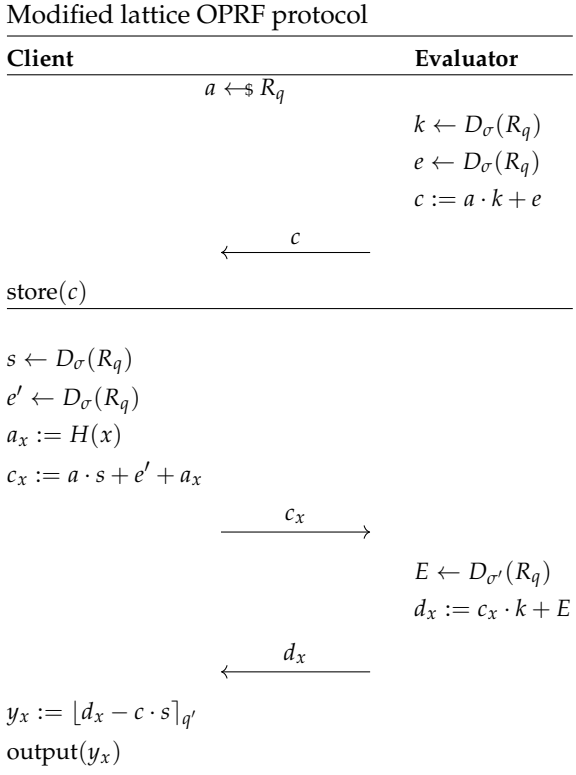

Figure E.6: Modified OPRF protocol based on [12] using the truncated PRF.

produces values in a range $[-B, B]$, where B is a large power of two smaller than q.

The final step, *rounding*, produces the Client's output, which is the polynomial y_x. If the rounding is implemented correctly and the protocol has been successfully executed, this rounded value will be equal to the rounded value $\lfloor a_x \cdot k \rceil_{q'}$. This is known as the unblinding operation, which allows the Client to receive the computation of $a_x \cdot k$ without learning the Evaluator's key k, while the Evaluator does not learn the value of a_x. Additionally, before rounding, it is necessary to represent the values that are to be rounded in $(-\frac{q-1}{2}, ..., \frac{q-1}{2})$.

The principle behind the validity of the rounding mechanism is shown in the following equations based on [12], which depict the total amount of noise that is accrued through the protocol. Firstly, we introduce the M-LWE samples c, d_x

PQ-BRAKE enrolment protocol

Client	**Server**	**Evaluator**
t reference template	$\texttt{ssk} \in \mathcal{K}$	$c \in R_q$
$\texttt{id}$ verified identity	$\texttt{spk} \in \mathcal{P}$	$k \in R_q$
$c \in R_q$		
$(f, V) \leftarrow \texttt{lock}(t)$		
$s \leftarrow D_\sigma(R_q)$		
$e' \leftarrow D_\sigma(R_q)$		
$a_x := H(f)$		
$c_x := a \cdot s + e' + a_x$		
$\xrightarrow{c_x}$		$\xrightarrow{c_x}$
		$E \leftarrow D_{\sigma'}(R_q)$
		$d_x := c_x \cdot k + E$
$\xleftarrow{d_x}$		$\xleftarrow{d_x}$
$y_x := \lfloor d_x - c \cdot s \rceil_{q'}$		
$\texttt{csk}_t \leftarrow \texttt{ExpKDF}(y_x)$		
$\texttt{cpk}_t \leftarrow \texttt{pkGen}(\texttt{csk}_t)$		
$\xrightarrow{V, \texttt{cpk}_t, \texttt{id}}$		
	store $(V, \texttt{cpk}_t, \texttt{id})$	

Figure E.7: PQ-BRAKE enrolment protocol instantiated with modified lattice OPRF and Kyber KEM.

