BERGISCHE
UNIVERSITÄT
WUPPERTAL

Cumulative Habilitation Thesis

Theoretically-Sound Real-World Cryptography

Dr.-Ing. Kai Gellert

December 15, 2023

Submitted to the
School of Electrical, Information and Media Engineering
University of Wuppertal

Kai Gellert
Place of birth: Bochum, Germany
Author's contact information:
gellert@uni-wuppertal.de

First examiner:	**Prof. Dr.-Ing. Tibor Jager** University of Wuppertal, Germany
Second examiner:	**Prof. Dr. rer. nat. Christopher Brzuska** Aalto University, Finland
Third examiner:	**Prof. Dr. rer. nat. Kathrin Hövelmanns** Eindhoven University of Technology, Netherlands
Thesis submitted:	December 15, 2023
Last revision:	December 15, 2023
Habilitation colloquium:	November 12, 2024

Acknowledgements

I am grateful to many people whose support, mentorship, and encouragement have been vital in shaping my academic career and the completion of this habilitation thesis.

I owe special thanks to Tibor Jager, without whom I might never have embarked on an academic career. Fortunately, he invested significant time and energy into encouraging me to pursue this path, presenting the right arguments at the perfect moment, which ultimately convinced me that academia is a world where I could thrive. I am deeply grateful for his unwavering support and the numerous opportunities he has provided.

I would also like to express my gratitude to my co-authors, Nimrod Aviram, Manuel Barbosa, Colin Boyd, Fynn Dallmeier, Gareth Davies, David Derler, Denis Diemert, Jan Drees, Sebastian Faller, Dennis Funke, Kristian Gjøsteen, Tobias Handirk, Julia Hesse, Máté Horváth, Tibor Jager, Stanislaw Jarecki, Håkon Jacobsen, Bor de Kock, Jonas Klauke, Lin Lyu, Lise Millerjord, Robert Merget, Tom Neuschulten, Timo Renzelmann, Jörg Schwenk, Daniel Slamanig, Christoph Striecks, and Rudi Wolf. Each of you has contributed something invaluable to my academic journey, and I would not be where I am today without our work together. Thank you for the long, productive nights with rapidly approaching deadlines and for all the lively discussions (often sparked by Reviewer C). It has been an absolute pleasure working with each of you.

I am grateful to my colleagues, both past and present, for fostering a collaborative and supportive atmosphere in our group. Working with Peter Chvojka, Denis Diemert, Gareth Davies, Amin Faez, Dennis Funke, Raphael Heitjohann, Tobias Handirk, Christian Holler, Tibor Jager, Saqib Kakvi, Jan Drees, Sebastian Lauer, Lin Lyu, Jutta Maerten, Tom Neuschulten, David Niehues, Sebastian Overhage, Moritz Schmidt, Marloes Venema, and Jonas von der Heyden has been a deeply rewarding experience that I will never forget. Special thanks also to my students, whose curiosity and engagement constantly inspire and

motivate me.

Lastly, I want to thank my brother Len for his support, Babette for her emotional reassurance during moments of doubt, Arian for guiding me through the jungle of linguistic dilemmas, and my two cats, Neo and Bella, for their comforting presence.

Abstract

Cryptography is an indispensable aspect of our daily lives. From securing online banking and protecting personal information to enabling secure connections over the Internet and protecting critical infrastructures, it permeates every fiber of our modern society. Understanding the security guarantees of these cryptographic schemes is crucial to ensuring their effectiveness in the real world.

The theoretical foundation of modern cryptography was laid in a seminal work by Goldwasser and Micali in 1984. Their pioneering work applied well-established techniques from theoretical computer science to conduct a formal security analysis of a cryptographic scheme, allowing for the first time a precise analysis of the security guarantees achieved by a cryptographic scheme. This approach is today called the *provable security* paradigm. Nowadays, it is considered good practice to not only propose a new cryptographic scheme, but also to present a security proof, ensuring a rigorous analysis of its security guarantees.

However, despite the theoretical rigor of provable security, challenges still persist in applying its principles to real-world cryptographic schemes. The idealized nature of security models, which often make assumptions about how adversaries may attack a cryptographic scheme, can diverge from the complexities of the real world. While these idealizations help to build a universal theory of secure cryptography, the real-world meaningfulness of security proofs diminishes when they fail to accurately model the real world.

This thesis explores the gap between the idealized treatment of security in cryptography and its real-world applications. The overarching goal is to advance our comprehension of theoretically-sound real-world cryptography, addressing foundational research questions that aim to further bridge the gap between theory and practice.

Contents

Abstract iii

1 Introduction 1
- 1.1 A Gap Between Theory and Practice . . . 3
- 1.2 Publication Overview and Contributions . . . 4
- 1.3 Editorial Remarks . . . 7

2 Theoretically-Sound Instantiation 9
- 2.1 Signatures with Tight Multi-User Security . . . 12
  - 2.1.1 Results . . . 14
  - 2.1.2 Possible Future Research Directions . . . 15
- 2.2 Signatures with Memory-Tight Multi-Challenge Security . . . 16
  - 2.2.1 Results . . . 18
  - 2.2.2 Possible Future Research Directions . . . 19
- 2.3 Key Exchange with Optimal Tightness . . . 20
  - 2.3.1 Results . . . 21
  - 2.3.2 Possible Future Research Directions . . . 21

3 Cryptography for the Real World 23
- 3.1 Forward-Secure Symmetric Key Exchange . . . 23
  - 3.1.1 Results . . . 26
  - 3.1.2 Possible Future Research Directions . . . 27
- 3.2 A Security Analysis of the WhatsApp Backup Protocol . . . 27
  - 3.2.1 Results . . . 28
  - 3.2.2 Possible Future Research Directions . . . 29
- 3.3 A New Perspective on Length-Hiding Encryption . . . 30
  - 3.3.1 Results . . . 32
  - 3.3.2 Possible Future Research Directions . . . 33

Bibliography 33

1 Introduction

The Art of Cryptography. The word *cryptography* stems from the Ancient Greek words κρυπτός (romanized "kryptós"; meaning "hidden, secret") and γράφειν (romanized "gráphein"; meaning "to write") and could be translated literally into "secret writing." In close relation to this etymology, early versions of the Oxford English Dictionary defined cryptography as

> "*the art of writing or solving codes.*"[1]

This definition is indeed historically accurate for two reasons.

First, cryptography was originally used to deliver confidential messages during times of war, *i.e.*, if a messenger were to be intercepted, the message should be incomprehensible to anyone but the intended receiver. Hence, messages were often written in a "code," concealing their contents. Since then, cryptography has evolved into a broad field encompassing a wide range of applications beyond military purposes. Today cryptography is an omnipresent companion. It permeates our daily lives, securing Internet connections and communication through messaging services and enabling secure, contactless payments during shopping. Even intricate systems such as automobiles and "smart" energy networks rely on cryptography. A world without it is nearly unimaginable.

Second, creating a secure cryptographic scheme could have been considered a form of art—a scheme was postulated secure until an attack was found, while both creating a scheme and finding an attack were limited only by the creator's and attacker's creativity. This led to a perpetual game of cat-and-mouse where creators and attackers continuously tried to outsmart the other. However, the lack of a solid theoretical foundation in cryptography left a crucial question unanswered. What does security actually mean and how can it be achieved in practice?

[1] See, *e.g.*, `https://www.oed.com/dictionary/cryptography_n`.

The Science of Cryptography. In 1984, Goldwasser and Micali published a seminal work that revolutionized the field of cryptography [GM84]. They were the first to rigorously *define* what security for encryption could mean and, using well-established techniques from the area of complexity theory, to carefully *prove* that their scheme satisfies the security definition. Their work marks the advent of what is now known as *provable security* and "transformed" the *art* of cryptography into the *science* of cryptography. Both were awarded the Turing Award in 2012 for their pioneering work that now serves as the foundation of modern cryptography. Charles Rackoff pointedly described their achievements as

> "*Julius Caesar may have used cryptography, but now we were finally beginning to understand it.*"[2]

The Provable Security Paradigm. When creating a cryptographic scheme, it is nowadays considered good practice to give a *security proof* as well. A security proof usually consists of three parts.

1. A formal *security definition*, which precisely defines what security means for a given primitive. To this end, it defines (i) the *attacker model*, defining the capabilities of an adversary (*e.g.*, if an attacker can manipulate messages that are sent over a network), and (ii) the *attacker goal*, defining when the attacker has successfully broken security (*e.g.*, by recovering a secret key or a secret message). Both the attacker model and the attacker goal have to be part of any meaningful security definition.

2. A *computational problem*, which is assumed to be "hard," meaning that there should not exist an efficient algorithm that can solve the problem, or that it is infeasible to efficiently find such an algorithm. Examples for hard problems are well-studied mathematical problems, such as computing discrete logarithms in suitable algebraic groups, or computing the prime factors of a product consisting of two "large" primes.

3. A description of an *efficient reduction*, which transforms *any* efficient adversary $\mathcal{A}$ successfully breaking the security according to the security definition into an efficient algorithm that solves the computational problem.

[2] See `https://amturing.acm.org/award_winners/goldwasser_8627889.cfm`.

Hence, breaking the security of the cryptographic scheme is at least as hard as solving the computational problem. Since we assume that no algorithm $\mathcal{R}$ efficiently solving the computational problem can exist (or can be efficiently found), the adversary $\mathcal{A}$ breaking the security according to the security definition cannot exist either. Naturally, this relation is only "meaningful" if the assumption that the computational problem is "hard" holds true.

In contrast to the historical approach of designing cryptographic schemes, the provable security paradigm allows for a derivation of precise security statements, the falsifiability of security claims, and the verifiability of protocol designs.

1.1 A Gap Between Theory and Practice

Even though the provable security paradigm provides a theoretically-sound strategy to rigorously argue statements of security, it still comes with a few caveats, which sometimes make the applicability of security statements difficult for real-world cryptographic schemes. In particular, this is due to the subsequent reasons.

Asymptotic Security Guarantees. The provable security paradigm is based on concepts from complexity theory where the efficiency of algorithms is defined *asymptotically*. That is, an algorithm $\mathcal{A}$ is called *efficient* (or *polynomial-time*), if there exists a polynomial p such that the computation $\mathcal{A}(x)$ terminates for all inputs $x \in \{0,1\}^*$ within at most $p(|x|)$ steps. When proving the security of a cryptographic primitive with respect to this notion of asymptotic security, we also only get asymptotic security guarantees, *i.e.*, that sufficiently large parameters exist for the scheme to be secure. How those parameters need to be chosen to achieve security in the real-world is not clear. This leads to the following general research question.

How can we instantiate real-world cryptographic schemes in a theoretically-sound manner?

In Chapter 2, we delve deeper into the intricacies of this foundational research question, provide an overview of previous works on this topic, and describe the contributions of this thesis to this field of research.

Idealized Security Models. Security definitions of cryptographic primitives are often idealized, *e.g.*, by making assumptions on the distribution of certain values. For instance, standard security definitions for encryption-based primitives usually define the attacker goal as an efficient adversary not being able to distinguish the encryptions of two (distinct) messages with *equal length*. While this approach allows to keep the security model simple and to develop a general and universal theory of secure encryption, it does not reflect the real world where messages are rarely of equal length. In general, if the security model does not accurately reflect the real world, the "real-world meaningfulness" of security proofs conducted in such a model declines. Consequently, cryptographic schemes in the real world are often not fully covered by the idealized security proofs. This observation establishes another general research question treated in this thesis.

How can we define security models such that we can rigorously analyze the real-world security guarantees of a cryptographic scheme?

