Scalable Cryptography

Dennis Hofheinz[1] and Eike Kiltz[2]

[1] ETH Zürich, Switzerland hofheinz@inf.ethz.ch
[2] Ruhr-Universität Bochum, Germany eike.kiltz@rub.de

Abstract. In our modern digital society, cryptography is vital to protect the secrecy and integrity of transmitted and stored information. Settings like digital commerce, electronic banking, or simply private email communication already rely on encryption and signature schemes.
However, today's cryptographic schemes do not scale well, and thus are not suited for the increasingly large sets of data they are used on. For instance, the security guarantees currently known for RSA encryption—one of the most commonly used type of public-key encryption scheme—degrade linearly in the number of users and ciphertexts. Hence, larger settings (such as cloud computing, or simply the scenario of encrypting all existing email traffic) may enable new and more efficient attacks. To maintain a reasonable level of security in larger scenarios, RSA keylengths must be chosen significantly larger, and the scheme becomes very inefficient. Besides, a switch in RSA keylengths requires an update of the whole public key infrastructure, an impossibility in truly large scenarios. Even worse, when the scenario grows beyond an initially anticipated size, we may lose all security guarantees.
This problematic is the motivation for our project "Scalable Cryptography", which aims at offering a toolbox of cryptographic schemes that are suitable for huge sets of data. In this overview, we summarize the approach, and the main findings of our project. We give a number of settings in which it is possible to indeed provide scalable cryptographic building blocks. For instance, we survey our work on the construction of scalable public-key encryption schemes (a central cryptographic building block that helps secure communication), but also briefly mention other settings such as "reconfigurable cryptography". We also provide first results on scalable *quantum-resistant* cryptography, i.e., scalable cryptographic schemes that remain secure even in the presence of a quantum computer.

Keywords: public-key cryptography, security reductions

1 Introduction and motivation

Motivation: Public-key cryptography... Public-key cryptography, introduced by Diffie and Hellman [13] in 1976, is at the heart of modern cryptography. A public-key encryption (PKE) scheme can be used to transmit messages securely by encrypting them. The main feature that distinguishes PKE schemes from earlier encryption schemes (and in particular from symmetric encryption schemes such as AES) is the existence of two separate keys: the encryption (or, public) key is used to encrypt messages, while the decryption (or, secret) key is used to decrypt ciphertexts.

Among the first suggested PKE schemes were the RSA scheme of Rivest, Shamir, and Adleman [35], and the scheme of Merkle and Hellman [31]. Later on, many more followed, e.g., [20,9,15,32,6,12]. Today, PKE schemes are crucially used to protect large-scale systems. For instance, PKE schemes secure Internet browsers [37] (including e-banking applications such as HBCI, the home banking computer interface standard), Internet auctions [10], or simply email [39]. We stress that such applications cannot be solved with more classical methods of encryption (like symmetric encryption) alone. However, symmetric encryption schemes like AES do play a role in making such applications more efficient.

It has become a standard requirement that a cryptographic scheme (and in particular a PKE scheme) should come with provable security guarantees. Indeed, the *in*security of a cryptographic scheme can have catastrophic consequences (think of an electronic voting scheme), and is usually not immediately detectable. Hence, security cannot be achieved using a trial-and-error method, and should be argued beforehand.

The dangers of a missing security proof are best demonstrated by the PKCS Internet browser encryption standard [36,37]. This de facto standard defines how browsers should encrypt their communication when accessing sensitive websites, e.g., for e-banking, or e-commerce. An older version of that standard [36] used a PKE scheme *without* security proof, and was subsequently broken by Bleichenbacher [8]. This caused massive media attention, and made expensive updates necessary. As a result, the updated standard [37] relies upon a variant of the RSA PKE scheme *with* (heuristic) security proof.

