

Bloom Filter Encryption and Applications to Efficient Forward-Secret 0-RTT Key Exchange

David Derler^{1,‡}, Kai Gellert², Tibor Jager², Daniel Slamanig³, and Christoph Striecks³

¹ DFINITY

david@dfinity.org

² University of Wuppertal, Germany

{kai.gellert, tibor.jager}@uni-wuppertal.de

³ AIT Austrian Institute of Technology

{daniel.slamanig, christoph.striecks}@ait.ac.at

Abstract. Forward secrecy is considered an essential design goal of modern key establishment (KE) protocols, such as TLS 1.3, for example. Furthermore, efficiency considerations such as zero round-trip time (0-RTT), where a client is able to send cryptographically protected payload data along with the very first KE message, are motivated by the practical demand for secure low-latency communication.

For a long time, it was unclear whether protocols that simultaneously achieve 0-RTT and full forward secrecy exist. Only recently, the first forward-secret 0-RTT protocol was described by Günther et al. (EUROCRYPT 2017). It is based on Puncturable Encryption. Forward secrecy is achieved by “puncturing” the secret key after each decryption operation, such that a given ciphertext can only be decrypted once (cf. also Green and Miers, S&P 2015). Unfortunately, their scheme is completely impractical, since one puncturing operation takes between 30 seconds and several minutes for reasonable security and deployment parameters, such that this solution is only a first feasibility result, but not efficient enough to be deployed in practice.

In this paper, we introduce a new primitive that we term Bloom Filter Encryption (BFE), which is derived from the probabilistic Bloom filter data structure. We describe different constructions of BFE schemes, and show how these yield new puncturable encryption mechanisms with extremely efficient puncturing. Most importantly, a puncturing operation only involves a small number of very efficient computations, plus the deletion of certain parts of the secret key, which outperforms previous constructions by orders of magnitude. This gives rise to the first forward-secret 0-RTT protocols that are efficient enough to be deployed in practice. We believe that BFE will find applications beyond forward-secret 0-RTT protocols.

Keywords: Bloom filter encryption \diamond Bloom filter \diamond 0-RTT \diamond forward secrecy \diamond key exchange \diamond puncturable encryption

This is a major extension of a paper which appears in Advances in Cryptology - EUROCRYPT 2018 - 37th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Tel Aviv, Israel, April 29-May 3, 2018, Proceedings.

[‡] Work done while with Graz University of Technology.

1 Introduction

One central ingredient to secure today’s Internet are key exchange (KE) protocols with the most prominent and widely deployed instantiations thereof in the Transport Layer Security (TLS) protocol [45]. Using a KE protocol, two parties (e.g., a server and a client) are able to establish a shared secret (session key) which afterwards can be used to cryptographically protect data to be exchanged between those parties. The process of arriving at a shared secret requires the exchange of messages between client and server, which adds latency overhead to the protocol. The time required to establish a key is usually measured in round-trip times (RTTs). A novel design goal, which was introduced by Google’s QUIC protocol [47] and is also adopted in TLS version 1.3 [45], aims at developing zero round-trip time (0-RTT) protocols with strong security guarantees. So far, quite some effort was made in the cryptographic literature, e.g. [49,35], and, indeed, 0-RTT protocols are probably going to be used heavily in the future Internet as TLS version 1.3 adoption is growing rapidly. Besides TLS 1.3, Google’s QUIC protocol is used on Google webservers and within the Chrome and Opera browsers to support 0-RTT. Unfortunately, none of the above mentioned protocols are enjoying 0-RTT and full forward secrecy at the same time. Only recently, Günther, Hale, Jager, and Lauer (GHJL henceforth) [33] made progress and proposed the first 0-RTT key exchange protocol with full forward secrecy for all transmitted payload messages. However, although their 0-RTT protocol offers the desired features, their construction is not yet practical.

In more detail, GHJL’s forward-secret 0-RTT key-exchange solution is based on puncturable encryption (PE), which they showed can be constructed in a black-box way from any selectively secure hierarchical identity-based encryption (HIBE) scheme. Loosely speaking, PE is a public-key encryption primitive which provides a Puncture algorithm that, given a secret key and a ciphertext, produces an updated secret key that is able to decrypt all ciphertexts except the one it has been punctured on. PE has been introduced by Green and Miers [31] (GM henceforth) who provide an instantiation relying on a binary-tree encryption (BTE) scheme—or selectively secure HIBE—together with a key-policy attribute-based encryption (KP-ABE) [30] scheme for non-monotonic (NM) formulas with specific properties. In particular, the KP-ABE needs to provide a non-standard property to enhance existing secret keys with additional NOT gates, which is satisfied by the NM KP-ABE in [44]. Since then, PE has proved to be a valuable tool to construct public-key watermarking schemes [20], forward-secret proxy re-encryption [24]⁴, or to achieve chosen-ciphertext security for fully-homomorphic encryption [17]. However, the mentioned PE instantiations from [17,20] are based on indistinguishability obfuscation and, thus, do not yield practical schemes at all.