PQ-BRAKE verification protocol

Client		**Server**		**Evaluator**
t' probe feature vector		$\mathtt{ssk} \in \mathcal{K}$		$c \in R_q$
$\mathtt{spk} \in \mathcal{P}$		$\mathtt{spk} \in \mathcal{P}$		$k \in R_q$
biometric claim $\mathtt{id}$		$(V, \mathtt{cpk}_t, \mathtt{id})$		
$c \in R_q$				
	$\xrightarrow{\mathtt{id}}$			
	$\xleftarrow{V}$			
$f' \leftarrow \mathtt{unlock}(V, t')$				
$(\mathtt{csk}_e, \mathtt{cpk}_e) \leftarrow \mathtt{KeyGen}()$		$(\mathtt{ssk}_e, \mathtt{spk}_e) \leftarrow \mathtt{KeyGen}()$		
$s \leftarrow D_\sigma(R_q)$				
$e' \leftarrow D_\sigma(R_q)$				
$a_x := H(f')$				
$c_x := a \cdot s + e' + a_x$				
	$\xrightarrow{c_x, \mathtt{cpk}_e}$		$\xrightarrow{c_x}$	
		$(\mathtt{ctx}, \gamma) \leftarrow \mathtt{encap}(\mathtt{cpk}_t)$		$E \leftarrow D_{\sigma'}(R_q)$
		$\rho \leftarrow \mathtt{KDF}(\mathtt{cpk}_t, \mathtt{cpk}_e,$		$d_x := c_x \cdot k + E$
		$\mathtt{spk}, \mathtt{spk}_e, \gamma)$		
	$\xleftarrow[\mathtt{spk}_e, ctx]{d_x, H(\rho)}$		$\xleftarrow{d_x}$	
$y_x := \lfloor d_x - c \cdot s \rceil_{q'}$				
$\mathtt{csk}_{t'} \leftarrow \mathtt{ExpKDF}(y_x)$				
$\mathtt{cpk}_{t'} \leftarrow \mathtt{pkGen}(\mathtt{csk}_{t'})$				
$\gamma' \leftarrow \mathtt{decap}(\mathtt{ctx}, \mathtt{csk}_{t'})$				
$\rho' \leftarrow \mathtt{KDF}(\mathtt{cpk}_{t'}, \mathtt{cpk}_e,$				
$\mathtt{spk}, \mathtt{spk}_e, \gamma')$				
return $H(\rho') = H(\rho)$				

Figure E.8: PQ-BRAKE verification protocol instantiated with modified lattice OPRF and Kyber KEM.

and c_x, which form the total noise value. These are elements of R_q and are transmitted between the Client and Evaluator during the protocol. We recall their definitions as given in Figure E.6:

$$\begin{aligned} c &= a \cdot k + e \\ d_x &= c_x \cdot k + E \\ c_x &= a \cdot s + e' + a_x. \end{aligned}$$

Next, we recall the computation of the polynomial y on the Client's side, which includes the values d_x, c and s before they are summed and rounded in y_x:

$$\begin{aligned} y &= d_x - c \cdot s \\ &= c_x \cdot k + E - (a \cdot k + e) \cdot s \\ &= (a \cdot s + e' + a_x) \cdot k + E - a \cdot k \cdot s + e \cdot s \\ &= e' \cdot k + a_x \cdot k + E - e \cdot s. \end{aligned}$$

Then, as the polynomial y_x can be obtained from y as:

$$y_x = \left\lfloor \frac{q'}{q} \cdot (d_x - c \cdot s) \right\rceil = \left\lfloor \frac{q'}{q} \cdot a_x \cdot k \right\rceil.$$

In the expanded equation for y, we notice that it contains the polynomial $a_x \cdot k$ and and a noise polynomial $e' \cdot k - e \cdot s + E$. Therefore, the last equation, showing the value of y_x, is correct with all but a negligible probability if the noise polynomial $\left| \frac{q'}{q} \cdot (e' \cdot k - e \cdot s + E) \right|$ is small enough for each coefficient to achieve acceptable correctness after rounding. In other words:

$$\left| \frac{q'}{q} \cdot (e' \cdot k - e \cdot s + E) \right|_\infty < \frac{1}{2}.$$