We stress that a security proof always refers to a formal security model which covers the possible attacks in practice. Goldwasser and Micali [20] gave the first formal security notion, and proved a simple (but comparatively inefficient) PKE scheme secure in this sense. Later on, more efficient provably secure PKE schemes were devised (e.g., [9,12]), and the considered security notions were refined (e.g., [32,14,33,38]).

... *in a Big Data scenario.* Now consider the following simple but realistic example scenario. Namely, imagine that every owner of a smartphone encrypts all of his/her Internet communication (using a state-of-the art PKE scheme). Such an encryption already takes place for selected Internet connections, and usually for communication with email servers. However, for this example, we will assume that all communication is encrypted. This leads to a large-scale setting in which both the number of users and the number of encryptions is in the (large) millions. For simplicity, let us assume that there are $n_U = 2^{30}$ users, each performing $n_C = 2^{30}$ (i.e., about one billion) encryptions.[3]

We would like to derive provable security guarantees for the used encryption in this setting. This means that we would like to have a formal statement that the only way to break *any instance* of the used encryption scheme is to solve a (preferably well-understood) mathematical problem. Unfortunately, most existing PKE schemes do not scale well in this setting. For instance, the best known security guarantees for the PKCS encryption standard [37] degrade linearly in the number of users and ciphertexts. This

[3] Of course, many practical settings may actually involve fewer users or encryptions. To derive meaningful universal security guarantees, however, we are assuming what seems plausible in *some* realistic applications (like browser encryption or messaging apps).

means that while the scheme—implemented with current parameters and keylengths—is believed to be secure against attacks of complexity 2^{80}, the best guarantees we can currently derive for the same scheme in a 2^{30}-user, 2^{30}-ciphertext setting are almost trivial. (Namely, in that setting, we can only guarantee that any attack on the scheme must have complexity at least 2^{20}, i.e., the equivalent of less than a second of computing time on a modern desktop PC.)[4]

Goals of the "Scalable Encryption" project. The central goal of the "Scalable Encryption" project is to provide security models and cryptographic schemes that do scale well to Big Data scenarios. In particular, we provide cryptographic constructions that feature a "tight security proof" (i.e., a security reduction which gives guarantees that do not degrade in the size of the application setting). In the following, we will present and highlight the main contributions of the project.

2 Tightly secure cryptography

Our first and central concrete goal was to construct cryptographic schemes (and in particular PKE and signature schemes) with security guarantees that do not degrade in larger settings. Technically, we have aimed at constructing cryptographic schemes with a tight security reduction to a standard computational assumption. Several of our works prepared in the course of the "Scalable Cryptography" project have dealt with this topic.

At the core of all of these techniques lies the observation that some computational problems (such as computing discrete logarithms in a cyclic group) are *re-randomizable*. That means that one problem instance I can be re-randomized to obtain many problem instances $I_1, \ldots, I_n$. The solution of any instance I_i will then also yield a solution for the original instance I. To show scalable security of, say, a PKE scheme, one would then start from a single instance I, and seek to embed many re-randomized problem instances I_i in different instances of the PKE scheme. (For instance, a problem instance I_i might correspond to the public key of a PKE instance, while the corresponding problem solution might correspond to the secret key.) If an adversary breaks any PKE instance, this leads to a solution for I_i, which in turn yields a solution for I. In other words, breaking *any* PKE scheme instance (from a selection of many PKE instances) is no easier than breaking a single given problem instance I.

There are a number of interesting computational problems (which are known to be cryptographically useful) with this re-randomizability property. However, the difficulty in executing the aforementioned strategy is to deal with *active* adversaries (that may, e.g., send maliciously formed ciphertexts to an honest user of the encryption scheme to see how this user reacts). Such adversaries may require a security reduction as above to also exhibit at least partial knowledge about the *secret key* of honest users. This makes embedding a given challenge (with an *unknown* solution) into PKE instances much harder (since the embedded problem instance might also be easier to solve given that partial knowledge about the secret key).