When looking at the two most efficient PE schemes available, i.e., GM and GHJL, they still come with severe drawbacks. In particular, puncturing in GHJL is highly inefficient and takes several seconds to minutes on decent hardware for reasonable deployment parameters. In the GM scheme, puncturing is more efficient, but the cost of decryption is very significant and increases with the number of puncturings. More precisely, cost of decryption requires a number of pairing evaluations that depends on the number of puncturings, and can be in the order of 2^{10} to 2^{20} for realistic deployment

⁴ We note that [24] uses the same techniques as in GHJL.

parameters. These issues make both of them especially unsuitable for the application in forward-secret 0-RTT key exchange in a practical setting.

Contributions. In this paper, we introduce Bloom filter encryption (BFE), which can be considered as a variant of PE [31,20,17,33]. The main difference to other existing PE constructions is that in case of BFE, we tolerate a non-negligible correctness error.⁵ This allows us to construct PE with highly efficient puncturing and in particular where puncturing only requires a few very efficient operations, i.e., to *delete* parts of the secret key, but no further expensive cryptographic operations. Altogether, this makes BFE a very suitable building block to construct practical forward-secret 0-RTT key exchange. In more detail, our contributions are as follows:

- We formalize the notion of BFE by presenting a suitable security model. The intuition behind BFE is to provide highly efficient decryption and puncturing. Interestingly, puncturing mainly consists of *deleting* parts of the secret key. This approach is in contrast to existing puncturable encryption schemes, where puncturing and/or decryption is a very expensive operation.
- We propose efficient constructions of BFE. First, we present a direct construction which uses ideas from the Boneh-Franklin identity-based encryption (IBE) scheme [12]. This construction allows us to achieve constant size public keys. Second, we present a black-box construction from a ciphertext-policy attribute-based encryption (CP-ABE) scheme that only needs to be small-universe (i.e., bounded) and to support threshold policies, which allows us to achieve constant size ciphertexts. Third, we describe a generic construction from identity-based broadcast encryption (IBBE), which is efficiently instantiable with the IBBE scheme by Delerablée [22]. This construction allows us to simultaneously achieve compact public keys and constant size ciphertexts. Finally, we propose time-based BFE (TB-BFE), an enhancement of BFE which additionally provides forward secrecy and thus prevents message suppression attacks, and provide a generic construction of TB-BFE from selectively-secure HIBEs.
- We adapt the Fujisaki-Okamoto (FO) transformation [25] to obtain CCA security in the random oracle model (ROM) to the BFE setting. This is technically non-trivial, and therefore we consider it as another interesting aspect of this work. In particular, the original FO transformation [25] works only for schemes with *perfect* correctness. Recently, Hofheinz et al. [37] described a variant which works also for schemes with *negligible* correctness error. We formalize additional properties that are required to apply the FO transform, and show that our CPA-secure constructions satisfy them. This serves as a template that allows an easy application of the FO transform in a black-box manner to BFE schemes. Moreover, we also discuss how to achieve CCA security in the standard model.
- We provide a construction of a forward-secret 0-RTT key exchange protocol (in the sense of GHJL) from TB-BFE. Furthermore, we give a detailed comparison of (TB-)BFE with other PE schemes and discuss the efficiency in the context of the proposed application to forward-secret 0-RTT key exchange. In particular, our

⁵ We discuss below why this is not only tolerable, but actually a very reasonable approach for applications like 0-RTT key exchange.

construction of forward-secret 0-RTT key-exchange from TB-BFE has none of the drawbacks mentioned in the introduction (at the cost of a somewhat larger secret key, that, however, shrinks with the number of puncturings). Consequently, our forward-secret 0-RTT key exchange can be seen as a significant step forward to construct very *practical* forward-secret 0-RTT key exchange protocols.

On tolerating a non-negligible correctness error for 0-RTT. The huge efficiency gain of our construction stems partially from the relaxation of allowing a non-negligible correctness error, which, in turn, stems from the potentially non-negligible false-positive probability of a Bloom filter. While this is unusual for classical public-key encryption schemes, we consider it as a reasonable approach to accept a small, but non-negligible correctness error for the 0-RTT mode of a key exchange protocol, in exchange for the huge efficiency gain.

For example, a $1/10000$ chance that the key establishment fails allows to use 0-RTT in 9999 out of 10000 cases on average, which is a significant practical efficiency improvement. Furthermore, the communicating parties can implement a fallback mechanism which immediately continues with running a standard 1-RTT key exchange protocol with perfect correctness, if the 0-RTT exchange fails. Thus, the resulting protocol can have the same worst-case efficiency as a 1-RTT protocol, while most of the time 0-RTT is already sufficient to establish a key and full forward secrecy is *always* achieved.