In Chapter 3, we further explore this foundational research question, provide an overview of previous works on this topic, and describe how this thesis contributed to this field of research.

Outline of This Thesis. This thesis explores both foundational research questions and advances the state of the art with six publications at leading conferences on cryptography. Since all publications consider their own important research question, we will present each result in its own section, all containing a detailed motivation, a summary of results, and possible future research directions. We conclude the introduction of this thesis with a publication overview and a brief description of the author's contributions.

1.2 Publication Overview and Contributions

This thesis is based on the following six results, which all have been published in renowned conference proceedings with peer-review. The collaborative nature of these publications is reflected in the equal contributions of all authors, since substantial portions of each publication evolved through extensive mutual discussions. We remark that, in the field of cryptography, authors of a paper

are usually listed alphabetically, not allowing any conclusion on who were the leading authors of a work. However, to provide transparency, we emphasize the specific areas in which the author of this thesis has mainly worked on for each publication.

Digital Signatures with Memory-Tight Security in the Multi-Challenge Setting.

[DGJL21a] Denis Diemert, Kai Gellert, Tibor Jager, and Lin Lyu. Digital signatures with memory-tight security in the multi-challenge setting. In Mehdi Tibouchi and Huaxiong Wang, editors, *ASIACRYPT 2021, Part IV*, volume 13093 of *LNCS*, pages 403–433, Singapore, December 6–10, 2021. Springer, Heidelberg, Germany. doi:10.1007/978-3-030-92068-5_14.

The author of this thesis worked on formalizing canonical reductions for digital signature schemes with memory-tight security, the proof of the main theorems [DGJL21a, Theorems 10, 14], and the applications section.

The result has been published at the *27th Annual International Conference on the Theory and Applications of Cryptology and Information Security – ASIACRYPT 2021*, one of the world's flagship conferences on cryptography.

Symmetric Key Exchange with Full Forward Security and Robust Synchronization.

[BDdK+21] Colin Boyd, Gareth T. Davies, Bor de Kock, Kai Gellert, Tibor Jager, and Lise Millerjord. Symmetric key exchange with full forward security and robust synchronization. In Mehdi Tibouchi and Huaxiong Wang, editors, *ASIACRYPT 2021, Part IV*, volume 13093 of *LNCS*, pages 681–710, Singapore, December 6–10, 2021. Springer, Heidelberg, Germany. doi:10.1007/978-3-030-92068-5_23.

The author of this thesis has contributed ideas to the identification of the "synchronization robustness" property, to the design of all protocols, especially the idea to use puncturable pseudorandom functions for protocols with non-linear key evolution, and has worked on the security proofs for the protocols with non-linear key evolution [BDdK+21, Theorems 18,22,23,24].

The result has been published at the *27th Annual International Conference on the Theory and Applications of Cryptology and Information Security – ASIACRYPT 2021*, one of the world's flagship conferences on cryptography.

More Efficient Digital Signatures with Tight Multi-User Security.

[DGJL21b] Denis Diemert, Kai Gellert, Tibor Jager, and Lin Lyu. More efficient digital signatures with tight multi-user security. In Juan Garay, editor, *PKC 2021, Part II*, volume 12711 of *LNCS*, pages 1–31, Virtual Event, May 10–13, 2021. Springer, Heidelberg, Germany. doi:10.1007/978-3-030-75248-4_1.

The author of this thesis worked on the construction, the proof of the main theorem [DGJL21b, Theorem 9], and possible instantiations of the construction.

The result has been published at the *24th International Conference on Practice and Theory of Public-Key Cryptography – PKC 2021*, the world's leading conference specialized on public-key cryptography.

On Fingerprinting Attacks and Length-Hiding Encryption.

[GJLN22] Kai Gellert, Tibor Jager, Lin Lyu, and Tom Neuschulten. On fingerprinting attacks and length-hiding encryption. In Steven D. Galbraith, editor, *CT-RSA 2022*, volume 13161 of *LNCS*, pages 345–369, Virtual Event, March 1–2, 2022. Springer, Heidelberg, Germany. doi:10.1007/978-3-030-95312-6_15.

The author of this thesis worked on the analysis of real-world message distributions and on the validation of the pencil-and-paper analysis [GJLN22, Sections 3,4].

The result has been published at the *RSA Conference, Cryptographers' Track – CT-RSA 2022*, the world's leading conference at the intersection of industry and academia.

On Optimal Tightness for Key Exchange with Full Forward Secrecy via Key Confirmation.

[GGJJ23] Kai Gellert, Kristian Gjøsteen, Håkon Jacobsen, and Tibor Jager. On optimal tightness for key exchange with full forward secrecy via key confirmation. In Helena Handschuh and Anna Lysyanskaya, editors, *Advances in Cryptology – CRYPTO 2023*, pages 297–329, Cham, 2023. Springer Nature Switzerland

The author of this thesis worked on the discovery of the flaw in [CCG+19] and the main meta reduction theorem [GGJJ23, Theorem 6.4].

The result has been published at the *43rd Annual International Cryptology Conference – CRYPTO 2023*, one of the world's flagship conferences on cryptography.

Security Analysis of the WhatsApp End-to-End Encrypted Backup Protocol.

[DFG+23] Gareth T. Davies, Sebastian Faller, Kai Gellert, Tobias Handirk, Julia Hesse, Máté Horváth, and Tibor Jager. Security analysis of the whatsapp end-to-end encrypted backup protocol. In Helena Handschuh and Anna Lysyanskaya, editors, *Advances in Cryptology – CRYPTO 2023*, pages 330–361, Cham, 2023. Springer Nature Switzerland

The author of this thesis has worked on the full protocol description, the formalization of the ideal functionality for password-protected key retrieval, and the formal security analysis of the protocol (especially the corrupt client cases of [DFG+23, Theorem 1]).

The result has been published at the *43rd Annual International Cryptology Conference – CRYPTO 2023*, one of the world's flagship conferences on cryptography.

1.3 Editorial Remarks

We remark that the target audience of this thesis are readers with a background in computer science and modern cryptography on the level equivalent of an advanced undergraduate student in computer science. Due to its cumulative character, this thesis focuses on providing intuitions on how the gap between ideal models and real-world requirements is further closed. For a more technical approach, we refer to the original publications.

Gender Neutrality. In support of gender neutrality, we use the generic pronouns "they/them" when referring to users of cryptographic schemes and "it/its" when referring to algorithms or devices. Note that we sometimes give algorithms human-sounding names such as "adversary," but use neuter pronouns to highlight their algorithmic, non-human nature.

Notation. We denote the security parameter as λ. For any value $n \in \mathbb{N}$, let 1^n be the unary representation of n, and let $[n] = \{1, \ldots, n\}$ be the set of integers between 1 and n. Furthermore, let $|x|$ denote the length of a bit string x, while $|\mathcal{S}|$ denotes the size of a set $\mathcal{S}$. For bit strings a, b, we define $a \parallel b$ to be the concatenation of a and b. We write $x \xleftarrow{\$} \mathcal{S}$ to indicate that the element x is chosen uniformly at random from set $\mathcal{S}$. For a probabilistic polynomial-time algorithm $\mathcal{A}$, we define $y \xleftarrow{\$} \mathcal{A}(a_1, \ldots, a_n)$ as the execution of $\mathcal{A}$ with fresh random coins on the inputs $a_1, \ldots, a_n$, and assigning the output to y.

2 Theoretically-Sound Instantiation

In modern cryptography it is common to not only propose a cryptographic scheme, but to also provide a rigorous security proof alongside the scheme's description. The security proof is meant to establish confidence in the security guarantees the scheme provides. Such a security proof usually is a reduction (in a complexity theoretic sense), where the difficulty of breaking the scheme is related to the difficulty of breaking a hardness assumption. That is, the security proof aims to transform any efficient black-box attacker against the security of the scheme into an efficient algorithm solving an assumed-to-be-hard problem. However, if the hardness assumptions holds, *i.e.*, if it is impossible to efficiently solve the underlying problem, then the black-box attacker cannot exist and the scheme must be secure. This conceptual approach is illustrated in Figure 2.1.

Concrete Security. We can measure the *quality* of a cryptographic reduction by setting the running time and the success probability of both the black-box attacker and the reduction in relation. To this end, let $t_{\mathcal{A}}$ and $t_{\mathcal{R}}$ be the running times of an adversary $\mathcal{A}$ and a reduction $\mathcal{R}$, and let $\epsilon_{\mathcal{A}}$ and $\epsilon_{\mathcal{R}}$ be their respective success probabilities. We define the *security loss* ℓ as the smallest ℓ, such that

$$\ell \cdot \frac{\epsilon_{\mathcal{R}}}{t_{\mathcal{R}}} \geq \frac{\epsilon_{\mathcal{A}}}{t_{\mathcal{A}}},$$

and call the fractions $\epsilon_{\mathcal{A}}/t_{\mathcal{A}}$ and $\epsilon_{\mathcal{R}}/t_{\mathcal{R}}$ the *work factors* of $\mathcal{A}$ and $\mathcal{R}$, respectively. This is the standard approach to measure concrete security, which was originally defined by Bellare and Ristenpart in 2009 [BR09].

If we view cryptography from a classical complexity theory perspective, it would be sufficient to show that the security loss ℓ is polynomially bounded in some security parameter λ. Note that this asymptotic approach only implies that sufficiently large parameter sizes *exist* for the scheme to be secure, but it does not offer any insight on how they are chosen. That is, if we want to measure the *concrete* security a scheme provides, we need to make the bounds

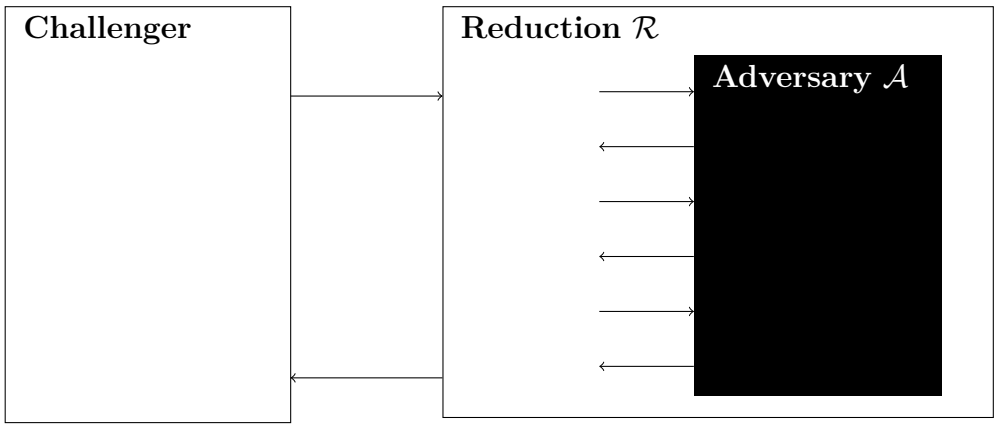

Figure 2.1: Conceptual approach of a security proof by reduction, where a black-box adversary $\mathcal{A}$ is transformed into a reduction $\mathcal{R}$ that aims to solve some hard problem. The challenger refers to an algorithmic representation of a hard problem.

more explicit.