E.4.3.2 CRYSTALS Kyber Key Encapsulation Mechanism

We exchange the Diffie-Hellman key exchange with a lattice-based KEM: the recently standardised CRYSTALS-Kyber [49]. Kyber is based on the M-LWE problem described in Section E.2.5 and provides IND-CCA2 security [21]. The main parameters of Kyber, $N = 256$ and $q = 3329$, were specifically chosen for the ability to use the Number Theoretic Transform (NTT) providing an efficient way to perform multiplications in R_q [21]. In our work, the parameter set of Kyber768 was chosen due to its optimal performance while providing more than 128 bits of security [21]. While no significant changes were applied to Kyber on a theoretical basis, we give further details on the integration of Kyber into the implementation of the BRAKE protocol in Section E.5. In particular, we note that the security of the session key established through BRAKE is given through the security guarantees of Kyber.

E.4.3.3 PQ-BRAKE

Combining the introduced modified lattice OPRF and the Kyber KEM, we can define the PQ-BRAKE protocol as described in Figures E.7 and E.8.

E.4.4 PQ-BRAKE Security Proofs

The security proofs for PQ-BRAKE follow directly from the proofs given for the DL-BRAKE instantiation given in Section E.4.2 through the hardness of M-LWE and M-SIS.

E.4.5 Improved Security using NIZK

The protocol can be further secured by the addition of non-interactive zero-knowledge proofs (NIZKs) using the established construction by Chaum and Pedersen [62] together with a Fiat-Shamir transform [113]. The NIZK is added to prove the honest evaluation of the OPRF. Thereby, a client can verify that the evaluator computed the evaluation honestly. In the case of an unsuccessful authentication attempt, the client therefore gains more knowledge about the reason of failure, and can potentially reveal a corrupted evaluator. We note that above this additional information, the passively secure protocol already allows for the protection of the biometric data even in the presence of malicious adversaries, as long as at least one of the parties remains honest as given by the security definitions above. However, in the lattice-based instantiation, a malicious client may be able to learn the OPRF key, facilitating a similar attack as in the case of a colluding client and signer. Therefore, the lattice-based instantiation can only be considered in the semi-honest adversary model.

E.5 Experimental Evaluation

We evaluated our protocol instantiated with elliptic curves presented in Figure E.5 and lattices presented in Figure E.8 experimentally and show the results in this Section. Our experiments were run on a commodity notebook with Intel Core i7-8565U CPU@1.80GHz and 8GB RAM. Our code is available at `https://github.com/dasec/DL-BRAKE` and `https://github.com/dasec/PQ-BRAKE` and includes automated installation scripts with all dependencies in order to support the reproducibility of our work.

To begin, we give a more detailed comparison of our work with closely related work in Table E.2 by extending Table 1 in [260] with our protocol. In terms of round efficiency, our protocol compares well to [100] and [106] with two rounds

Scheme	Technique	Rounds	Communication Cost	Compatibility	ISO/IEC 24745 [149]
fPAKE-1 [100]	Garbled Circuits	5	N/A		✗
fPAKE-2 [100]	PAKE + Secret Sharing	2	N/A	iris, fixed-length fingerprint	✗
fuzzy aPAKE-1 [106]	ECC + OT	2	∼ 700 KB		✗
fuzzy aPAKE-2 [106]	Generic k-parallel aPAKE	2	∼ 1 MB		✗
BAKE-1 [260]	Random Linear Codes	1	5-8.4 KB	minutiae-based fingerprint	✗
BAKE-2 [260]	Secret Sharing + Polynomial Interpolation	1	1.7-96.6 KB	iris	✗
iPAKE [48]	ECC + PAKE	1	N/A	iris, fixed-length fingerprint	✗
DL-BRAKE (ours)	Fuzzy Vault + DL-OPRF + DL-KEM	2	0.3 KB	minutiae-based fingerprint, iris, face	✓
PQ-BRAKE (ours)	Fuzzy Vault + lattice OPRF + lattice KEM	2	60.2 KB	minutiae-based fingerprint, iris, face	✓