[4] We are also cautious when making assumptions about attacker complexity, and will typically assume liberal upper bounds. It should be noted, however, that current (publicly known) supercomputers are known to achieve almost 2^{60} floating-point operations *per second*.

In our work, we have found various technical ways to embed problem instances into PKE and other cryptographic schemes. Namely, in our work [5] (published at the TCC 2015 conference), we have presented a general framework for constructing PKE, signature, and key exchange schemes with tight security proofs even in the face of *adaptive* corruptions. We note that the emphasis of this work does not lie in practical schemes. We merely describe a general paradigm to achieve an additional security property (security against adaptive corruptions) in large scenarios.

Our work [28] (published at the PKC 2015 conference) presents an identity-based encryption (IBE) scheme secure in large scenarios. While there are previous IBE schemes whose security does not degrade in the number of *users*, our scheme is the first IBE scheme whose security properties do not degrade in the number of *ciphertexts*. Hence, our scheme is the first IBE scheme suitable for the (very realistic) scenario of a large number of encryptions per user. The techniques developed in this work could furthermore be used in our next work, [16] (published at the EUROCRYPT 2016 conference) to develop a tightly secure PKE scheme. Our scheme is the first PKE scheme for large scenarios that does not require a mathematical pairing. As a consequence, our scheme is based upon a very standard computational assumption (the Decisional Diffie-Hellman assumption), and very efficient. This work has been awarded the "Best Paper" at the EUROCRYPT 2016 conference.

Most tightly secure encryption schemes (including the ones from [28] and [16]) share the disadvantage of a large public key. The work [25] (published at the TCC 2016 conference) presents a technique to obtain tightly secure encryption and signature schemes with small public keys (and ciphertexts, resp. signatures). Indeed, we could show that the concepts introduced in [25] lead not only to tightly secure public-key encryption schemes with short public keys (published in [17] at the CRYPTO 2017 conference), but also to tightly secure *structure-preserving* signature schemes (published in [1,18] at the CRYPTO 2017 and EUROCRYPT 2018 conferences), and identity-based encryption schemes [27] (published at ASIACRYPT 2018).

At this point, it might be interesting to explain the importance of *structure-preserving* cryptographic building blocks (like our signature schemes from [1,18]). Informally, a structure-preserving building block is one in which all public operations are algebraic (in a formally defined sense). As a consequence, it is possible to efficiently conduct non-interactive zero-knowledge proofs about their execution (e.g., using the highly efficient proof system of Groth and Sahai [21]). In other words, it is possible to efficiently and publicly prove, e.g., knowledge of a signature without releasing that signature. This enables applications like anonymous credentials (i.e., secure digital identities) which rely on *not* releasing all available information publicly. Our tightly secure structure-preserving signature schemes are the first of their kind, and form highly flexible and universal components for scalable such systems.

Our work [7] (published at the PKC 2015 conference) provides a new framework for obtaining digital signatures with a tight security reduction from standard hardness assumptions. Concretely, we show that any Chameleon Hash function can be transformed into a tree-based signature scheme with tight security. Our framework explains and generalizes most of the existing schemes as well as providing a generic means for constructing tight signature schemes based on arbitrary assumptions, which improves

the standard Merkle tree transformation. Moreover, we obtain the first tightly secure signature scheme from the SIS assumption and several schemes based on Diffie-Hellman in the standard model.

Our paper [23] (also published at the PKC 2015 conference) considers security notions for public-key encryption in a slightly more realistic multi-challenge model. We show that two well-known and widely employed public-key encryption schemes—RSA Optimal Asymmetric Encryption Padding (RSA-OAEP) and Diffie-Hellman Integrated Encryption Standard (DHIES)—are secure in this model. Surprisingly, our reductions are optimal in terms of tightness in the sense that they are as tight as the ones for standard security. In the follow-up work [24] (to be published at the ASIACRYPT 2016 conference) we derive new and tight bounds for the composition of symmetric and asymmetric primitives. In particular, we consider the realistic cases where the symmetric part consists of popular modes of operation like CTR, CBC, CCM, and GCM.