Compared to other practical 0-RTT solutions, note that both TLS 1.3 [45] and QUIC [47] have similar fallback mechanisms. Furthermore, to achieve at least a very weak form of forward secrecy, they define so called *tickets* [45] or *server configuration (SCFG) messages* [47], which expire after a certain time. Forward secrecy is only achieved after the ticket/SCFG message has expired and the associated secrets have been erased. Therefore the lifetime should be kept short. If a client connects to a server after the ticket/SCFG message has expired, then the fallback mechanism is invoked and a full 1-RTT handshake is performed. In particular for settings where a client connects only occasionally to a server, and for reasonably chosen parameters and a moderate life time of the ticket/SCFG message, which at least guarantees some weak form of forward secrecy, this requires a full handshake more often than with our approach.

Finally, note that puncturable encryption with perfect (or negligible) correctness error inherently seems to require secret keys whose size at least grows linearly with the number of puncturings. This is because any such scheme inherently must (implicitly or explicitly) encode information about the list of punctured ciphertexts into the secret key, which lower-bounds the size of the secret key [41]. By tolerating a non-negligible correctness error, we are also able to restrict the growth of the secret key to a limit which seems tolerable in practice.

Remark on forward secrecy and time-based constructions. In the literature, time-based puncturable encryption schemes are often termed puncturable *forward-secure* encryption schemes [31,33], which may seem confusing as the puncturable encryption schemes already provide mechanisms to achieve forward secrecy. The motivation for why time-based constructions was initially introduced is along the same lines and goes back to Green and Miers [31]. They described a *message suppression attack* against the forward secrecy of puncturable encryption. An adversary that suppresses message de-

livery can break forward secrecy of the primitive by compromising the receiving party’s secret at a later point in time, and retroactively decrypting all suppressed messages.

Hence, Green and Miers proposed to construct a time-based construction where the attack is only feasible until both parties move to the next time slot, achieving a form of *delayed* forward secrecy. As the time-based constructions were inspired by *forward-secure* encryption, the qualifier “forward-secure” was added to the primitive’s name. For a detailed discussion on the meaning of forward secrecy in non-interactive settings such as 0-RTT, we refer to a recent work by Gellert and Boyd [14].

We believe a distinction between time-based and non-time-based constructions is meaningful. It makes explicit that the non-time-based constructions puncture out *ciphertexts*, in order to remove decryption capability for this ciphertext. In contrast, time-based constructions additionally allow to puncture time slots, which removes decryption capability for *all possible* ciphertexts from previous time slots. For our constructions this also makes it possible to keep the size of secret keys smaller, as we explain in Section 4.

Differences to the conference version [23]. In contrast to the conference version [23], this extended version contains some additions and updates. First, we chose to present all constructions explicitly as Bloom Filter Key Encapsulation Mechanisms (BFKEMs) instead of referring to them as Bloom Filter Encryption (cf. Section 2 for a discussion). Second, we provide an additional generic construction of a BFKEM from identity-based broadcast encryption (IBBE) in Section 3.4. Furthermore, we have corrected some ambiguities and minor issues within the definitional framework. Third, we provide a more elaborate discussion on the choice of parameters to provide more insights and decision support for the practical application of our proposals.

Follow-up work. After the conference version of this paper, there was some follow-up work which we want to mention for completeness. Aviram et al. in [3] study practical forward secrecy for 0-RTT in TLS 1.3 and in particular the session resumption feature of TLS 1.3. Lauer et al. [40] introduce a single-pass circuit construction protocol with forward secrecy for Tor, called Tor 0-RTT (T0RTT), which they construct from BFE. Dallmeier et al. [21] use BFE to implement the first fully forward-secret 0-RTT key exchange in Google’s QUIC protocol and analyze its performance. Finally, there is follow-up work on puncturable encryption from Sun et al. [?] providing further constructions with negligible correctness error and different trade-offs.

Outline. The remainder of this paper is organized as follows. In Section 2, we introduce the concept of Bloom filter encryption including a discussion on how to choose suitable Bloom filter parameters for our schemes. In Section 3, we present three constructions of Bloom filter encryption alongside with a modified Fujisaki-Okamoto transformation to achieve CCA security for our schemes. In Section 4, we formally define time-based Bloom filter encryption and present a generic construction based on hierarchical identity-based encryption. Section 5 explains that our time-based construction can be used to construct forward-secret 0-RTT key exchange. In Section 6, we compare computational efficiency and parameter size of our constructions with existing constructions in literature. Section 7 concludes the results of our work.