Table E.2: Summary of our protocol compared to previous published protocols as described in Table 1 of [260].

of communication. In order to prevent offline attacks, a minimum number of two rounds of communication is necessary. Therefore, [100], [106], and our protocol can be considered optimal in terms of number of rounds. As [260] constructed a one-round protocol, this leaves them open to offline attacks. In terms of the protection of the biometric data compliant with ISO/IEC 24745 [149], our protocol is the only compliant one: we inherit unlinkability, renewability, and irreversibility from the fuzzy vault schemes. Moreover, we show that our protocol is efficient in terms of execution times given in Table E.3 and as well as in terms of biometric performance shown in Figure E.9. In comparison, fPAKE [100] does not achieve irreversibility as templates are disclosed to the server in plaintext, fuzzy aPAKE [106] does not achieve computational efficiency, and [260] does not achieve an acceptable biometric performance, as we show in Appendix E.6.

E.5.1 Fuzzy Vault Implementation

For the fingerprint fuzzy vault instantiation, we used the open-source implementation provided by [251] with all original parameter settings, in particular, the minutiae quantisation and encoding into a product of finite field $\mathbb{F}_{2^{18}} \times \mathbb{F}_{2^{18}}$ which accommodates a unique encoding of at most $t_{max} = 44$ genuine minutiae as described in [251]. Keeping the parameter choices evaluated in the work of [251] ensures perfect replaceability with other state-of-the-art fuzzy vault instantiations, such as [220] for iris and [218] for face. In particular, we run our implementation on the same fingerprint database MCYT-330 [199] and same feature extractor, Digital Persona's FingerJetFX open source edition minutiae extractor[1]. This means that all evaluations of biometric performance can be compared directly to the original paper of [251] and papers that compare their work with the latter [218, 220].

The only modification applied to the implementation of [251] is in the unlocking function. Here, [251] use the stored hash $H(f)$ of the secret polynomial f corresponding to a reference template t, which allows for offline brute force attacks. Our protocol prevents offline attacks by removing the hash and using highest-frequency decoding in its place (see Section E.3.2). As discussed above, this does not impact the security in terms of the false-match rate of our protocol.

E.5.2 DL-BRAKE Implementation

Our implementation of the OPRF and Diffie-Hellman key exchange is based on OpenSSL. For all cryptographic operations, we used P-256 [25] as the elliptic

[1] `http://www.digitalpersona.com/fingerjetfx`

	Polynomial degree $\tau - 1$					
	6	8	10	12	14	16
Feature extraction and preprocessing	200.59					
`lock`	2.38					
`unlock`	112.24	185.99	276.37	385.26	511.91	694.87
`DL−OPRF`	0.21					
`PQ−OPRF`	31.81					
`DL−KeyGen`	0.05					
`PQ−KeyGen`	0.21					
`DL−encap`	0.16					
`PQ−encap`	0.08					
`DL−decap`	0.15					
`PQ−decap`	0.03					
DL-Verification (Figure E.5)	313.4	387.15	477.53	586.42	713.07	896.03
PQ-Verification (Figure E.8)	347.34	421.09	511.91	620.36	747.01	929.97
FMR (%)	1.04%	0.04%	0.00%	0.00%	0.04%	0.09%
1− FNMR (%)	92.88%	88.79%	81.97%	73.18%	60.45%	44.09%
Estimated security in bits based on [251]	17	23	29	36	44	—

Table E.3: Execution times in milliseconds for the DL-BRAKE and PQ-BRAKE protocols using the fingerprint fuzzy vault by [251].

	DL-BRAKE	PQ-BRAKE
Locked fuzzy vault	99 B	
OPRF	128 B	114 KB
KEM	64 B	4672 B
Hash digest	32 B	
Total	0.3 KB	60.2 KB

Table E.4: Communication cost for DL-BRAKE and PQ-BRAKE.

curve and SHA-256 as the hash function.