We also investigate a similar generic encryption technique, the "Fujusaki-Okamoto" method to achieve secure encryption. Namely, in [26] (published at the TCC 2017 conference), we show that variants of this method achieve tight security or security against quantum computers. Similarly, and even more generically, the work [19] (published at the PKC 2018 conference), investigates the tightness of the generic "KEM-DEM" paradigm to achieve efficient public-key encryption schemes.

In the paper [29] (published at the CRYPTO 2016 conference), we perform a concrete security treatment of digital signature schemes obtained from canonical identification schemes via the Fiat-Shamir transform. If the identification scheme is random self-reducible and satisfies the weakest possible security notion (hardness of key-recoverability), then the signature scheme obtained via Fiat-Shamir is unforgeable against chosen-message attacks in the multi-user setting. Previous reductions incorporated an additional multiplicative loss of N, the number of users in the system. As an important application of our framework, we obtain a concrete security treatment for Schnorr signatures in the multi-user setting.

In the work [3] (published at the CRYPTO 2017 conference), we consider the "memory-tightness" of security reductions, as opposed to the "runtime-tightness" more commonly considered (in particular in most of the works from the previous subsection). Interestingly, this work finds that sometimes, security reductions have an inherent *intrinsic memory usage* (i.e., the reduction necessarily requires a significant amount of memory to perform its job), and that sometimes this memory usage grows with the size of the application setting. This yields another dimension of relations between different problems (and the security of certain cryptographic schemes), and shows that the scalability of cryptographic schemes can be a multi-faceted question.

The work [4] (published at the EUROCRYPT 2020 conference) which does not consider security guarantees (as given, e.g., by a security reduction), but instead investigates how the best concrete attacks on cryptographic schemes scale to larger scenarios. As a result, this work gives lower bounds (and thus also security guarantees) by more directly considering all possible attacks in a generalized setting, the generic group model.

The results we have surveyed so far are concerned with the quality of a security reduction as a measure of scalability. This is a very important factor when deriving concrete security guarantees, but not the only one. For instance, in our work [22] (pub-

lished at the TCC 2016 conference), we have formalized the notion of reconfigurable cryptographic schemes. A reconfigurable scheme allows to adapt its security parameter (i.e., the quantitative level of given security guarantees) on the fly, without changing all registered user public keys (e.g., for encryption or signature schemes). Hence, reconfigurable cryptographic schemes avoid an expensive update of potentially huge public key databases.

This work also contains proof-of-concept PKE and signature schemes. In these schemes, every user has a long-term public key and secret key. The security of these long-term keys is based on very weak assumptions from the realm of secret-key cryptography: in our PKE scheme, for instance, the public key is the image of the secret key under a generic pseudo-random generator. These long-term keys are not directly used to encrypt or decrypt. Instead, they are used to derive short-term keys (e.g., for the RSA PKE scheme) of any desired bitlength that are then used for encryption or decryption.

3 Post-Quantum Cryptography

The security of all currently used asymmetric (public-key) cryptography relies on the intractability of only two number-theoretic intractability problems, namely the factoring problem and the discrete logarithm problem over elliptic curves and finite fields. This "monoculture" poses a dangerous security threat as, in the not too unlikely scenario of scalable quantum computers, Shor's algorithm will render all the asymmetric cryptosystems in current use immediately insecure: All data transmitted over encrypted channels – past and present – will immediately become public. This in particular also holds for the cryptography considered in the previous section. Leading international tech companies like Google and Microsoft are currently investing in building quantum computers. It can only be speculated whether large intelligence agencies are already in possession of a cryptologically useful quantum computer. For that reason, a number of standardization bodies (such as NIST) are currently selecting quantum-secure asymmetric cryptosystems. Promising candidates for building quantum-resistant asymmetric cryptosystems are, amongst others, based on finding solutions to certain difficult problems regarding codes and lattices. In this project we also worked on the foundations to find truly practical, and at the same time, provably secure encryption schemes, key exchange protocols, signature schemes, and more complex protocols based on well understood and meaningful hard mathematical problems over codes and lattices.