If we the work factor $\epsilon_{\mathcal{A}}/t_{\mathcal{A}}$ is an upper bound for all adversaries $\mathcal{A}$, then the primitive offers security equivalent to $-\log(\epsilon_{\mathcal{A}}/t_{\mathcal{A}})$ bits. We usually desire a security level equivalent to at least 128 bits (often called "128-bit security"), which intuitively corresponds to the difficulty of correctly guessing a uniformly random 128-bit key. Given a reduction with work factor $\epsilon_{\mathcal{R}}/t_{\mathcal{R}}$, we can now bound

$$-\log(\epsilon_{\mathcal{A}}/t_{\mathcal{A}}) \geq -\log(\ell \cdot \epsilon_{\mathcal{R}}/t_{\mathcal{R}}) \geq 128.$$

Note that this bound depends on the security loss ℓ and the work factor $\epsilon_{\mathcal{R}}/t_{\mathcal{R}}$. That is, if we want to achieve the desired security equivalent to at least 128 bits, we need to ensure that the loss ℓ and the work factor $\epsilon_{\mathcal{R}}/t_{\mathcal{R}}$ are small enough such that $-\log(\ell \cdot \epsilon_{\mathcal{R}}/t_{\mathcal{R}}) \geq 128$ holds. Consequently, a large loss ℓ would need to be compensated by a sufficiently small work factor $\epsilon_{\mathcal{R}}/t_{\mathcal{R}}$. This can be achieved by instantiating the scheme with "large" parameter sizes, but this has an undesirable impact on the computational efficiency of the cryptographic primitive. Hence, we would like to have a so-called *tight* security proof, where the loss ℓ is small enough, or even constant in the security parameter, such that ideally no compensation in parameter sizes is necessary.

The Impact of a Security Loss. We illustrate the impact of a security loss via an example considering the famous Transport Layer Security (TLS) proto-

col [Res18], which is implemented in any modern web browser and the most widely used protocol to establish a secure connection over the Internet. While we do know that TLS 1.3 can indeed be proven secure in a tight manner [DJ21, DG21], many previous proofs of (candidate versions of) TLS 1.3 suffer from a loss at least *quadratic* in the number of sessions n_s [DFGS15, DFGS16, FG17, Gün20].

According to Statistia, a provider of market and consumer data, the number of worldwide Internet users in 2023 is around 5 billion[1] and the number of devices connected to the Internet of Things (IoT) is around 15 billion, with a strong upward trend[2]. Assuming more than 2^{15} sessions per user and device, we get a total number of sessions $n_s \geq 2^{15} \cdot 20 \cdot 10^9 \approx 2^{50}$, which yields a security loss of $\ell \approx 2^{2 \cdot 50} = 2^{100}$.

We can compensate this security loss by choosing the parameter sizes of cryptographic primitives used within the TLS protocol accordingly. Concretely, the elliptic curve group would need to be of order $\approx 2^{256+2 \cdot 100} = 2^{456}$ rather than of order $\approx 2^{256}$, and the Rivest–Shamir–Adleman (RSA) modulus (for the commonly used RSA-based signatures) would need to be over 10,000 bits long rather than the usually recommended 3,072 bits.[3] This increase of parameter sizes has a direct impact on the computational efficiency of the TLS protocol and is therefore undesirable.

We observe that the tightness of a security proof provides insight on how to instantiate cryptographic schemes in a theoretically-sound manner. It is hence an important field of research to investigate (i) if real-world cryptographic schemes can be proven tightly-secure, (ii) if certain classes of protocols are impossible to prove tightly-secure, and (iii) how impossibility results can potentially be evaded.

[1] See *Number of internet and social media users worldwide as of April 2023*, `https://www.statista.com/statistics/617136/digital-population-worldwide/`.

[2] See *Number of Internet of Things connected devices worldwide from 2019 to 2023, with forecasts from 2022 to 2030*, `https://www.statista.com/statistics/1183457/iot-connected-devices-worldwide/`.

[3] We remark that the compensation of parameter sizes needs to take attacks on the underlying algebraic structures into account as well. For example, the Pollard's Rho algorithm [Pol78] is applicable to elliptic curve groups and requires a *per se* doubling of the group order. Hence, we need a group of order $2^{2 \cdot (128+100)}$ instead of $2^{128+100}$. Similarly, the general number field sieve [Pom96] needs to be taken into account for RSA-based primitives.

Outline of This Chapter. We discuss how this thesis advances the state of the art in the domain of theoretically-sound instantiation of cryptographic schemes. We consider two fundamental cryptographic primitives, digital signature schemes and key exchange protocols. In Section 2.1, we discuss digital signatures with tight multi-user security, i.e., where the number of users does not impact parameter sizes. Next, in Section 2.2, we discuss memory-tight signatures, where the memory consumption of the adversary and the reduction are also taken into account. We conclude with Section 2.3, in which we investigate the optimal tightness of key exchange protocols with explicit authentication.

Each of subsequent sections motivates a fundamental research problem and details how this problem was solved. Additionally, we propose future research directions based on our results.

2.1 Signatures with Tight Multi-User Security

The commonly accepted standard security notion for digital signatures is the existential unforgeability under adaptive chosen-message attacks (EUF-CMA). Intuitively, this notion guarantees that an adversary can forge a fresh and valid signature only with negligible probability, even if it has access to a signing oracle providing it with valid signatures of its choice. Standard EUF-CMA security considers a *single-user setting*, where the adversary is given a single public key at the start of the security experiment, which it has to forge a signature for.

This notion can be strengthened to a multi-user setting with adaptive corruptions, so-called existential unforgeability under adaptive chosen-message attacks in the multi-user setting with adaptive corruptions ($\mathsf{MU\text{-}EUF\text{-}CMA}^{\mathsf{corr}}$) security. Here, the adversary receives not only a single public key at the start of the experiment, but public keys of multiple users. The users may get corrupted by the adversary, allowing it to obtain their respective secret keys. The attacker goal of the adversary is to forge a fresh signature that is valid with respect to an uncorrupted public key.

It is easy to show that EUF-CMA security implies $\mathsf{MU\text{-}EUF\text{-}CMA}^{\mathsf{corr}}$ security via a straightforward guessing argument. That is, given an adversary $\mathcal{A}$ against the $\mathsf{MU\text{-}EUF\text{-}CMA}^{\mathsf{corr}}$ security of some signature scheme, we construct a reduction $\mathcal{R}$ as follows. The reduction $\mathcal{R}$ guesses for which user i^* the adversary $\mathcal{A}$ will forge a signature and embeds the public key of its EUF-CMA challenger as

the i^*th public key, while generating and simulating all other users by itself. Note that if the guess by the reduction $\mathcal{R}$ is correct, it will perfectly simulate the adversary $\mathcal{A}$. Since the guess is independent of the behavior of the adversary $\mathcal{A}$, the reduction wins the security game against the EUF-CMA challenger with advantage $\epsilon_{\mathcal{A}}/q_N$, where q_N is the number of users and $\epsilon_{\mathcal{A}}$ is the success probability of $\mathcal{A}$. This yields a security loss linear in the number of users q_N and the reduction is hence non-tight.

The Need for Tight MU-EUF-CMA$^{\text{corr}}$-Secure Signatures. There are different reasons why the construction of digital signature schemes with tight multi-user security is interesting. First of all, digital signatures are often used as a building block in more complex protocols, with key exchange protocols being the prime example with the highest real-world relevance. Security of key exchange protocols is often argued in Bellare–Rogaway-style [BR94] or Canetti–Krawczyk-style [CK01] models, where the adversary may adaptively compromise the secret key material of users. Since this key material often comprises of the secret signing key of a signature scheme, these adaptive corruptions directly translate to adaptive corruptions of the signature scheme in a multi-user setting. Note that real-world key exchange protocols are executed by a vast number of users and hence it is desirable to avoid losses in the number of users whenever possible. The construction of MU-EUF-CMA$^{\text{corr}}$-secure signatures helps in this endeavor by offering a convenient-to-use interface for the reduction, when the security of the key exchange protocol is reduced to the security of the digital signature scheme.

Apart from this practically motivated reason, we deem the question how to construct such digital signature schemes as a general and foundational research question in cryptography, which helps us to better understand the general theory of digital signature schemes with strong security guarantees.

The Difficulty of Constructing Tightly MU-EUF-CMA$^{\text{corr}}$-Secure Signatures. We have already discussed that EUF-CMA security implies MU-EUF-CMA$^{\text{corr}}$ security via a straightforward guessing argument. This approach incurs a loss linear in the number of users and is hence a non-tight proving strategy. Any tight reduction for such an implication needs to avoid this kind of guessing argument. We observe that any such reduction must

1. know the secret keys of *all* users, such that it can correctly answer a corruption query for any user, but without guessing which users will not get compromised;

2. needs to extract a solution to some assumed-to-be-hard computational problem, while knowing the secret key for this instance in order to correctly answer signing queries.

On first sight, these observations seem to contradict each other and indeed one can formally prove that such a reduction does not exist under non-interactive hardness assumptions[4] and certain additional properties of the digital signature scheme [BJLS16]. While the impossibility result covers a large class of widely-used schemes, there also exist some constructions that manage to avoid the impossibility result.

Known Constructions with Tight MU-EUF-CMA$^{\text{corr}}$ Security. There are only few known constructions of digital signature schemes with tight MU-EUF-CMA$^{\text{corr}}$ security under non-interactive hardness assumptions. One scheme is given by Bader *et al.* and achieves a constant security loss, but the size of the signatures are linear in the security parameter. They also describe another scheme with constant-size signatures, but with a linear security loss in the security parameter instead. Another scheme was presented by Gjøsteen and Jager, which has a constant security loss in the programmable random oracle model [GJ18]. This motivates the leading research question of this section:

Is it possible to construct a digital signature scheme with tight security in the multi-user setting, which is more practical than previous schemes?

2.1.1 Results

This section summarizes our results described in [DGJL21b]. We present a new generic construction of strongly MU-EUF-CMA$^{\text{corr}}$-secure digital signature schemes based on lossy identification schemes (LIDs) [AFLT12, AFLT16] and

[4] Using an interactive hardness assumption "shifts" the difficulty of constructing such schemes into assumption. This makes the assumption stronger and a tight security proof trivial. Since we want to prove security under well-known, weak assumptions, relying on interactive hardness assumptions is undesirable.

sequential OR-proofs [AOS02, FHJ20]. Our scheme is described in [DGJL21b, Section 4] and has the following properties:

- It is the *first generic construction* of $\mathsf{MU\text{-}EUF\text{-}CMA}^{\mathsf{corr}}$-secure digital signatures. It can be instantiated from any LID with suitable properties, which are met by, *e.g.*, the LID based on the decisional Diffie–Hellman (DDH) problem [CEv88], or the scheme based on the ϕ-hiding problem [ABP13]. A direct instantiation via the LIDs based on, *e.g.*, the decisional short discrete logarithm problem, the ring learning with errors problem, or the subset sum problem (all due to Abdalla *et al.* [AFLT12, AFLT16]) does not directly yield a tightly-secure digital signature scheme. This is due to the requirement of a tight "multi-instance security" of the underlying assumption. While this is met by the self-reducibility property of the LID based on the DDH problem, it is not obvious how this property can be achieved for hardness problems based on lattices or subset sum.

- It is the first construction of $\mathsf{MU\text{-}EUF\text{-}CMA}^{\mathsf{corr}}$-secure digital signatures with *strong* existential unforgeability if the underlying LID supports the uniqueness property in the sense of [AFLT12, AFLT16].

- When our construction is instantiated via the LID based on the DDH problem, a signature only consists of three elements of $\mathbb{Z}_q$ and the public key consists of four group elements, where q is the order of the underlying algebraic group. This makes it the most efficient digital signature scheme with tight $\mathsf{MU\text{-}EUF\text{-}CMA}^{\mathsf{corr}}$ security to this date.

2.1.2 Possible Future Research Directions

The digital signature schemes discussed in this section can be instantiated from any LID with suitable properties. Naturally, it would be interesting to find novel construction of LIDs, which directly give rise to new tightly-secure signature schemes. This leads us to the first broad research question.