Regarding the computational performance and recognition accuracy of our protocol, we give timings for increasing polynomial degrees $\tau - 1$ in Table E.3, where τ is the biometric decision threshold. At the same time, we give the biometric performance in FMR and FNMR along with the estimated false-accept security in bits as evaluated in [251]. As these security levels are derived from the FMR and our modified unlocking function does not impact the FMR, we are able to refer to the evaluation performed in [251] directly. For an acceptable recognition accuracy at $\tau - 1 = 8$, the execution of the protocol DL-BRAKE given in Figure E.5 takes 387.15 milliseconds. To compare, the fastest setting reported in Table 2 in [260] also achieves 387 milliseconds, but at significantly lower accuracy (see Appendix E.6).

The execution times are dominated by the constant cost of feature extraction (200.59 milliseconds) and the cost for unlocking, which is dependent on the polynomial degree. We note that timing for the enrolment part of the protocol given in Figure E.4 is 203.23 milliseconds, where feature extraction dominates compared to the locking at 2.38 milliseconds. However, the enrolment step is a one-time effort when setting up the system, and does not affect verification performance.

Accordingly, Figure E.9 shows the trade-off between FMR and FNMR for our protocol. To conclude the efficiency evaluation of our protocol, we report that the communication cost of objects transferred between the parties during the verification step of the protocol is 32 bytes for any point on the elliptic curve P-256 [25] (i.e., $\texttt{cpk}_e$, $\texttt{spk}_e$, B' and S'), 99 bytes for a locked fuzzy vault of degree at most 43 and coefficients in $\mathbb{F}_{2^{18}}$, and 32 bytes for the hash digest.

E.5.3 PQ-BRAKE Implementation

For the lattice-based instantiation of our protocol, we utilised the OpenSSL implementation of the SHA-256 hash function, Open Quantum Safe's `liboqs` C library [244] through its C++ wrapper, `liboqscpp`, for the CRYSTALS Kyber [49] implementation. To support key generation from a designated input (i.e., the fuzzy vault secret polynomial f), we extended the C++ wrapper to include the functionalities required for BRAKE. The documentation can be found in our repository at `https://github.com/dasec/PQ-BRAKE`.

For the OPRF part of the protocol, parameter choice is crucial for both communication and computation complexity along with security, and needs to be carefully evaluated. We therefore tested our parameter validity using the established `lwe-estimator` [13]. As a result, we chose the parameters $N = 4096$,

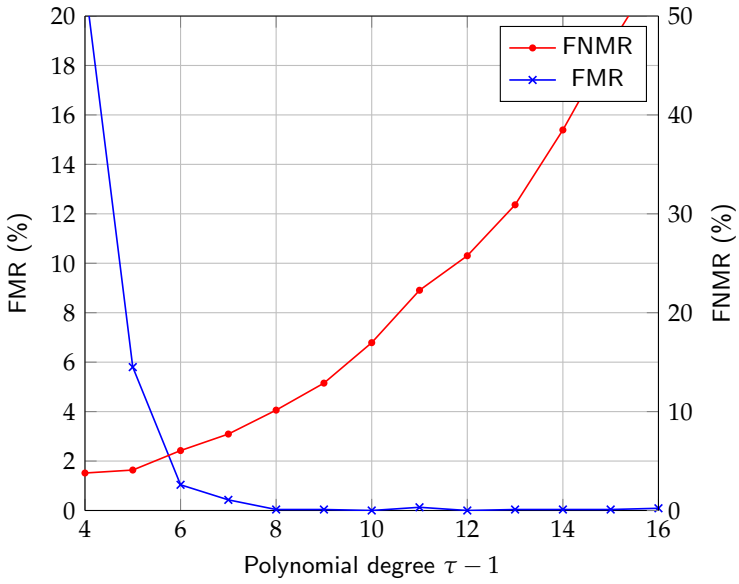

Figure E.9: Biometric performance for the DL-BRAKE protocol instantiated with fingerprint fuzzy vault [251].