In the context of cryptography, a lattice is a (full-rank) discrete subgroup of R^n, commonly described by a basis. Basic lattice-based cryptosystems have already existed for almost two decades and are arguably among the most promising candidates for quantum-resilience. They are simple and efficient in that their algorithms consist mostly of matrix operations, and they currently resist sub-exponential and quantum attacks. Drawing on the seminal work of Ajtai in 1996 [2], we are able to connect the average-case complexity of lattice problems (upon which the security of our schemes is based) to their complexity in the worst case. The latter property is unique among all known hardness assumptions and is one of the many reasons why people believe in its intractability. In this context the "learning with errors" (LWE) problem emerged as a suitable abstraction for a hard problem on lattices since it was shown that solving this

problem would imply breaking a few well-studied lattice-problems in the worst case, such as the approximate shortest vector problem.

In [11] (published at EuroS&P 2018) we proposed Kyber, a simple and fast encryption scheme. The design of Kyber has its roots in the seminal LWE-based encryption scheme of Regev [34]. Since Regev's original work, the practical efficiency of passively secure LWE encryption schemes has been improved by observing that the secret key can come from the same distribution as the noise and also noticing that "LWE-like" schemes can be built by using a square (rather than a rectangular) matrix as the public key. Kyber does some further efficiency improvements such as dropping several bits from the public-keys and ciphertexts to save bandwidth. At the core of its security analysis lies the security reduction of the Fujusaki-Okamoto transformation [26] already mentioned in Section 2, which transforms any passively secure encryption scheme into one withstanding active adversaries. The key feature here is that the security reduction is tight, i.e., it does not degrade with the number of evaluations of the hash function. This, together with Kyber's extremely fast performace, makes it very suitable for big-data scenarios. As of 2020, Kyber has been selected by the NIST as one of the finalists of its Post-Quantum Cryptography Standardization process for public-key encryption.[5]

4 Open Questions

Although the project significantly advanced our understanding of scalable security (and in particular scalable security *guarantees*), many questions remain. First, we are still missing technical tools to tackle the tight security of all cryptographic building blocks: the tight security (and thus the scalability) of *hierarchically organized* schemes (such as HIBE or hierarchical signature schemes) is not well-understood, and most known results (such as [30]) are negative. Besides, there are few results about the scalability of new and modern cryptographic building blocks such as obfuscation or functional or homomorphic encryption schemes. Even though these building blocks are extremely powerful (and imply a multitude of other building blocks and tasks), their scalability is currently unclear.

Moreover, the interplay between cryptanalytic attacks and the guarantees given by security reductions is generally not well-understood. The work of [4] is a promising step in this direction, but there remains a lot to be done.

References

1. Abe, M., Hofheinz, D., Nishimaki, R., Ohkubo, M., Pan, J.: Compact structure-preserving signatures with almost tight security. In: CRYPTO 2017, Part II. pp. 548–580. Springer, Heidelberg (2017). https://doi.org/10.1007/978-3-319-63715-0_19
2. Ajtai, M.: Generating hard instances of lattice problems (extended abstract). In: 28th ACM STOC. pp. 99–108. ACM Press (1996). https://doi.org/10.1145/237814.237838
3. Auerbach, B., Cash, D., Fersch, M., Kiltz, E.: Memory-tight reductions. In: CRYPTO 2017, Part I. pp. 101–132. Springer, Heidelberg (2017). https://doi.org/10.1007/978-3-319-63688-7_4