Research Question 1. *How can we construct new LIDs, which give rise to new signature schemes with tight multi-user security?*

On a more specific level, it would be interesting to find a suitable LID, which is based on post-quantum-secure assumptions, such as lattice-based assumptions. Even though LIDs based on lattices exist (*e.g.*, the schemes by Abdalla

et al. [AFLT12, AFLT16]), they do not support a tight "multi-instance security," such as the self-reducibility property of DDH-based LID. Thus, it is not obvious how we can use lattice-based LIDs to construct digital signatures with tight multi-user security.

Research Question 2. *How can we construct lattice-based digital signature schemes with tight multi-user security?*

The last research question of this section will be motivated from a practical perspective. All digital signature schemes, which are currently standardized for use in widely-deployed key exchange protocols, such as TLS 1.2 and TLS 1.3, do not achieve tight multi-user security, or tightness at all. This makes a theoretically-sound instantiation of the digital signature schemes without a compensation of parameter sizes very difficult. Since our signature scheme is competitive when considering computational efficiency and signature sizes, it would be interesting to see how well it would perform when implemented in real-world protocols. Hence, we propose the following research question.

Research Question 3. *How well do our digital signature schemes perform when implemented in real-world key exchange protocols, such as, e.g., TLS 1.3?*

2.2 Signatures with Memory-Tight Multi-Challenge Security

Tightness with respect to the work factor (i.e., the ratio of running time and success probability) has been considered the standard notion for tightness since its introduction in 2009. In 2017, Auerbach *et al.* [ACFK17] showed that the *memory consumption* of a reduction is another relevant dimension, which is often overlooked, especially if the security is reduced to a hardness assumption, where the efficiency of a solving algorithm depends on how much memory is available. We call such problems *memory-sensitive*.

Memory-Sensitive Problems. We know that several memory-sensitive problems with relevance to cryptography exist. Well-known examples include k-way collision resistance (*e.g.*, due to the Joux–Lucks algorithm [JL09]), the learning parity with noise problem (*e.g.*, due to the Blum–Kalai–Wassermann

algorithm [BKW03]), the shortest vector problem (*e.g.*, due to the Herold–Kirshanova algorithm [HK17]), the discrete logarithm problem in prime fields or the factoring problem (both due to the number field sieve, see [LLMP93]). For a more detailed discussion on these problems, we refer the reader to [ACFK17, Section 6].

Memory Tightness. This gap was addressed by Auerbach *et al.* [ACFK17] who were the first to formally define *memory tightness*, a tightness notion that takes the consumed memory of a reduction into account. Their work presents techniques how memory tightness can be achieved for certain reductions (*e.g.*, by implementing a random oracle via a pseudorandom function rather than via lazy sampling), but also gives impossibility results that cover large classes of natural reductions. Since then, many follow-up works on this topic have emerged, considering the memory tightness of digital signature schemes [WMHT18, GGJT22, Xag23], public-key encryption schemes [GT20, GGJT22, JK22], key encapsulation mechanisms [Bha20], and symmetric key cryptography [WMHT18, GJT20, GGJT22].

Memory-Tight Signatures. The standard notion of EUF-CMA security for signatures gives an adversary access to a signing oracle. That is, it may adaptively query the oracle for arbitrary messages m and expects a valid signature σ in return. The challenger usually stores all queried messages m in a list $\mathcal{Q}$. Eventually, the adversary outputs a candidate forgery (m^*, σ^*) and wins if it is indeed valid and if $m^* \notin \mathcal{Q}$. This is called the *single-challenge setting*, since the adversary gets only one attempt to output a forgery. This notion can be strengthened by allowing the adversary to output multiple candidate forgeries. It would then win if at least one of those candidate forgeries is indeed a valid forgery. This security variant is in the stronger *multi-challenge setting*.

The Difficulty of Constructing Memory-Tight Signatures in the Multi-Challenge Setting. The single-challenge and multi-challenge notions are tightly equivalent if work-factor tightness is considered. Given an adversary in the multi-challenge setting, a reduction can store all message-signature pairs it has given the adversary via the signing oracle. Whenever the adversary outputs a forgery candidate, the candidate is checked against the message-signature pairs.

If it is not contained, the candidate is a valid forgery in the single-setting challenge as well. Note that this reduction is not memory-tight, since it requires storing a message-signature pair for each request to the signing oracle, incurring a loss linear in the number of signing queries.

Auerbach *et al.* showed that it is impossible to prove a memory-tight equivalence of these notions for a large class of natural reductions [ACFK17]. Their result was subsequently revisited and extended by Wang *et al.* [WMHT18]. We note that Auerbach *et al.* give a memory-tight security proof for the RSA full-domain hash signature scheme under the RSA assumption. However, their reduction only achieves memory tightness in the single-challenge setting. Furthermore, their security proof is also not work-factor-tight, since the lower bounds of previous impossibility results still apply, making a loss linear in the number of signing queries impossible to avoid [Cor02, KK12, BJLS16, KK18].

There are two major difficulties when constructing digital signature schemes with memory-tight security in the multi-challenge setting. First, the reduction cannot store a table of random oracle queries by the adversary. While the table can in principle be replaced via a pseudorandom function [ACFK17], one needs to be careful if the random oracle requires "programming" by the reduction. Second, the reduction needs to avoid storing a list of message-signature tuples that are given to the adversary via the signing query. This motivates the leading research question of this section:

> *Is it possible to evade existing impossibility results and construct a digital signature scheme that achieves full tightness in a multi-challenge setting?*

2.2.1 Results

This section summarizes our results described in [DGJL21a].

- We present a sequence of transformations that gives rise to the first digital signature schemes that achieve tightness in all three dimensions, *i.e.*, in running time, success probability, and memory consumption. Our scheme exploits a new notion of *canonical reductions*, where we enforce a certain behavior onto the reduction, which is met by many well-known security proofs [DGJL21a, Sections 3,4]. This "non-black-box perspective" allows us to avoid previous impossibility results by Auerbach *et al.* [ACFK17] and

Wang *et al.* [WMHT18], showing that such a construction is impossible for a large class of black-box reductions.

- We demonstrate that our approach works for the generic construction of signature schemes via LIDs by Abdalla *et al.* [AFLT12, AFLT16], for RSA full-domain hash signatures [BR96], and for Boneh–Lynn–Shacham signatures [BLS01]. We further apply the Katz–Wang technique [KW03] and sign the message with an additional random bit, avoiding the linear loss in the number of signing queries (which is inherent for a large class of digital signature schemes with unique signatures [Cor02]) for the latter two schemes [DGJL21a, Section 5].

2.2.2 Possible Future Research Directions

In this section, we discussed digital signature schemes with multi-challenge security that achieve tightness in the dimensions running time, success probability, and memory consumption. A first interesting research question is to investigate, whether we can achieve fully tight digital signature schemes in a *multi-user* setting. This notion is especially interesting, since schemes with tight multi-user security are often used as building blocks in more complex protocols. The result of this research could either be constructive, by achieving such a security notion, or an impossibility result, showing that a certain class of protocols cannot achieve such a security notion.

Research Question 4. *Is it (im)possible to construct fully tight digital signatures with multi-user security?*

Building upon the above research question, we can use our tight signature schemes as building block in more involved protocols and investigate whether it is still possible to achieve full tightness. An interesting candidate could be key exchange protocols, since they often rely on digital signature schemes as building blocks.

When considering key exchange protocols, we will face another new challenge apart from the multi-user security issue described above. There exist several impossibility results, which show that a loss linear in the number of users is impossible to avoid for a large class of protocols [BJLS16, CCG+19, GGJJ23]. Since this loss cannot be avoided in the "classical" work-factor tightness, an

additional consideration of memory consumption as dimension will not help to avoid this loss. However, it is not clear if we can construct a key exchange protocol with optimal work-factor tightness and memory tightness, or if there is yet another unavoidable loss when additionally considering the memory consumption of a reduction.

Research Question 5. *Is it (im)possible to construct a key exchange protocol with an optimal work-factor tightness and memory tightness?*

2.3 Key Exchange with Optimal Tightness

Authenticated key exchange (AKE) protocols allow two parties two establish a shared secret over an insecure channel such as the Internet. Such protocols are often proven to provide implicit authentication, meaning that the intended partner can derive that same shared secret, but not guaranteeing that they have actually participated in the protocol execution. If a key exchange protocol only provides implicit authentication, an adversary can impersonate a user, such that the intended partner believes it is establishing a key with the user. However, the adversary still cannot distinguish the established session keys from random, preserving the standard security notion of indistinguishable keys.

It is a natural question how an implicitly authenticated protocol can be tightly upgraded into one that provides explicit authentication, guaranteeing that both parties have indeed participated in the protocol run. A well-known approach is to add key confirmation messages derived from the implicitly-authenticated session key to the protocol[5], which has been studied in several previous works [BPR00, Kra05, Yan13, FGSW16, CCG+19, dSGFW20].

Key Confirmation and the Preservation of Tightness. In 2019, Cohn-Gordon *et al.* have presented nearly-tight reductions to implicitly authenticated Diffie–Hellman protocols. Their reductions have a security loss, which is linear in the number of users and cannot be avoided for a large class of natural reductions [CCG+19]. Furthermore, they argued that their protocol can be tightly

[5] We note that key confirmation messages can also be used to upgrade a weakly forward-secure protocol to a fully forward-secure protocol, but remark that implicit/explicit authentication and weak/full forward security are separate notions. In fact, protocols with explicit authentication but no forward security [BR94] and protocols with forward security but only implicit authentication [BN11, CF15] exist.

upgraded to provide explicit authentication, thus preserving the optimal tightness of the underlying protocol [CCG+19, Theorem 6]. This claim does not hold true[6], which motivates the leading research question of this section:

Is it possible to tightly upgrade an implicitly AKE protocol to achieve explicit authentication, and can the result of Cohn-Gordon et al. be restored?

2.3.1 Results

This section summarizes our results described in [GGJJ23].

- We present a meta reduction [Cor02] showing that no security proof for adding key confirmation to a weakly forward-secure key exchange protocol can avoid a loss linear in the number of protocol users [GGJJ23, Section 6]. Our impossibility result holds for a large class of protocols, including, *e.g.*, the generic compiler by Cohn-Gordon *et al.* [CCG+19], and the MAC-based approach transforming the HMQV protocol into the HMQV-C protocol [Kra05].
- We present an alternative proof strategy to the one used in [CCG+19]. That is, we weaken the security notion of the underlying key exchange protocol such that (i) the underlying protocol can be proven secure in a tight manner, and (ii) the compiler adding key confirmation messages has a security loss linear in the number of users. This proving strategy is optimal, since any such key exchange protocol must lose a factor that is at least linear in the number of users [GGJJ23, Section 4].
- We give a tight proof of the underlying protocol used by Cohn-Gordon *et al.* [CCG+19] under our weaker notion of security, overall restoring their main result [GGJJ23, Section 5].

2.3.2 Possible Future Research Directions

In this section we have discussed an impossibility result, which shows that a large class of compilers upgrading an implicitly-authenticated protocol into an

[6] We decided against an explanation of the actual flaw, since it is quite subtle and requires some in-depth knowledge of key exchange security models, which are known for their complexity. We hence refer the interested reader to [GGJJ23, Section 1] for a more technical explanation of the flaw.

explicitly-authenticated one inherently cause a loss linear in the number of users. The impossibility result holds if certain criteria of the underlying protocol are met. These include, *e.g.*, that messages of the underlying protocol are independent of the secret long-term keys, that the compiler adds key confirmation messages derived from the output key of the underlying protocol and that the underlying protocol has unique and efficiently verifiable secret keys. A natural question is whether we can avoid the impossibility result by constructing a protocol, which avoids at least one of these requirements. If this seems difficult, one could also try to extend the impossibility result to capture more general requirements.