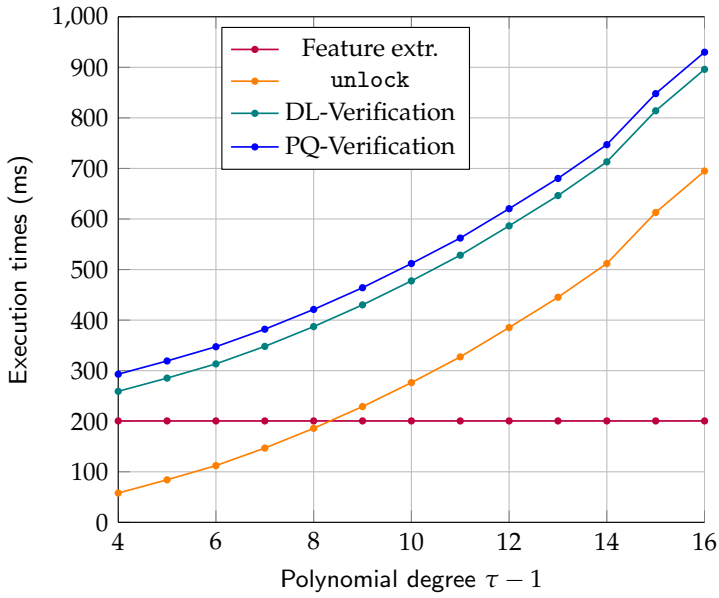

Figure E.10: Execution times in milliseconds for the DL-BRAKE and PQ-BRAKE protocols instantiated with fingerprint fuzzy vault [251].

$q \approx 2^{75}$, and $B = 2^{53}$ with security of 188 bits. In comparison, the Kyber KEM is instantiated with $N = 256$ and $q = 3329$.

Using these parameters, it is also possible to calculate a probability of the rounding step failing, which would result in a decryption failure in practice, due to noise wrapping the value around $\mathbb{Z} + 1/2$ and causing a rounding to the wrong value. As demonstrated in Section E.4.3.1, the upper bound on the noise is given as: $2N + B \leq \frac{q}{4}$. We consider the probability of one coefficient of the output polynomial y_x being wrongly decrypted to be: $\frac{2N+B}{q}$, and its complement situation, the probability of no error occurring as $1 - \frac{2N+B}{q}$. With this in mind, we claim that the probability of at least one decryption error occurring during the rounding of N polynomial coefficients and thus the protocol failing in the OPRF step, to be

$$1 - \left(1 - \frac{2N + B}{q}\right)^N. \tag{E.3}$$

Applying this formula, we set the parameters so that the failure rate is significantly smaller than the false-accept security of the biometric component, i.e., the improved fuzzy vault scheme. A success rate of 99.9% was chosen for this benchmark.

The computational performance of the PQ-BRAKE protocol can be seen in Table E.3. Compared to DL-BRAKE, the most significant change is the lattice-based OPRF, which has a significantly higher computational workload of 31.81 milliseconds compared to the classically secure OPRF at only 0.21 milliseconds. However, compared to the overwhelming cost of feature extraction, preprocessing, and the unlocking step of the fuzzy vault, the lattice OPRF cost can still be considered feasible. A visual comparison of the execution times for both the Dl-BRAKE and PQ-BRAKE protocols as well as the fixed costs of feature extraction and the individual effort of the fuzzy vault unlocking step is given in Figure E.10.

The communication cost for PQ-BRAKE can be determined as 99 bytes for a locked fuzzy vault as before, 114KB for the OPRF, covering a total of three R-LWE samples, a total of 4672 bytes for the Kyber key exchange, and 32 bytes for the has digest. A comparison of the communication cost for DL-BRAKE, PQ-BRAKE, and the original lattice VOPRF by Albrecht et al. [12] can be seen in Table E.4.

E.6 Conclusions

In this work, we constructed biometric resilient authenticated key exchange from fuzzy vaults and proved its security in compliance with ISO/IEC 24745. Our protocol is efficient both in terms of execution times and biometric performance.