2 Bloom Filter Encryption

Notation. Let $\lambda \in \mathbb{N}$ be the security parameter. For a finite set \mathcal{S} , we denote by $s \leftarrow^{\$} \mathcal{S}$ the process of sampling s uniformly from \mathcal{S} . For an algorithm A , let $y \leftarrow^{\$} A(\lambda, x)$ be the process of running A on input (λ, x) with access to uniformly random coins and assigning the result to y . (We may omit to mention the λ -input explicitly and assume that all algorithms take λ as input.) To make the random coins r explicit, we write $A(\lambda, x; r)$. We say an algorithm A is probabilistic polynomial time (PPT) if the running time of A is polynomial in λ . A function f is negligible if its absolute value is smaller than the inverse of any polynomial (i.e., if $\forall c \exists k_0 \forall \lambda \geq k_0 : |f(\lambda)| < 1/\lambda^c$). Furthermore, for $n \in \mathbb{N}$, let $[n] := \{1, \dots, n\}$ and let BilGen be an algorithm that, on input a security parameter 1^λ , outputs $(q, e, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, g_1, g_2) \leftarrow^{\$} \text{BilGen}(1^\lambda)$, where $\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T$ are groups of prime order q with bilinear map $e : \mathbb{G}_1 \times \mathbb{G}_2 \rightarrow \mathbb{G}_T$ and generators $g_i \in \mathbb{G}_i$ for $i \in \{1, 2\}$. Finally, we will use square brackets to access the individual bits of bitstrings, i.e., $T[i]$ denotes the i -th bit of a bitstring $T = \{0, 1\}^m$, for $m \in \mathbb{N}$.

Bloom Filter Encryption. The key idea behind Bloom Filter Encryption (BFE) is that the key pair of such a scheme is associated to a Bloom filter (BF) [10], a probabilistic data structure for the approximate set membership problem with a non-negligible false-positive probability in answering membership queries. A BF initially represents a bit array of m bits, all set to 0. Insertion takes an element and inputs it to k different hash functions each mapping the element to one of the m array positions, which are then set to 1. When querying the BF on an element, it is considered to be in the BF if all positions obtained by evaluating the hash evaluations are set to 1. The initial secret key sk output by the key generation algorithm of a BFE scheme corresponds to an empty BF. Encryption takes a message M and the public key pk , samples a random element s (acting as a tag for the ciphertext) corresponding to the universe \mathcal{U} of the BF and encrypts a message using pk with respect to the k positions set in the BF by s . A ciphertext is then basically identified by s and decryption works as long as at least one index pointed to by s in the BF is still set to 0. Puncturing the secret key with respect to a ciphertext (i.e., the tag s of the ciphertext) corresponds to inserting s in the BF (i.e., updating the corresponding indices to 1 and deleting the corresponding parts of the secret key). This basically means updating sk such that it no longer can decrypt any position indexed by s .

A note on modeling BFE. For 0-RTT key establishment, our prime application in this paper, we do not need a full-blown encryption scheme, but only a key-encapsulation mechanisms (KEM) to transport a symmetric encryption key. Consequently, we chose to focus on what we call Bloom Filter Key Encapsulation Mechanisms (BFKEMs). We stress that defining BFKEM instead of BFE does not represent any limitation, as any KEM can generically be converted into a secure full-blown encryption scheme [25]. Conversely, any secure encryption scheme trivially yields a secure KEM. For the reasons mentioned before we, henceforth, may thus use the terms BFE and BFKEM interchangeably. Nonetheless, for completeness, we give stand-alone definitions of BFE tolerating a non-negligible correctness error in Appendix A.

2.1 Formal Definition of Bloom Filters

A Bloom filter (BF) [10] is a probabilistic data structure for the approximate set membership problem. It allows a succinct representation T of a set \mathcal{S} of elements from a large universe \mathcal{U} . For elements $s \in \mathcal{S}$ a query to the BF always answers 1 (“yes”), i.e., its false-negative probability is 0. Ideally, a BF would always return 0 (“no”) for elements $s \notin \mathcal{S}$, but the succinctness of the BF comes at the cost that for any query to $s \notin \mathcal{S}$ the answer can be 1, too, but only with small probability (called the *false-positive probability*).

We will only be interested in the original construction of Bloom filters [10], and omit a general abstract definition. Instead we describe the construction from [10] directly. For a general definition we refer to [43].

Definition 1 (Bloom Filter). A Bloom filter B for set \mathcal{U} consists of algorithms $B = (\text{BFGen}, \text{BFUpdate}, \text{BFCheck})$, which are defined as follows.

BFGen(m, k): This algorithm takes as input two integers $m, k \in \mathbb{N}$. It first samples k universal hash functions H_1, \dots, H_k , where $H_j : \mathcal{U} \rightarrow [m]$, defines $H := (H_j)_{j \in [k]}$ and $T := 0^m$, and outputs (H, T) .