[5] `https://csrc.nist.gov/projects/post-quantum-cryptography`

4. Auerbach, B., Giacon, F., Kiltz, E.: Everybody's a target: Scalability in public-key encryption. In: EUROCRYPT 2020, Part III. pp. 475–506. Springer, Heidelberg (2020). https://doi.org/10.1007/978-3-030-45727-3_16
5. Bader, C., Hofheinz, D., Jager, T., Kiltz, E., Li, Y.: Tightly-secure authenticated key exchange. In: TCC 2015, Part I. pp. 629–658. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-46494-6_26
6. Bellare, M., Rogaway, P.: Optimal asymmetric encryption. In: EUROCRYPT'94. pp. 92–111. Springer, Heidelberg (1995). https://doi.org/10.1007/BFb0053428
7. Blazy, O., Kakvi, S.A., Kiltz, E., Pan, J.: Tightly-secure signatures from chameleon hash functions. In: PKC 2015. pp. 256–279. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-46447-2_12
8. Bleichenbacher, D.: Chosen ciphertext attacks against protocols based on the RSA encryption standard PKCS #1. In: CRYPTO'98. pp. 1–12. Springer, Heidelberg (1998). https://doi.org/10.1007/BFb0055716
9. Blum, M., Goldwasser, S.: An efficient probabilistic public-key encryption scheme which hides all partial information. In: CRYPTO'84. pp. 289–302. Springer, Heidelberg (1984)
10. Bogetoft, P., Christensen, D.L., Damgård, I., Geisler, M., Jakobsen, T., Krøigaard, M., Nielsen, J.D., Nielsen, J.B., Nielsen, K., Pagter, J., Schwartzbach, M.I., Toft, T.: Secure multiparty computation goes live. In: FC 2009. pp. 325–343. Springer, Heidelberg (2009)
11. Bos, J.W., Ducas, L., Kiltz, E., Lepoint, T., Lyubashevsky, V., Schanck, J.M., Schwabe, P., Seiler, G., Stehlé, D.: CRYSTALS - kyber: A cca-secure module-lattice-based KEM. In: 2018 IEEE European Symposium on Security and Privacy, EuroS&P 2018, London, United Kingdom, April 24-26, 2018. pp. 353–367. IEEE (2018). https://doi.org/10.1109/EuroSP.2018.00032, https://doi.org/10.1109/EuroSP.2018.00032
12. Cramer, R., Shoup, V.: A practical public key cryptosystem provably secure against adaptive chosen ciphertext attack. In: CRYPTO'98. pp. 13–25. Springer, Heidelberg (1998). https://doi.org/10.1007/BFb0055717
13. Diffie, W., Hellman, M.E.: New directions in cryptography. IEEE Transactions on Information Theory **22**(6), 644–654 (1976)
14. Dolev, D., Dwork, C., Naor, M.: Non-malleable cryptography (extended abstract). In: 23rd ACM STOC. pp. 542–552. ACM Press (1991). https://doi.org/10.1145/103418.103474
15. ElGamal, T.: A public key cryptosystem and a signature scheme based on discrete logarithms. IEEE Transactions on Information Theory **31**, 469–472 (1985)
16. Gay, R., Hofheinz, D., Kiltz, E., Wee, H.: Tightly CCA-secure encryption without pairings. In: EUROCRYPT 2016, Part I. pp. 1–27. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-49890-3_1
17. Gay, R., Hofheinz, D., Kohl, L.: Kurosawa-desmedt meets tight security. In: CRYPTO 2017, Part III. pp. 133–160. Springer, Heidelberg (2017). https://doi.org/10.1007/978-3-319-63697-9_5
18. Gay, R., Hofheinz, D., Kohl, L., Pan, J.: More efficient (almost) tightly secure structure-preserving signatures. In: EUROCRYPT 2018, Part II. pp. 230–258. Springer, Heidelberg (2018). https://doi.org/10.1007/978-3-319-78375-8_8
19. Giacon, F., Kiltz, E., Poettering, B.: Hybrid encryption in a multi-user setting, revisited. In: PKC 2018, Part I. pp. 159–189. Springer, Heidelberg (2018). https://doi.org/10.1007/978-3-319-76578-5_6
20. Goldwasser, S., Micali, S.: Probabilistic encryption. Journal of Computer and System Sciences **28**(2), 270–299 (1984)
21. Groth, J., Sahai, A.: Efficient non-interactive proof systems for bilinear groups. In: EUROCRYPT 2008. pp. 415–432. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-78967-3_24