The combination of asymmetric, secure, and efficient biometric authenticated key exchange has not been achieved in prior works. Related protocols are either symmetric, and thus does not provide protection of the biometric data on the server side, or inefficient in terms of computational speed due to their generality, or else insufficient in terms of recognition accuracy, allowing for zero-effort imposter and low-effort brute-force attacks. The accuracy deficiencies of the latter cannot be addressed by exchanging the biometric comparison subsystem, as the construction is specific to the imprecise comparator used.

In our protocol, we enforce communication for every adversarial guess through OPRFs. Using established and interchangeable improved fuzzy vault schemes for different biometric modalities, the key exchange is only successful if the two biometric samples were close. Furthermore, we show that our protocol can be instantiated both with classical primitives, namely discrete logarithm based OPRFs and Diffie-Hellman key exchange, as well as with lattice-based OPRFs and KEMs.

Future works may focus on addressing the necessary pre-alignment processes of minutiae-based fingerprint representations. A promising approach both with regard to rotation and entropy is the use of four-finger captures, where four fingerprints are captured within one image. Through the relative position of the fingers, pre-alignment can be realised more efficiently than based on minutiae, and the intra-identity independence of fingerprint patterns yield the fourfold entropy of the biometric data. Notably, the implementation of the minutiae fuzzy vault evaluated in our work includes the option of combining four fingerprints into one fuzzy vault. However, auxiliary alignment data required for pre-alignment are not yet discussed in this context.

Appendix E.6: Biometric Performance Analysis

In this Appendix, we give the experimental evaluation of the recent work on biometrics-authenticated key exchange proposed by [260]. Specifically, we show the biometric performance of their construction for fingerprint and discuss its shortcomings.

For this evaluation, we implemented Algorithm 2 in [260] according to the description available in the paper. According to the description, we set the number of neighbours for each minutia at $\mu = 4$ and, iterating through the minutiae in the template, construct the vectors $v_{j,\rho}$ from the minutia's x- and y-coordinates which are given in pixels (i.e., integers) from the upper left corner. The calculation of the Euclidean distances $d_{j,1}, ..., d_{j,4}$ therefore result in floating point numbers, whereas the angles $\phi_{j,\rho,1}, ..., d_{j,\rho,6}$ remain as integer values. In Section 6.2.2 in [260], the authors state that the number of neighbours $\mu = 4$ originates an encoding of the values $d_{j,\rho}$ and $\phi_{j,\rho,\omega}$ into $\mu = 4$ bits each. This relation is not clear to us and we were not able to satisfactorily follow the reasoning given by the authors of [260] during an email exchange. Therefore, we give the evaluation of the biometric performance for the original float and integer values, which can be considered an upper bound for the performance of a binary encoding. As comparison function, we determined the set difference by mapping minutiae based on their minimal Hamming distance.

We evaluated our implementation of Algorithm 2 in [260] on the FVC2004 DB-1 [252], which is the least challenging out of the four databases used in [260] in terms of image quality and rotation of the fingerprint images. We compare the performance against a state-of-the art rotation invariant minutiae comparator, SourceAFIS [259]. From the evaluation, it becomes evident that the fingerprint comparison algorithm proposed by [260] does not have an acceptable performance (see Table E.5).

	FVC2004 DB-1 [252]		CASIA-FPV5[3]	
	FMR	FNMR	FMR	FNMR
BAKE [260]	27.8%	25.4%	27.6%	30.90%
SOTA[2]	1.01%	17.29%	1.13%	9.85%

Table E.5: Biometric performance of BAKE [260] compared to state-of-the-art (SOTA) performance.