BFUpdate(H, T, u): Given $H = (H_j)_{j \in [k]}$, $T \in \{0, 1\}^m$, and $u \in \mathcal{U}$, this algorithm defines the updated state T' by first assigning $T' := T$. Then, it sets $T'[H_j(u)] := 1$ for all $j \in [k]$, and finally returns T' .

BFCheck(H, T, u): Given $H = (H_j)_{j \in [k]}$, $T \in \{0, 1\}^m$, and $u \in \mathcal{U}$, this algorithm returns a bit $b := \bigwedge_{j \in [k]} T[H_j(u)]$.

Relevant properties of Bloom filters. Let us summarize the properties of Bloom filters relevant to our work.

Perfect completeness. A Bloom filter always “recognizes” elements that have been added with probability 1. More precisely, let $\mathcal{S} = (s_1, \dots, s_n) \in \mathcal{U}^n$ be any vector of n elements of \mathcal{U} . Let $(H, T_0) \leftarrow^{\$} \text{BFGen}(m, k)$ and define

$$T_i = \text{BFUpdate}(H, T_{i-1}, s_i) \text{ for } i \in [n].$$

Then for all $s^* \in \mathcal{S}$ and all $(H, T_0) \leftarrow^{\$} \text{BFGen}(m, k)$ with $m, k \in \mathbb{N}$, it holds that

$$\Pr [\text{BFCheck}(H, T_n, s^*) = 1] = 1,$$

where the probability is taken over the random coins of BFGen .

Compact representation of \mathcal{S} . Independent of the size of the set $\mathcal{S} \subset \mathcal{U}$ and the representation of individual elements of \mathcal{U} , the size of representation T is a constant number of m bits. A larger size of \mathcal{S} increases only the false-positive probability, as discussed below, but not the size of the representation.

Bounded false-positive probability. The probability that an element which has not yet been added to the Bloom filter is erroneously “recognized” as being contained in the filter can be made arbitrarily small, by choosing m and k adequately, given (an upper bound on) the size of \mathcal{S} .

More precisely, let $\mathcal{S} = (s_1, \dots, s_n) \in \mathcal{U}^n$ be any vector of n elements of \mathcal{U} . Then for any $s^* \in \mathcal{U} \setminus \mathcal{S}$, the false positive probability μ is bounded by

$$\mu := \Pr[\text{BFCheck}(H, T_n, s^*) = 1] \leq \left(1 - e^{-\frac{(n+1/2)k}{m-1}}\right)^k,$$

where $(H, T_0) \leftarrow^{\$} \text{BFGen}(m, k)$, $T_i = \text{BFUpdate}(H, T_{i-1}, s_i)$ for $i \in [n]$, and the probability is taken over the random coins of BFGen . See Goel and Gupta [29] for a proof of this bound.

Discussion on the choice of parameters. In order to provide a first intuition on the concrete selection of Bloom filter parameters and their impact on the size of ciphertexts, public- and secret keys for BFE, we subsequently give some examples.

Suppose we are given an upper bound n on the number of elements inserted into the Bloom filter, and an upper bound p on the false positive probability for this number of elements that we can tolerate. Our goal is to determine the size m of the Bloom filter and the number k of hash functions to achieve a false positive probability of $\mu \leq p$ with respect to n . As already mentioned above, Goel and Gupta [29] proved that the false positive probability μ of a Bloom filter is strictly bounded by

$$\mu \leq \left(1 - e^{-\frac{(n+1/2)k}{m-1}}\right)^k.$$

Hence, if we set

$$m := \left\lceil \frac{-(n+1/2)\log_2 p}{\ln 2} \right\rceil + 1 \quad \text{and} \quad k := \left\lceil \frac{(m-1)\ln 2}{n+1/2} \right\rceil, \quad (1)$$

then due to our choice of k we obtain

$$\frac{(n+1/2)k}{m-1} \leq \frac{(n+1/2)\frac{(m-1)\ln 2}{n+1/2}}{m-1} = \ln 2$$

and therefore

$$\mu \leq \left(1 - e^{-\frac{(n+1/2)k}{m-1}}\right)^k \leq \frac{1}{2^k}.$$

Furthermore, due to the choice of m in (1), we obtain a bound on k as

$$k \geq \frac{(m-1)\ln 2}{n+1/2} \geq \frac{\left(\frac{-(n+1/2)\log_2 p}{\ln 2}\right)\ln 2}{n+1/2} = -\log_2 p$$

which yields the desired bound $\mu \leq 2^{-k} \leq p$ on the false positive probability of the Bloom filter.

Table 1 lists m and k for different values of p and n . In order to give an intuition of the impact of different choices of p and n , Table 1 also lists the size of ciphertexts, secret-, and public keys, when the BFE construction in Section 3.1 is instantiated for these parameters using the pairing-friendly BLS12-381 curve, which provides a security level of about “120-bit”.