Research Question 6. *Is it possible to avoid the impossibility result? If not, can the impossibility result be extended to an even more general class of protocols?*

Another potential research question could further widen the scope of AKE protocols with optimal tightness. For example, it could be interesting to explore tightness notions in password-authenticated key exchange [BM92]. While these protocols are usually proven in the universal composability framework [Can00], which is due to the composition theorem highly non-tight, there also exists a game-based security model by Bellare, Pointcheval and Rogaway [BPR00] that could serve as a foundation to study the tightness of password-authenticated key protocols. Hence, we propose the following research question.

Research Question 7. *Can password-authenticated key exchange protocols achieve tight security, or can we show under which conditions it is impossible to achieve tightness?*

3 Cryptography for the Real World

This chapter is about cryptographic schemes with a high real-world relevance. First, we explore symmetric key exchange, a variant of AKE, which only deploys symmetric cryptographic primitives. Symmetric key exchange is especially relevant for low-performance devices that do not have the computational capacity to use public-key-based cryptography. In Section 3.1, we explore what security guarantees we expect from such protocols and how they can be achieved in practice.

Second, we analyze the security guarantees of the message backup protocol that was recently deployed by WhatsApp in late 2021. Their protocol aims to provide strong security guarantees, requiring that even a malfeasant service provider must not get access to backups stored by users. We present the findings of our security analysis in Section 3.2.

Last, we investigate secure symmetric encryption, which is often idealized in security models. That is, classical security models for secure encryption require that all messages are of equal length, which does not accurately model the real world. We extend a previous approach called *length-hiding encryption* and develop a novel methodology to quantify which security guarantees can be achieved for real-world message distributions. Our results can be found in Section 3.3.

3.1 Forward-Secure Symmetric Key Exchange

AKE protocols enable two parties to mutually authenticate and to derive a session key over an insecure channel such as the Internet. A subclass of such protocols is based on pre-shared long-term symmetric keys, which were previously established, *e.g.*, via an out-of-band communication. A prominent example of such a protocol is the TLS 1.3 protocol in the pre-shared key mode [Res18]. However, this mode does not only rely on symmetric cryptographic primitives,

but also uses the public-key-based Diffie–Hellman key exchange for key derivation. Since public-key techniques are several orders of magnitude more expensive than symmetric-key techniques, pre-shared key protocols exclusively relying on symmetric-key techniques can be much more efficient than classical AKE protocols. We call such protocols *symmetric AKE protocols*.

The Need for Symmetric AKE. Symmetric AKE protocols are desirable for low-performance devices (*e.g.*, battery-powered wireless IoT devices), for which every computation and every transmitted bit has a negative impact on its battery life. Especially in an industrial context, where it is often easier to establish pre-shared keys than to rely on public-key infrastructures. Furthermore, protocols that exclusively rely on symmetric-key primitives provide an easy hardening against quantum adversaries by adjusting key size accordingly.

Forward Security in Symmetric AKE Protocols. Forward security is considered a standard goal of modern key exchange. We call a protocol *forward-secure* if the session keys of *past* sessions remain secure, even if the secret long-term key material gets compromised [BG20]. Note that forward security can only be achieved if session keys cannot be efficiently derived from the long-term key material. This property is often easy to achieve, when public-key techniques are used, *e.g.*, by introducing ephemeral keys for key establishment and long-term keys for authentication.

The only technique exclusively relying on symmetric-key techniques is often called the *derive-then-evolve approach*. That is, a session key is derived from the long-term key material, and then the long-term key is evolved. If the evolution of the key cannot be efficiently reversed, then compromise of the long-term key material also does not expose prior session keys. There are two common variants to this technique:

- *Synchronized Key Evolution.* In this case, long-term keys are evolved in epochs, *e.g.*, once per day. Note that this epoch-based evolution cannot achieve "full" forward security, but only a weaker "delayed" notion. Prior session keys are only secure if the long-term key used to derive them has evolved. This approach also requires both parties to have synchronized clocks, which may be difficult to achieve in a setting for constrained

devices. Hence, this approach seems impractical for the setting where symmetric AKE protocols are needed.

- *Triggered Key Evolution.* In this case, long-term keys are evolved during the protocol execution. Note that this ensures "full" forward security, since keys are always evolved after a protocol run has successfully concluded. However, triggered key evolution is much more difficult to realize on a technical level since both parties must remain "in sync," especially if transmitted messages are lost, or if an active attack on the communication is launched.

Concurrency and Key-Evolving Protocols. A standard correctness requirement for AKE protocols is that the protocol supports running concurrent sessions. This is reflected in the common Bellare–Rogaway [BR94] and Canetti–Kraczyk [CK01] security models, where the adversary may arbitrarily interleave concurrent protocol sessions. In the case of symmetric AKE, the main challenge is to preserve concurrent correctness while achieving "full" forward security as well.

Even in a setting where the adversary remains passive and all messages are transmitted reliably (*i.e.*, without message loss), is achieving concurrent correctness non-trivial. The main difficulty is that one party may advance its long-term key too early for another concurrent session to be completed. This issue does not appear in classical AKE protocols using public-key techniques since long-term keys are usually static and used for authentication only. A protocol achieving both "full" forward security and concurrent correctness does not yet exist, which raises the question if it is possible to construct such a protocol.

If we are now to consider a setting where messages may not be delivered, or where an adversary may actively interfere with the communication, the situation gets even more complicated. For example, an adversary may try to get two parties "out-of-sync," such that they cannot efficiently recover. It is an unsolved foundational problem to develop techniques, which ensure that such a "synchronization robustness" property is achieved in stateful key exchange protocols. This motivates the leading research question of this section:

How can we build efficient symmetric key exchange protocols based on pre-shared keys that achieve a meaningful security notion? Furthermore, is it

even possible to achieve synchronization robustness in stateful key exchange protocols?

3.1.1 Results

This section summarizes our results described in [BDdK+21].

- We describe a security model for forward-secure symmetric AKE capturing entity authentication (both unilateral and mutual), indistinguishability of established keys, and forward security [BDdK+21, Section 3]. Our model is based on the standard approach of modeling AKE by Bellare and Rogaway [BR94], adapted to the requirements of symmetric protocols with evolving keys.

- We identify a key property for symmetric key exchange, which we call *synchronization robustness* [BDdK+21, Section 3.3]. Intuitively, synchronization robustness guarantees that parties can efficiently re-synchronize their states in order to complete a successful protocol run, even if an adversary interferes with their communication. This property is trivially achieved by traditional AKE protocols with fixed long-term keys, but requires more care for symmetric key exchange with evolving keys.

- We define the notion of linear key evolution, which is based on the classical *derive-then-evolve* approach. We argue that protocols with linear key evolution can only achieve a weak form of synchronization robustness, where the "target" protocol session must be executed without adversarial influence. We present three different protocols based on linear key evolution. Two of them are lightweight and achieve unilateral authentication (resp. mutual authentication) with a communication complexity of only one (resp. two) message authentication codes and one (resp. two) counter values. The third protocol additionally achieves a "bounded gap" property, where both parties can always re-synchronize in one key evolution step [BDdK+21, Section 4].

- We define the notion of *non-linear key evolution* and present two such protocols based on puncturable pseudorandom functions (PPRFs), which achieve full synchronization robustness. Both protocols are lightweight,

since PPRFs can be instantiated from cryptographic hash functions, and achieve unilateral authentication (resp. mutual authentication) with a communication complexity of only one (resp. two) message authentication codes and one counter value [BDdK+21, Section 5].

3.1.2 Possible Future Research Directions

It remains to verify that our solutions indeed perform well in practice. To this end, it would be interesting to implement our five protocols and analyze how well they perform on low-performance devices. Specifically, the protocols based on non-linear key evolution pose a challenge, since it involves an implementation of a PPRF. A suitable instantiation could be the PPRF based on the Goldwasser–Goldreich–Micali tree as it only employs hash function derivations, which can be evaluated efficiently, even on low-performance devices. An alternative construction could be the PPRF based on the strong RSA assumption from [AGJ19, AGJ21]. However, this construction relies on modular exponentiation, which is expensive on low-performance devices, and only supports polynomially many evaluations.

Research Question 8. *How well do our protocols perform when implemented on low-performance devices?*

3.2 A Security Analysis of the WhatsApp Backup Protocol

WhatsApp is the most popular instant messenger application with over 2.7 billion unique users since 2020[1] and over 100 billion messages sent per day[2]. The messenger provides end-to-end encryption (E2EE), where only the sender and intended receiver should be able to read messages. Specifically, WhatsApp themselves should not be able to peak into messages sent by their users, preserving confidentiality of messages even if the WhatsApp service providers were to be compromised. Nowadays, E2EE is seen as a standard security goal for modern messengers and many security analyses on the messaging protocol

[1] https://www.statista.com/statistics/1306022/whatsapp-global-unique-users/
[2] https://techcrunch.com/2020/10/29/whatsapp-is-now-delivering-roughly-100-billion-messages-a-day/

deployed by WhatsApp or other messengers, such as Signal, have been conducted [BSJ+17, ACD19, JMM19, CPZ20, CCD+20, BFG+22, CJSV22].

Bypassing End-to-End Encryption with Backups. Secure messaging is not the only service WhatsApp offers its users. Additionally, users may create a backup of their messages, such that the messages can be recovered if, *e.g.*, their mobile device got lost or stolen. While this feature allows for a convenient recovery of messages, it also imposes new challenges to the security needs of a messenger. For example, if messages were to be stored in plain on the external storage servers (*e.g.*, on Google Drive or iCloud in the case of WhatsApp), the E2EE guarantee of the messaging protocol would be undermined. Encrypting the messages with a random symmetric key seems infeasible, since such keys would need to be stored on the user's mobile device and could get lost as well. Instead, a more sophisticated solution is necessary.

WhatsApp End-to-End Encrypted Backups. In late 2021, WhatsApp deployed a new backup protocol, with the aim to extend E2EE guarantees from messaging to backups as well [Wha21]. The backup protocol is password-based, *i.e.*, the user is not required to securely store a symmetric key but only needs to remember a password of its own choice. When a user loses their mobile device, they shall be able to restore their backup from their password, but any potentially malicious party (including the service provider) shall not be able to retrieve the backup without knowing the password. The backup protocol utilizes hardware security modules (HSMs) as a trusted building block. Intuitively, an HSM is a hardware device that can be programmed once and then "locked" such that the code cannot be retroactively changed. This approach guarantees that parts of the protocol remain trusted, even if the owner of the HSMs (here, WhatsApp) gets compromised. This motivates the leading research question of this section:

Which security guarantees do we formally expect from a password-protected key retrieval protocol and which of these guarantees are achieved by the WhatsApp backup protocol?

3.2.1 Results

This section summarizes our results described in [DFG+23].

- We formalize the security properties expected by a password-protected key retrieval protocol, where users may use a password to store and retrieve cryptographic key material with an untrusted storage provider [DFG+23, Section 4].
- We provide a full protocol description of the cryptographic core of the WhatsApp backup protocol [DFG+23, Section 3]. Our description is based on a whitepaper published by WhatsApp [Wha21], a public security assessment of the backup protocol conducted by the NCC Group [DLS21], and personal correspondence with the designers of the backup protocol [Lew23].
- We present the first security analysis of the WhatsApp backup protocol in the universal composability (UC) framework [Can01], which (due to its simulation-based nature) enables us to conveniently capture user-chosen passwords with low entropy. Our security analysis confirms several prior statements about the security guarantees of the WhatsApp backup protocol [DFG+23, Section 5].
- We identify a mechanism how a compromised server could get more than the claimed [DLS21] ten password guesses per backup. Concretely, we show that a corrupted server can get ten password guesses per *backup initialization* [DFG+23, Section 3.6].