22. Hesse, J., Hofheinz, D., Rupp, A.: Reconfigurable cryptography: A flexible approach to long-term security. In: TCC 2016-A, Part I. pp. 416–445. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-49096-9_18
23. Heuer, F., Jager, T., Kiltz, E., Schäge, S.: On the selective opening security of practical public-key encryption schemes. In: PKC 2015. pp. 27–51. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-46447-2_2
24. Heuer, F., Poettering, B.: Selective opening security from simulatable data encapsulation. In: ASIACRYPT 2016, Part II. pp. 248–277. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-53890-6_9
25. Hofheinz, D.: Algebraic partitioning: Fully compact and (almost) tightly secure cryptography. In: TCC 2016-A, Part I. pp. 251–281. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-49096-9_11
26. Hofheinz, D., Hövelmanns, K., Kiltz, E.: A modular analysis of the Fujisaki-Okamoto transformation. In: TCC 2017, Part I. pp. 341–371. Springer, Heidelberg (2017). https://doi.org/10.1007/978-3-319-70500-2_12
27. Hofheinz, D., Jia, D., Pan, J.: Identity-based encryption tightly secure under chosen-ciphertext attacks. In: ASIACRYPT 2018, Part II. pp. 190–220. Springer, Heidelberg (2018). https://doi.org/10.1007/978-3-030-03329-3_7
28. Hofheinz, D., Koch, J., Striecks, C.: Identity-based encryption with (almost) tight security in the multi-instance, multi-ciphertext setting. In: PKC 2015. pp. 799–822. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-46447-2_36
29. Kiltz, E., Masny, D., Pan, J.: Optimal security proofs for signatures from identification schemes. In: CRYPTO 2016, Part II. pp. 33–61. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-53008-5_2
30. Lewko, A.B., Waters, B.: Why proving HIBE systems secure is difficult. In: EUROCRYPT 2014. pp. 58–76. Springer, Heidelberg (2014). https://doi.org/10.1007/978-3-642-55220-5_4
31. Merkle, R.C., Hellman, M.E.: Hiding information and signatures in trapdoor knapsacks. IEEE Trans. Information Theory **24**(5), 525–530 (1978)
32. Naor, M., Yung, M.: Public-key cryptosystems provably secure against chosen ciphertext attacks. In: 22nd ACM STOC. pp. 427–437. ACM Press (1990). https://doi.org/10.1145/100216.100273
33. Rackoff, C., Simon, D.R.: Non-interactive zero-knowledge proof of knowledge and chosen ciphertext attack. In: CRYPTO'91. pp. 433–444. Springer, Heidelberg (1992). https://doi.org/10.1007/3-540-46766-1_35
34. Regev, O.: On lattices, learning with errors, random linear codes, and cryptography. In: 37th ACM STOC. pp. 84–93. ACM Press (2005). https://doi.org/10.1145/1060590.1060603
35. Rivest, R.L., Shamir, A., Adleman, L.M.: A method for obtaining digital signatures and public-key cryptosystems. Commun. ACM **21**(2), 120–126 (1978)
36. RSA Laboratories: PKCS #1: RSA Encryption Standard, Version 1.5 (1993)
37. RSA Laboratories: PKCS #1: RSA Cryptography Standard, Version 2.1 (2002)
38. Shoup, V.: Why chosen ciphertext security matters. Tech. Rep. RZ 3076, IBM Zurich Research Laboratory (1998)
39. Zimmermann, P.R.: PGP: Pretty good privacy (1991)