$\lceil \log_2(p) \rceil$	$\lceil \log_2(n) \rceil$	$\lceil \log_2(m) \rceil$	k	$ C $	$ pk $	$ sk $
-7	16	20	8	215 B	215 B	30.14 MB
-7	20	24	8	215 B	215 B	482.22 MB
-7	24	28	8	215 B	215 B	7.53 GB
-7	30	34	8	215 B	215 B	482.22 GB
-10	16	20	11	260 B	260 B	43.06 MB
-10	20	24	11	260 B	260 B	688.89 MB
-10	24	28	11	260 B	260 B	10.76 GB
-10	30	34	11	260 B	260 B	688.89 GB
-16	16	21	17	350 B	350 B	68.89 MB
-16	20	25	17	350 B	350 B	1.08 GB
-16	24	29	17	350 B	350 B	17.22 GB
-16	30	35	17	350 B	350 B	1.08 TB
-20	16	21	21	410 B	410 B	86.11 MB
-20	20	25	21	410 B	410 B	1.35 GB
-20	24	29	21	410 B	410 B	21.53 GB
-20	30	35	21	410 B	410 B	1.35 TB

Table 1. Bloom filter parameters and size of public keys, secret keys, and ciphertexts of the construction from Section 3.1 (with 120-bit security level and using the pairing-friendly BLS12-381 curve) for different choices of the false positive probability p and the number of inserted elements n .

Here we need to emphasize that initially the secret key (representing the empty BF) has its maximum size, but every puncturing (i.e., addition of an element to the BF), reduces the size of the secret key. Moreover, we stress that the false-positive probability represents an upper bound as it assumes that all n elements are added to the BF. The false positive probability *before* n insertions, as a function of the number of inserted elements and for given parameters m and k , is discussed below. Finally, we note that when we use our time-based BFE approach (TB-BFE) from Section 4, we can even reduce the secret key size by reducing the maximum number of puncturings at the cost of switching the time intervals more frequently.

False-positive probability p before n insertions. So far we have argued that we can bound the probability of a non-inserted element being recognized by a BF after n elements have been added to the BF. However, we stress that this probability is far lower if only a fraction of the n elements have been added. We can illustrate this by computing the false-positive probability of an element after only $\alpha < n$ insertions.

Lemma 1. *Let (H, T_α) be a Bloom filter where α random elements have been added. The false-positive probability of a random element $u \in \mathcal{U}$ being recognized by the Bloom filter is*

$$\Pr[\text{BFCheck}(H, T_n, u)] = \left(1 - \left(1 - \frac{1}{m} \right)^{\alpha k} \right)^k.$$

We prove the above lemma in Appendix B.

To give some intuition how the false-positive probability evolves over time, we plot the above function for $n = 2^{20}$ and $k \in \{8, 11, 17, 21\}$ in Figure 1. Furthermore, we provide plots for $n \in \{2^{16}, 2^{24}, 2^{30}\}$ in Appendix C. It is clearly visible that the false-positive probability is overwhelmingly low if only a fraction of the n elements have been added to the Bloom filter.

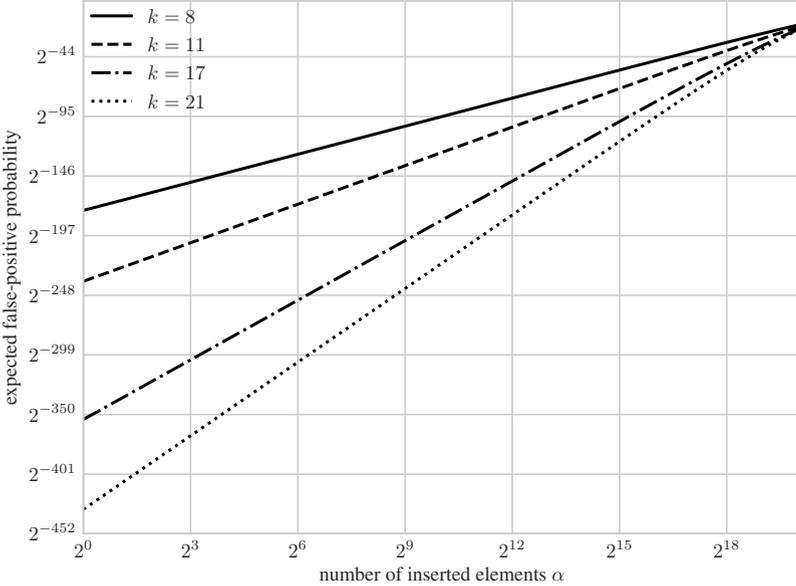


Fig. 1. The false-positive probability of a random element after α elements have been added to a Bloom filter with $n = 2^{20}$ for $k \in \{8, 11, 17, 21\}$.