All of our findings have been reported to WhatsApp and, even though the attack described above was never demonstrated in practice, its feasibility was acknowledged.

3.2.2 Possible Future Research Directions

We discuss potential future research directions. So far, we have only analyzed what security guarantees the cryptographic core of the WhatsApp backup protocol offers. There are multiple possible directions to widen the scope of this analysis. For example, we have not yet formally analyzed the protocol, which is used by the HSM to outsource storage to an untrusted entity. Likewise, it would be interesting to consider additional models of corruption and investigate if, *e.g.*, users can recover from adaptive corruptions when being compelled to reveal their passwords at border control.

Research Question 9. *Which additional security guarantees are achieved by the WhatsApp backup protocol if the scope of the security analysis is widened?*

Apart from widening the analysis of the current protocol, one could also think about further improvements. For example, the increased number of maximal password guesses is due to the client not being authenticated towards the HSM, but only towards the WhatsApp servers. That way, the WhatsApp server can impersonate the client towards an HSM, increasing its number of admissible password guesses. However, this issue seems difficult to fix in practice as it does not only require changes to the WhatsApp protocol, but also changes to the authentication infrastructure of the entire WhatsApp ecosystem. Another interesting question is how the efficiency of the protocol can be improved, *e.g.*, by carefully weakening used cryptographic building blocks such as the OPAQUE protocol. This may yield a much more efficient protocol without the sacrifice of too many security guarantees.

Research Question 10. *How can the WhatsApp backup protocol be improved?*

A more foundational research question would be to deepen our general understanding in secure password-protected key retrieval protocols. For example, it would be interesting to generically build such a protocol from password-authenticated key exchange (or from oblivious pseudorandom functions and AKEs). Likewise, it would be interesting to see whether some of the public-key primitives could be replaced with secret-key primitives, which would yield a much more efficient protocol.

Research Question 11. *How can we construct new password-protected key retrieval schemes?*

3.3 A New Perspective on Length-Hiding Encryption

The first formal security definition for public-key encryption was given by Goldwasser and Micali in 1984 [GM84]. Their notion of *indistinguishability of ciphertexts* has inspired many of today's standard definitions of secure encryption (*e.g.*, security against chosen-plaintext attacks [BDJR97] and security against chosen-ciphertext attacks [NY90, RS92]).

The attacker goal of such security models is often modeled as follows: The adversary chooses two messages m_0, m_1, which are sent to its challenger. The challenger then samples a random bit $b \xleftarrow{\$} \{0,1\}$ and encrypts $c \xleftarrow{\$} \mathsf{Enc}(k, m_b)$ with some random key k. The ciphertext c is sent to the adversary that now has to decide, which of the two messages m_0, m_1 was encrypted. This attacker goal is often idealized in that m_0 and m_1 need to be of *equal length*.

A General Theory of Secure Encryption. The idealization of the message length is widely accepted in theoretical cryptography. It has proven to be a solid and easy-to-use foundation to develop a general and universal theory of secure encryption. However, note that all schemes proven secure under this notion do not need to conceal the length of a message. This is often tolerated, since most widely-used schemes reveal an approximation of the message length. For example, stream ciphers often reveal the exact bit length of a message, while encryption via block ciphers reveals, how many blocks of length b a message is split into.

Attacks Based on the Message Length. In many applications, the message length already reveals sensitive information to an adversary. For example, a passive attacker monitoring encrypted network traffic can already deduce which webpages a user has visited [Hin02, LL06, GBKS12, MHJT14, WG16]. Likewise, language and phrases spoken can be identified by observing an encrypted voice-over-IP conversation [WBMM07, WBC+08]. In both examples, the adversary does not break the expected security of the used encryption scheme, but utilizes a side channel, the lengths of encrypted messages, instead. These attacks illustrate, that concealing the message length is of high real-world relevance.

Impossibility of Length Hiding in Cryptographic Theory. Tezcan and Vaudenay showed in 2011 that efficiently hiding the message length is impossible if arbitrary message distributions are considered [TV11]. That is, they used techniques from theoretical cryptography to show that, in order to hide a one bit difference of two messages in an asymptotic setting, a padding of size 2^λ is required, where λ is the security parameter. Note that a padding of such size is infeasible in practice, since we aim for security parameters $\lambda \geq 80$.

Length-Hiding Encryption. In 2011, Paterson, Ristenpart, and Shrimpton introduce the notion of length-hiding encryption (LHE) [PRS11] as an attempt to overcome the impossibility result by Tezcan and Vaudenay. In LHE, the encryption procedure gets an additional length-hiding parameter ℓ as input. This parameter is specified by the application and determines how a given message is padded before encryption. While [PRS11] shows that "secure" LHE can be achieved if the length-hiding parameter ℓ ensures that ciphertexts have equal sizes. Unfortunately, [PRS11] does not explain how ℓ can be chosen in practice, such that it provides meaningful security guarantees for real-world message distributions. In particular, they do not provide a methodology to quantify, which security guarantees are given for which parameter ℓ. This motivates the leading research question of this section:

How can we quantify the security of LHE in practice for real-world message distributions?

3.3.1 Results

This section summarizes the results of [GJLN22].

- We develop a new methodology, which concretely quantifies the effect of LHE on the security for a given application [GJLN22, Section 2]. Our security model uses the classical security model for secure encryption as a basis, but extends it to capture fingerprinting attacks based on observable ciphertext sizes.
- We demonstrate the feasibility of our approach by applying it to different types of fingerprinting attacks, covering simple webpage fingerprinting, Google search term fingerprinting, and DNS fingerprinting, all for their *specific* real-world message distributions [GJLN22, Section 3]. For each of these settings, we consider different block modes (*e.g.*, the cipher block chaining mode and the counter mode), and whether compression is enabled or disabled.
- We find that, even with small padding sizes with a bandwidth overhead of 2–5%, we can already increase the security against fingerprinting attacks significantly. Our results demonstrate, that LHE may serve as a valuable tool to better protect against fingerprinting attacks in practice.

3.3.2 Possible Future Research Directions

In this section, we presented a new perspective on LHE, which can been seen as the first stepping stone towards a meaningful application of LHE in practice. A possible next step could be to perform further experiments on public web servers with real-world message distributions. To this end, we could extend the widely-deployed OpenSSL library to allow appending TLS fragments for padding, if necessary. This approach could offer valuable insights on how well our approach to LHE works in practice.

Research Question 12. *How well does our approach to LHE work in practice, e.g., when implemented in real-world servers?*

Bibliography

[ABP13] Michel Abdalla, Fabrice Ben Hamouda, and David Pointcheval. Tighter reductions for forward-secure signature schemes. In Kaoru Kurosawa and Goichiro Hanaoka, editors, *PKC 2013*, volume 7778 of *LNCS*, pages 292–311, Nara, Japan, February 26 – March 1, 2013. Springer, Heidelberg, Germany. doi:10.1007/978-3-642-36362-7_19.

[ACD19] Joël Alwen, Sandro Coretti, and Yevgeniy Dodis. The double ratchet: Security notions, proofs, and modularization for the Signal protocol. In Yuval Ishai and Vincent Rijmen, editors, *EUROCRYPT 2019, Part I*, volume 11476 of *LNCS*, pages 129–158, Darmstadt, Germany, May 19–23, 2019. Springer, Heidelberg, Germany. doi:10.1007/978-3-030-17653-2_5.

[ACFK17] Benedikt Auerbach, David Cash, Manuel Fersch, and Eike Kiltz. Memory-tight reductions. In Jonathan Katz and Hovav Shacham, editors, *CRYPTO 2017, Part I*, volume 10401 of *LNCS*, pages 101–132, Santa Barbara, CA, USA, August 20–24, 2017. Springer, Heidelberg, Germany. doi:10.1007/978-3-319-63688-7_4.

[AFLT12] Michel Abdalla, Pierre-Alain Fouque, Vadim Lyubashevsky, and Mehdi Tibouchi. Tightly-secure signatures from lossy identification schemes. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 572–590, Cambridge, UK, April 15–19, 2012. Springer, Heidelberg, Germany. doi:10.1007/978-3-642-29011-4_34.

[AFLT16] Michel Abdalla, Pierre-Alain Fouque, Vadim Lyubashevsky, and Mehdi Tibouchi. Tightly secure signatures from lossy identification

schemes. *Journal of Cryptology*, 29(3):597–631, July 2016. doi:10.1007/s00145-015-9203-7.

[AGJ19] Nimrod Aviram, Kai Gellert, and Tibor Jager. Session resumption protocols and efficient forward security for TLS 1.3 0-RTT. In Yuval Ishai and Vincent Rijmen, editors, *EUROCRYPT 2019, Part II*, volume 11477 of *LNCS*, pages 117–150, Darmstadt, Germany, May 19–23, 2019. Springer, Heidelberg, Germany. doi:10.1007/978-3-030-17656-3_5.

[AGJ21] Nimrod Aviram, Kai Gellert, and Tibor Jager. Session resumption protocols and efficient forward security for TLS 1.3 0-RTT. *Journal of Cryptology*, 34(3):20, July 2021. doi:10.1007/s00145-021-09385-0.

[AOS02] Masayuki Abe, Miyako Ohkubo, and Koutarou Suzuki. 1-out-of-n signatures from a variety of keys. In Yuliang Zheng, editor, *ASIACRYPT 2002*, volume 2501 of *LNCS*, pages 415–432, Queenstown, New Zealand, December 1–5, 2002. Springer, Heidelberg, Germany. doi:10.1007/3-540-36178-2_26.

[BDdK+21] Colin Boyd, Gareth T. Davies, Bor de Kock, Kai Gellert, Tibor Jager, and Lise Millerjord. Symmetric key exchange with full forward security and robust synchronization. In Mehdi Tibouchi and Huaxiong Wang, editors, *ASIACRYPT 2021, Part IV*, volume 13093 of *LNCS*, pages 681–710, Singapore, December 6–10, 2021. Springer, Heidelberg, Germany. doi:10.1007/978-3-030-92068-5_23.

[BDJR97] Mihir Bellare, Anand Desai, Eric Jokipii, and Phillip Rogaway. A concrete security treatment of symmetric encryption. In *38th FOCS*, pages 394–403, Miami Beach, Florida, October 19–22, 1997. IEEE Computer Society Press. doi:10.1109/SFCS.1997.646128.

[BFG+22] Alexander Bienstock, Jaiden Fairoze, Sanjam Garg, Pratyay Mukherjee, and Srinivasan Raghuraman. A more complete analysis of the Signal double ratchet algorithm. In Yevgeniy Dodis and Thomas Shrimpton, editors, *CRYPTO 2022, Part I*, volume 13507

of *LNCS*, pages 784–813, Santa Barbara, CA, USA, August 15–18, 2022. Springer, Heidelberg, Germany. doi:10.1007/978-3-031-15802-5_27.

[BG20] Colin Boyd and Kai Gellert. A Modern View on Forward Security. *The Computer Journal*, 64(4):639–652, 08 2020. arXiv:https://academic.oup.com/comjnl/article-pdf/64/4/639/37161647/bxaa104.pdf, doi:10.1093/comjnl/bxaa104.