For the optimal threshold, the FMR is measured at 27.8% with a FNMR of 25.4%. Both of these values are not close to the required FMR of 0.1% [117] and FNMR

below 5%. Compared to the state-of-the-art, the performance that can be achieved in this dataset lies at a FMR of 1.01% at FNMR of 17.29% using the SourceAFIS comparison algorithm[2]. This shows the challenging nature of the dataset, which was collected as a fingerprint verification challenge with the goal of providing challenging fingerprint samples. Therefore, we also evaluated both algorithms on the less challenging CASIA-FPV5[3] database. However, the result are similar with a FMR of 27.6% and FNMR of 30.90% for BAKE-1 compared to a FMR of 1.13% and FNMR of 9.85% for SourceAFIS.

To conclude, the fingerprint comparison algorithm proposed for the construction in [260] is not able to distinguish between mated and non-mated comparison trials to a satisfactory degree.

[2] `https://sourceafis.machinezoo.com/`
[3] `http://biometrics.idealtest.org`

Appendix E.7: Notation

	Parameter	Explanation
Generic	t	Biometric lock feature set.
	t'	Biometric unlocking feature set.
	f	Secret random polynomial.
	τ	Correction capacity of $\mathcal{C}$.
	$\mathbb{F}_{p'}$	Finite field for minutiae encoding.
	$\mathcal{C}$	Error-correcting code.
	H	Cryptographic hash function.
	λ	Security level.
	V	Locked fuzzy vault.
	f_k	Pseudorandom function with key k.
	x	Secret client input for OPRF.
	r	Randomness sampled by client.
	B, B'	Blinded OPRF input.
	S, S'	OPRF evaluation.
	U, U'	Unblinded OPRF evaluation.
	k	Secret OPRF evaluation key.
	`pp`	Public parameters.
	`id`	Biometric claim.
	`csk`	Client secret key.
	`cpk`	Client public key.
	`ssk`	Static server secret key.
	`spk`	Static server public key.
	(`sk`, `pk`)	Ephemeral asymmetric keys.
	γ	Session pre-key.
	`ctx`	Encapsulation of session pre-key γ.
	γ'	Decapsulation of session pre-key γ.
	`KDF`	Key derivation function.
	ρ	Session key.
	f^{-1}	False-accept security.
	l	Rate limit enforced by the server.
	$\mathcal{A}$	Adversary.
	$l_{\mathcal{A}}$	Brute-force capacity of adversary.
	ϵ	Adversary advantage.
Group setting	p	Prime group order.
	$\mathbb{G}$	Cyclic group.
	$\mathbb{Z}_p$	Scalar field of order p.
	$H_{\mathbb{G}}$	Cryptographic hash function $H_{\mathbb{G}} : \{0,1\}^* \to \mathbb{G}$.
	$H_{\mathbb{Z}_p}$	Cryptographic hash function $H_{\mathbb{Z}_p} : \{0,1\}^* \to \mathbb{Z}_p$.
Lattice setting	q	Ciphertext modulus.
	$\mathcal{R}_q$	Cyclotomic ring $R_q = \mathbb{Z}_q[X]/\langle X^N + 1\rangle$.
	N	Ring dimension of cyclotomic ring.
	χ	Bounded distribution over R_q.
	d	Module dimension for M-LWE.
	s	M-LWE secret sampled from χ.
	e	M-LWE error sampled from χ.
	m	Number of M-SIS vectors.
	β	Bound for M-SIS solutions.
	$\mathcal{D}_\sigma$	Ternary distribution over $\mathcal{R}_q$.
	$\mathcal{D}_{\sigma'}$	Uniform distribution over $\mathcal{R}_q$ bounded by $[-B, B]$.
	B	Bound for $\mathcal{D}_{\sigma'}$.

Table E.6: Overview of parameters.

Paper F

Post-Quantum Secure Biometric Systems: An Overview

Pia Bauspieß

Submitted to IEEE Access, 2024

This paper is submitted for publication and is therefore not included.

ISBN 978-82-326-8044-3 (printed ver.)
ISBN 978-82-326-8043-6 (electronic ver.)
ISSN 1503-8181 (printed ver.)
ISSN 2703-8084 (online ver.)

NTNU
Norwegian University of
Science and Technology