A remark on Bloom filters in adversarial environments. For our bounds of the correctness error in the BFKEM, we assume that the puncturing inserts random elements (ciphertexts) into the BF. Now, an adversary could more efficiently exhaust a BF by a clever choice of the ciphertexts and thus violating our bounds. This would essentially represent a denial-of-service (DoS) attack on the scheme. We, however, stress that this class of attacks is hard to prevent in our application in general and thus we do not consider this as an attack vector. Nevertheless, one approach to counter such types of attacks on BFs is the concept of adversarial resilient Bloom-filters introduced by Naor and Yegorov in [43]. However, the efficient approach to construct such BFs in [43] requires a secret (unknown to the adversary) to evaluate BF queries, and thus would not be applicable in our setting. Naor and Yegorov additionally provide a construction secure against unbounded adversaries, which however requires to know the precise set \mathcal{S} up-

front. This is not the case in our application and we leave the study of BFs in adversarial environments for application in BFE for future work.

2.2 Formal Model of a BFKEM

Subsequently, we introduce the formal model for BFKEM, which is a KEM-variant of puncturable encryption (PE) [31,20,17,33] with the difference that with BFKEM we tolerate a non-negligible correctness error. Our Definition 2 below is a variant of the one in [33], except that we allow the key generation to take the additional parameters m and k (of the BF) as input, which specify the correctness error. As already mentioned in the introduction, resorting to present BFKEMs instead of BFE does not represent any limitation.

Definition 2 (BFKEM). *A Bloom Filter key encapsulation scheme (BFKEM) with key space \mathcal{K} is a tuple $(\text{KGen}, \text{Enc}, \text{Punc}, \text{Dec})$ of PPT algorithms:*

$\text{KGen}(1^\lambda, m, k)$: *Takes as input a security parameter λ , parameters m and k and outputs a secret and public key (sk, pk) (we assume that \mathcal{K} is implicit in pk , and that pk is implicit in sk).*

$\text{Enc}(\text{pk})$: *Takes as input a public key pk and outputs a ciphertext C and a symmetric key K .*

$\text{Punc}(\text{sk}, C)$: *Takes as input a secret key sk , a ciphertext C and outputs an updated secret key sk' .*

$\text{Dec}(\text{sk}, C)$: *Takes as input a secret key sk , a ciphertext C and deterministically computes and outputs a symmetric key K or \perp if decapsulation fails.*

Correctness. We start by defining correctness of a BFKEM scheme. Basically, here one requires that a ciphertext can always be decapsulated with unpunctured secret keys. However, we allow that if punctured secret keys are used for decapsulation then the probability that the decapsulation fails is bounded by some non-negligible function in the scheme's parameters m, k .

Definition 3 (Correctness). *We require that the following holds for all $\lambda, m, k \in \mathbb{N}$ and any $(\text{sk}, \text{pk}) \leftarrow^{\$} \text{KGen}(1^\lambda, m, k)$.*

For any (arbitrary interleaved) sequence of invocations of

$$\text{sk}_{j+1} \leftarrow^{\$} \text{Punc}(\text{sk}_j, C_j),$$

where $j \in \{1, \dots, n\}$, $\text{sk}_1 := \text{sk}$, and $(C_j, K_j) \leftarrow^{\$} \text{Enc}(\text{pk})$, it holds that

$$\Pr[\text{Dec}(\text{sk}_{n+1}, C^*) \neq K^*] \leq \left(1 - e^{-\frac{(n+1/2)k}{m-1}}\right)^k + \epsilon(\lambda),$$

where $(C^, K^*) \leftarrow^{\$} \text{Enc}(\text{pk})$ and $\epsilon(\cdot)$ is a negligible function in λ . The probability is over the random coins of KGen , Punc , and Enc .*

Remark 1. The bound $\left(1 - e^{-\frac{(n+1/2)k}{m-1}}\right)^k$ is motivated by the bound achievable by Bloom filters, cf. Equation (1) and the subsequent discussion.

2.3 Additional Properties of a BFKEM

In this section, we will define additional properties of a BFKEM that we will use for the application to 0-RTT key exchange from [33] and to construct a CCA-secure BFKEM via the Fujisaki-Okamoto (FO) transformation, as described in Section 3.2. We will show below that our constructions of CPA-secure BFKEMs satisfy these additional properties, and thus are suitable for our variant of the FO transformation, and to construct 0-RTT key exchange.

Extended correctness. Intuitively, we first require an extended variant of correctness which demands that (1) decapsulation always yields a failure when attempting to decapsulate under a secret key previously punctured for that ciphertext. This is analogous to [33]. Second, we additionally demand that (2) decapsulating an honest ciphertext with the unpunctured key does always succeed and (3) if decryption does *not* fail, then the decapsulated value must match the key returned by the Enc algorithm, for any key sk' obtained from applying any sequence of puncturing operations to the initial secret key sk .