[Bha20] Rishiraj Bhattacharyya. Memory-tight reductions for practical key encapsulation mechanisms. In Aggelos Kiayias, Markulf Kohlweiss, Petros Wallden, and Vassilis Zikas, editors, *PKC 2020, Part I*, volume 12110 of *LNCS*, pages 249–278, Edinburgh, UK, May 4–7, 2020. Springer, Heidelberg, Germany. doi:10.1007/978-3-030-45374-9_9.

[BJLS16] Christoph Bader, Tibor Jager, Yong Li, and Sven Schäge. On the impossibility of tight cryptographic reductions. In Marc Fischlin and Jean-Sébastien Coron, editors, *EUROCRYPT 2016, Part II*, volume 9666 of *LNCS*, pages 273–304, Vienna, Austria, May 8–12, 2016. Springer, Heidelberg, Germany. doi:10.1007/978-3-662-49896-5_10.

[BKW03] Avrim Blum, Adam Kalai, and Hal Wasserman. Noise-tolerant learning, the parity problem, and the statistical query model. *J. ACM*, 50(4):506–519, jul 2003. doi:10.1145/792538.792543.

[BLS01] Dan Boneh, Ben Lynn, and Hovav Shacham. Short signatures from the Weil pairing. In Colin Boyd, editor, *ASIACRYPT 2001*, volume 2248 of *LNCS*, pages 514–532, Gold Coast, Australia, December 9–13, 2001. Springer, Heidelberg, Germany. doi:10.1007/3-540-45682-1_30.

[BM92] Steven M. Bellovin and Michael Merritt. Encrypted key exchange: Password-based protocols secure against dictionary attacks. In *1992 IEEE Symposium on Security and Privacy*, pages 72–84. IEEE Computer Society Press, May 1992. doi:10.1109/RISP.1992.213269.

[BN11] Colin Boyd and Juanma González Nieto. On forward secrecy in one-round key exchange. In *IMACC*, volume 7089 of *Lecture Notes in Computer Science*, pages 451–468. Springer, 2011.

[BPR00] Mihir Bellare, David Pointcheval, and Phillip Rogaway. Authenticated key exchange secure against dictionary attacks. In Bart Preneel, editor, *EUROCRYPT 2000*, volume 1807 of *LNCS*, pages 139–155, Bruges, Belgium, May 14–18, 2000. Springer, Heidelberg, Germany. doi:10.1007/3-540-45539-6_11.

[BR94] Mihir Bellare and Phillip Rogaway. Entity authentication and key distribution. In Douglas R. Stinson, editor, *CRYPTO'93*, volume 773 of *LNCS*, pages 232–249, Santa Barbara, CA, USA, August 22–26, 1994. Springer, Heidelberg, Germany. doi:10.1007/3-540-48329-2_21.

[BR96] Mihir Bellare and Phillip Rogaway. The exact security of digital signatures: How to sign with RSA and Rabin. In Ueli M. Maurer, editor, *EUROCRYPT'96*, volume 1070 of *LNCS*, pages 399–416, Saragossa, Spain, May 12–16, 1996. Springer, Heidelberg, Germany. doi:10.1007/3-540-68339-9_34.

[BR09] Mihir Bellare and Thomas Ristenpart. Simulation without the artificial abort: Simplified proof and improved concrete security for Waters' IBE scheme. In Antoine Joux, editor, *EUROCRYPT 2009*, volume 5479 of *LNCS*, pages 407–424, Cologne, Germany, April 26–30, 2009. Springer, Heidelberg, Germany. doi:10.1007/978-3-642-01001-9_24.

[BSJ+17] Mihir Bellare, Asha Camper Singh, Joseph Jaeger, Maya Nyayapati, and Igors Stepanovs. Ratcheted encryption and key exchange: The security of messaging. In Jonathan Katz and Hovav Shacham, editors, *CRYPTO 2017, Part III*, volume 10403 of *LNCS*, pages 619–650, Santa Barbara, CA, USA, August 20–24, 2017. Springer, Heidelberg, Germany. doi:10.1007/978-3-319-63697-9_21.

[Can00] Ran Canetti. Universally composable security: A new paradigm

for cryptographic protocols. Cryptology ePrint Archive, Report 2000/067, 2000. https://eprint.iacr.org/2000/067.

[Can01] Ran Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *42nd FOCS*, pages 136–145, Las Vegas, NV, USA, October 14–17, 2001. IEEE Computer Society Press. doi:10.1109/SFCS.2001.959888.

[CCD+20] Katriel Cohn-Gordon, Cas Cremers, Benjamin Dowling, Luke Garratt, and Douglas Stebila. A formal security analysis of the signal messaging protocol. *Journal of Cryptology*, 33(4):1914–1983, October 2020. doi:10.1007/s00145-020-09360-1.

[CCG+19] Katriel Cohn-Gordon, Cas Cremers, Kristian Gjøsteen, Håkon Jacobsen, and Tibor Jager. Highly efficient key exchange protocols with optimal tightness. In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO 2019, Part III*, volume 11694 of *LNCS*, pages 767–797, Santa Barbara, CA, USA, August 18–22, 2019. Springer, Heidelberg, Germany. doi:10.1007/978-3-030-26954-8_25.

[CEv88] David Chaum, Jan-Hendrik Evertse, and Jeroen van de Graaf. An improved protocol for demonstrating possession of discrete logarithms and some generalizations. In David Chaum and Wyn L. Price, editors, *EUROCRYPT'87*, volume 304 of *LNCS*, pages 127–141, Amsterdam, The Netherlands, April 13–15, 1988. Springer, Heidelberg, Germany. doi:10.1007/3-540-39118-5_13.

[CF15] Cas Cremers and Michèle Feltz. Beyond eCK: perfect forward secrecy under actor compromise and ephemeral-key reveal. *Des. Codes Cryptogr.*, 74(1):183–218, 2015. URL: https://doi.org/10.1007/s10623-013-9852-1.

[CJSV22] Ran Canetti, Palak Jain, Marika Swanberg, and Mayank Varia. Universally composable end-to-end secure messaging. In Yevgeniy Dodis and Thomas Shrimpton, editors, *CRYPTO 2022, Part II*, volume 13508 of *LNCS*, pages 3–33, Santa Barbara, CA, USA, Au-

gust 15–18, 2022. Springer, Heidelberg, Germany. doi:10.1007/978-3-031-15979-4_1.

[CK01] Ran Canetti and Hugo Krawczyk. Analysis of key-exchange protocols and their use for building secure channels. In Birgit Pfitzmann, editor, *EUROCRYPT 2001*, volume 2045 of *LNCS*, pages 453–474, Innsbruck, Austria, May 6–10, 2001. Springer, Heidelberg, Germany. doi:10.1007/3-540-44987-6_28.

[Cor02] Jean-Sébastien Coron. Optimal security proofs for PSS and other signature schemes. In Lars R. Knudsen, editor, *EUROCRYPT 2002*, volume 2332 of *LNCS*, pages 272–287, Amsterdam, The Netherlands, April 28 – May 2, 2002. Springer, Heidelberg, Germany. doi:10.1007/3-540-46035-7_18.

[CPZ20] Melissa Chase, Trevor Perrin, and Greg Zaverucha. The signal private group system and anonymous credentials supporting efficient verifiable encryption. In Jay Ligatti, Xinming Ou, Jonathan Katz, and Giovanni Vigna, editors, *ACM CCS 2020*, pages 1445–1459, Virtual Event, USA, November 9–13, 2020. ACM Press. doi:10.1145/3372297.3417887.

[DFG+23] Gareth T. Davies, Sebastian Faller, Kai Gellert, Tobias Handirk, Julia Hesse, Máté Horváth, and Tibor Jager. Security analysis of the whatsapp end-to-end encrypted backup protocol. In Helena Handschuh and Anna Lysyanskaya, editors, *Advances in Cryptology – CRYPTO 2023*, pages 330–361, Cham, 2023. Springer Nature Switzerland.

[DFGS15] Benjamin Dowling, Marc Fischlin, Felix Günther, and Douglas Stebila. A cryptographic analysis of the TLS 1.3 handshake protocol candidates. In Indrajit Ray, Ninghui Li, and Christopher Kruegel, editors, *ACM CCS 2015*, pages 1197–1210, Denver, CO, USA, October 12–16, 2015. ACM Press. doi:10.1145/2810103.2813653.

[DFGS16] Benjamin Dowling, Marc Fischlin, Felix Günther, and Douglas Stebila. A cryptographic analysis of the TLS 1.3 draft-10 full and pre-

shared key handshake protocol. Cryptology ePrint Archive, Report 2016/081, 2016. https://eprint.iacr.org/2016/081.

[DG21] Hannah Davis and Felix Günther. Tighter proofs for the SIGMA and TLS 1.3 key exchange protocols. In Kazue Sako and Nils Ole Tippenhauer, editors, *ACNS 21, Part II*, volume 12727 of *LNCS*, pages 448–479, Kamakura, Japan, June 21–24, 2021. Springer, Heidelberg, Germany. doi:10.1007/978-3-030-78375-4_18.

[DGJL21a] Denis Diemert, Kai Gellert, Tibor Jager, and Lin Lyu. Digital signatures with memory-tight security in the multi-challenge setting. In Mehdi Tibouchi and Huaxiong Wang, editors, *ASIACRYPT 2021, Part IV*, volume 13093 of *LNCS*, pages 403–433, Singapore, December 6–10, 2021. Springer, Heidelberg, Germany. doi:10.1007/978-3-030-92068-5_14.

[DGJL21b] Denis Diemert, Kai Gellert, Tibor Jager, and Lin Lyu. More efficient digital signatures with tight multi-user security. In Juan Garay, editor, *PKC 2021, Part II*, volume 12711 of *LNCS*, pages 1–31, Virtual Event, May 10–13, 2021. Springer, Heidelberg, Germany. doi:10.1007/978-3-030-75248-4_1.

[DJ21] Denis Diemert and Tibor Jager. On the tight security of TLS 1.3: Theoretically sound cryptographic parameters for real-world deployments. *Journal of Cryptology*, 34(3):30, July 2021. doi:10.1007/s00145-021-09388-x.

[DLS21] Gérald Doussot, Marie-Sarah Lacharité, and Eric Schorn. End-to-End Encrypted Backups Security Assessment. https://research.nccgroup.com/wp-content/uploads/2021/10/NCC_Group_WhatsApp_E001000M_Report_2021-10-27_v1.2.pdf, October 2021.

[dSGFW20] Cyprien Delpech de Saint Guilhem, Marc Fischlin, and Bogdan Warinschi. Authentication in key-exchange: Definitions, relations and composition. In *2020 IEEE 33rd Computer Security Foundations Symposium (CSF)*, pages 288–303, 2020. doi:10.1109/CSF49147.2020.00028.

[FG17] Marc Fischlin and Felix Günther. Replay attacks on zero round-trip time: The case of the TLS 1.3 handshake candidates. In *2017 IEEE European Symposium on Security and Privacy, EuroS&P 2017, Paris, France, April 26-28, 2017*, pages 60–75, 2017. doi:10.1109/EuroSP.2017.18.

[FGSW16] Marc Fischlin, Felix Günther, Benedikt Schmidt, and Bogdan Warinschi. Key confirmation in key exchange: A formal treatment and implications for TLS 1.3. In *2016 IEEE Symposium on Security and Privacy*, pages 452–469, San Jose, CA, USA, May 22–26, 2016. IEEE Computer Society Press. doi:10.1109/SP.2016.34.

[FHJ20] Marc Fischlin, Patrick Harasser, and Christian Janson. Signatures from sequential-OR proofs. In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020, Part III*, volume 12107 of *LNCS*, pages 212–244, Zagreb, Croatia, May 10–14, 2020. Springer, Heidelberg, Germany. doi:10.1007/978-3-030-45727-3_8.

[GBKS12] Xun Gong, Nikita Borisov, Negar Kiyavash, and Nabil Schear. Website detection using remote traffic analysis. In Simone Fischer-Hübner and Matthew K. Wright, editors, *PETS 2012*, volume 7384 of *LNCS*, pages 58–78, Vigo, Spain, July 11–13, 2012. Springer, Heidelberg, Germany. doi:10.1007/978-3-642-31680-7_4.

[GGJJ23] Kai Gellert, Kristian Gjøsteen, Håkon Jacobsen, and Tibor Jager. On optimal tightness for key exchange with full forward secrecy via key confirmation. In Helena Handschuh and Anna Lysyanskaya, editors, *Advances in Cryptology – CRYPTO 2023*, pages 297–329, Cham, 2023. Springer Nature Switzerland.

[GGJT22] Ashrujit Ghoshal, Riddhi Ghosal, Joseph Jaeger, and Stefano Tessaro. Hiding in plain sight: Memory-tight proofs via randomness programming. In Orr Dunkelman and Stefan Dziembowski, editors, *EUROCRYPT 2022, Part II*, volume 13276 of *LNCS*, pages 706–735, Trondheim, Norway, May 30 – June 3, 2022. Springer, Heidelberg, Germany. doi:10.1007/978-3-031-07085-3_24.

[GJ18] Kristian Gjøsteen and Tibor Jager. Practical and tightly-secure digital signatures and authenticated key exchange. In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO 2018, Part II*, volume 10992 of *LNCS*, pages 95–125, Santa Barbara, CA, USA, August 19–23, 2018. Springer, Heidelberg, Germany. doi:10.1007/978-3-319-96881-0_4.

[GJLN22] Kai Gellert, Tibor Jager, Lin Lyu, and Tom Neuschulten. On fingerprinting attacks and length-hiding encryption. In Steven D. Galbraith, editor, *CT-RSA 2022*, volume 13161 of *LNCS*, pages 345–369, Virtual Event, March 1–2, 2022. Springer, Heidelberg, Germany. doi:10.1007/978-3-030-95312-6_15.

[GJT20] Ashrujit Ghoshal, Joseph Jaeger, and Stefano Tessaro. The memory-tightness of authenticated encryption. In Daniele Micciancio and Thomas Ristenpart, editors, *CRYPTO 2020, Part I*, volume 12170 of *LNCS*, pages 127–156, Santa Barbara, CA, USA, August 17–21, 2020. Springer, Heidelberg, Germany. doi:10.1007/978-3-030-56784-2_5.

[GM84] Shafi Goldwasser and Silvio Micali. Probabilistic encryption. *Journal of Computer and System Sciences*, 28(2):270–299, 1984.

[GT20] Ashrujit Ghoshal and Stefano Tessaro. On the memory-tightness of hashed ElGamal. In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020, Part II*, volume 12106 of *LNCS*, pages 33–62, Zagreb, Croatia, May 10–14, 2020. Springer, Heidelberg, Germany. doi:10.1007/978-3-030-45724-2_2.

[Gün20] Felix Günther. Modeling advanced security aspects of key exchange and secure channel protocols. *it-Information Technology*, 62(5-6):287–293, 2020.

[Hin02] Andrew Hintz. Fingerprinting websites using traffic analysis. In Roger Dingledine and Paul F. Syverson, editors, *PET 2002*, volume 2482 of *LNCS*, pages 171–178, San Francisco, CA, USA, April 14–15, 2002. Springer, Heidelberg, Germany. doi:10.1007/3-540-36467-6_13.

[HK17] Gottfried Herold and Elena Kirshanova. Improved algorithms for the approximate k-list problem in euclidean norm. In Serge Fehr, editor, *PKC 2017, Part I*, volume 10174 of *LNCS*, pages 16–40, Amsterdam, The Netherlands, March 28–31, 2017. Springer, Heidelberg, Germany. doi:10.1007/978-3-662-54365-8_2.

[JK22] Joseph Jaeger and Akshaya Kumar. Memory-tight multi-challenge security of public-key encryption. In Shweta Agrawal and Dongdai Lin, editors, *ASIACRYPT 2022, Part III*, volume 13793 of *LNCS*, pages 454–484, Taipei, Taiwan, December 5–9, 2022. Springer, Heidelberg, Germany. doi:10.1007/978-3-031-22969-5_16.

[JL09] Antoine Joux and Stefan Lucks. Improved generic algorithms for 3-collisions. In Mitsuru Matsui, editor, *ASIACRYPT 2009*, volume 5912 of *LNCS*, pages 347–363, Tokyo, Japan, December 6–10, 2009. Springer, Heidelberg, Germany. doi:10.1007/978-3-642-10366-7_21.

[JMM19] Daniel Jost, Ueli Maurer, and Marta Mularczyk. A unified and composable take on ratcheting. In Dennis Hofheinz and Alon Rosen, editors, *TCC 2019, Part II*, volume 11892 of *LNCS*, pages 180–210, Nuremberg, Germany, December 1–5, 2019. Springer, Heidelberg, Germany. doi:10.1007/978-3-030-36033-7_7.

[KK12] Saqib A. Kakvi and Eike Kiltz. Optimal security proofs for full domain hash, revisited. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 537–553, Cambridge, UK, April 15–19, 2012. Springer, Heidelberg, Germany. doi:10.1007/978-3-642-29011-4_32.

[KK18] Saqib A. Kakvi and Eike Kiltz. Optimal security proofs for full domain hash, revisited. *Journal of Cryptology*, 31(1):276–306, January 2018. doi:10.1007/s00145-017-9257-9.

[Kra05] Hugo Krawczyk. HMQV: A high-performance secure Diffie-Hellman protocol. In Victor Shoup, editor, *CRYPTO 2005*, volume 3621 of *LNCS*, pages 546–566, Santa Barbara, CA, USA, August 14–18, 2005. Springer, Heidelberg, Germany. doi:10.1007/11535218_33.

[KW03] Jonathan Katz and Nan Wang. Efficiency improvements for signature schemes with tight security reductions. In Sushil Jajodia, Vijayalakshmi Atluri, and Trent Jaeger, editors, *ACM CCS 2003*, pages 155–164, Washington, DC, USA, October 27–30, 2003. ACM Press. doi:10.1145/948109.948132.

[Lew23] Direct correspondences with Kevin Lewi and other members of the WhatsApp engineering team, 2022-2023.

[LL06] Marc Liberatore and Brian Neil Levine. Inferring the source of encrypted HTTP connections. In Ari Juels, Rebecca N. Wright, and Sabrina De Capitani di Vimercati, editors, *ACM CCS 2006*, pages 255–263, Alexandria, Virginia, USA, October 30 – November 3, 2006. ACM Press. doi:10.1145/1180405.1180437.

[LLMP93] A. K. Lenstra, H. W. Lenstra, M. S. Manasse, and J. M. Pollard. The number field sieve. In Arjen K. Lenstra and Hendrik W. Lenstra, editors, *The development of the number field sieve*, pages 11–42, Berlin, Heidelberg, 1993. Springer Berlin Heidelberg.

[MHJT14] Brad Miller, Ling Huang, Anthony D. Joseph, and J. D. Tygar. I know why you went to the clinic: Risks and realization of HTTPS traffic analysis. In Emiliano De Cristofaro and Steven J. Murdoch, editors, *PETS 2014*, volume 8555 of *LNCS*, pages 143–163, Amsterdam, The Netherlands, July 16–18, 2014. Springer, Heidelberg, Germany. doi:10.1007/978-3-319-08506-7_8.

[NY90] Moni Naor and Moti Yung. Public-key cryptosystems provably secure against chosen ciphertext attacks. In *Proceedings of the twenty-second annual ACM symposium on Theory of computing*, pages 427–437, 1990.

[Pol78] John M Pollard. Monte carlo methods for index computation (mod p). *Mathematics of computation*, 32(143):918–924, 1978.

[Pom96] Carl Pomerance. A tale of two sieves. *Notices of the American Mathematical Society*, 43(12):1473–1485, 1996.

[PRS11] Kenneth G. Paterson, Thomas Ristenpart, and Thomas Shrimpton. Tag size does matter: Attacks and proofs for the TLS record protocol. In Dong Hoon Lee and Xiaoyun Wang, editors, *ASIACRYPT 2011*, volume 7073 of *LNCS*, pages 372–389, Seoul, South Korea, December 4–8, 2011. Springer, Heidelberg, Germany. doi:10.1007/978-3-642-25385-0_20.

[Res18] E. Rescorla. The Transport Layer Security (TLS) Protocol Version 1.3. RFC 8446 (Proposed Standard), August 2018. URL: https://www.rfc-editor.org/rfc/rfc8446.txt, doi:10.17487/RFC8446.

[RS92] Charles Rackoff and Daniel R. Simon. Non-interactive zero-knowledge proof of knowledge and chosen ciphertext attack. In Joan Feigenbaum, editor, *CRYPTO'91*, volume 576 of *LNCS*, pages 433–444, Santa Barbara, CA, USA, August 11–15, 1992. Springer, Heidelberg, Germany. doi:10.1007/3-540-46766-1_35.

[TV11] Cihangir Tezcan and Serge Vaudenay. On hiding a plaintext length by preencryption. In Javier Lopez and Gene Tsudik, editors, *ACNS 11*, volume 6715 of *LNCS*, pages 345–358, Nerja, Spain, June 7–10, 2011. Springer, Heidelberg, Germany. doi:10.1007/978-3-642-21554-4_20.

[WBC+08] Charles V. Wright, Lucas Ballard, Scott E. Coull, Fabian Monrose, and Gerald M. Masson. Spot me if you can: Uncovering spoken phrases in encrypted VoIP conversations. In *2008 IEEE Symposium on Security and Privacy*, pages 35–49, Oakland, CA, USA, May 18–21, 2008. IEEE Computer Society Press. doi:10.1109/SP.2008.21.

[WBMM07] Charles V. Wright, Lucas Ballard, Fabian Monrose, and Gerald M. Masson. Language identification of encrypted VoIP traffic: Alejandra y roberto or alice and bob? In Niels Provos, editor, *USENIX Security 2007*, Boston, MA, USA, August 6–10, 2007. USENIX Association.

[WG16] Tao Wang and Ian Goldberg. On realistically attacking tor with

website fingerprinting. *PoPETs*, 2016(4):21–36, October 2016. doi:10.1515/popets-2016-0027.

[Wha21] WhatsApp. Security of End-to-End Encrypted Backups. https://www.whatsapp.com/security/WhatsApp_Security_Encrypted_Backups_Whitepaper.pdf, September 2021.

[WMHT18] Yuyu Wang, Takahiro Matsuda, Goichiro Hanaoka, and Keisuke Tanaka. Memory lower bounds of reductions revisited. In Jesper Buus Nielsen and Vincent Rijmen, editors, *EUROCRYPT 2018, Part I*, volume 10820 of *LNCS*, pages 61–90, Tel Aviv, Israel, April 29 – May 3, 2018. Springer, Heidelberg, Germany. doi:10.1007/978-3-319-78381-9_3.

[Xag23] Keita Xagawa. Signatures with memory-tight security in the quantum random oracle model. Cryptology ePrint Archive, Paper 2023/1734, 2023. https://eprint.iacr.org/2023/1734. URL: https://eprint.iacr.org/2023/1734.

[Yan13] Zheng Yang. Modelling simultaneous mutual authentication for authenticated key exchange. In *FPS*, volume 8352 of *Lecture Notes in Computer Science*, pages 46–62. Springer, 2013.