Definition 4 (Extended Correctness). *We require that the following holds for all $\lambda, m, k, n \in \mathbb{N}$ and any $(sk, pk) \leftarrow^{\$} \text{KGen}(1^\lambda, m, k)$.*

For any (arbitrary interleaved) sequence of invocations of

$$sk_{j+1} \leftarrow^{\$} \text{Punc}(sk_j, C_j)$$

where $j \in \{1, \dots, n\}$, $sk_1 := sk$, and $(C_j, K_j) \leftarrow^{\$} \text{Enc}(pk)$, it holds that:

1. **Impossibility of false-negatives:**
 $\text{Dec}(sk_{n+1}, C_j) = \perp$ for all $j \leq n$.
2. **Perfect correctness of the initial secret key:**
 $\text{Dec}(sk, C) = K$ for all $(C, K) \leftarrow^{\$} \text{Enc}(pk)$.
3. **Semi-correctness of punctured secret keys:**
 If $\text{Dec}(sk_{j+1}, C) \neq \perp$ then $\text{Dec}(sk_{j+1}, C) = \text{Dec}(sk, C)$.

Separable randomness. We require that the encapsulation algorithm Enc essentially reads the key K in $(C, K) \leftarrow^{\$} \text{Enc}(pk)$ directly from its random input tape. Intuitively, this will later enable us to make the randomness r used by the encapsulation algorithm Enc dependent on the key K computed by Enc.

Definition 5 (Separable Randomness). *Let $\text{BFKEM} = (\text{KGen}, \text{Enc}, \text{Punc}, \text{Dec})$ be a BFKEM. We say that BFKEM has separable randomness, if one can equivalently write the encapsulation algorithm Enc as*

$$(C, K) \leftarrow^{\$} \text{Enc}(pk) = \text{Enc}(pk; (r, K)),$$

for uniformly random $(r, K) \in \{0, 1\}^{\rho+\lambda}$, where $\text{Enc}(\cdot; \cdot)$ is a deterministic algorithm whose output is uniquely determined by pk and the randomness $(r, K) \in \{0, 1\}^{\rho+\lambda}$.

Remark. We note that one can generically construct a separable BFKEM from any non-separable BFKEM. Given a non-separable BFKEM with encapsulation algorithm Enc, a separable BFKEM with encryption algorithm Enc' can be obtained as follows:

$\text{Enc}'(\text{pk}; (r, K')) : \text{Run}(C, K) \stackrel{\$}{\leftarrow} \text{Enc}(\text{pk}; r)$, set $C' := (C, K \oplus K')$ return (C', K') .

We need separability in order to apply our variant of the FO transformation, which is the reason why we have to make it explicit. Alternatively, we could have started from a non-separable BFKEM and applied the above construction. However, this adds an additional component to the ciphertext, while the construction given in Section 3.1 will already be separable, such that we can avoid this overhead.

Publicly-checkable puncturing. Finally, we need that it is efficiently checkable whether the decapsulation algorithm outputs $\perp = \text{Dec}(\text{sk}, C)$, given *not* the secret key sk , but only the public key pk , the ciphertext C to be decrypted, and the sequence C_1, \dots, C_w at which the secret key sk has been punctured.

Definition 6 (Publicly-Checkable Puncturing). *Let $\mathcal{Q} = (C_1, \dots, C_w)$ be any list of ciphertexts. We say that BFKEM allows publicly-checkable puncturing, if there exists an efficient algorithm CheckPunct with the following correctness property.*

1. *Run $(\text{sk}, \text{pk}) \stackrel{\$}{\leftarrow} \text{KGen}(1^\lambda, m, k)$.*
2. *Compute $(C_i, K_i) \stackrel{\$}{\leftarrow} \text{Enc}(\text{pk})$ and $\text{sk} = \text{Punc}(\text{sk}, C_i)$ for $i \in [w]$.*
3. *Let C be any string. We require that*

$$\perp = \text{Dec}(\text{sk}, C) \iff \perp = \text{CheckPunct}(\text{pk}, \mathcal{Q}, C).$$

From a high-level perspective, this additional property will be necessary to simulate the decryption oracle properly in the CCA security experiment when our variant of the FO transformation is applied. Together with the second and third property of Definition 4, it replaces the perfect correctness property required in the original FO transformation.

Min-entropy of ciphertexts. Following [37], we require that ciphertexts of a randomness-separable BFKEM have sufficient min-entropy, even if K is fixed:

Definition 7 (γ -Spreadness). *Let $\text{BFKEM} = (\text{KGen}, \text{Enc}, \text{Punc}, \text{Dec})$ be a randomness-separable BFKEM with ciphertext space \mathcal{C} . We say that BFKEM is γ -spread, if for any honestly generated pk , any key K and any $C \in \mathcal{C}$*

$$\Pr_{r \stackrel{\$}{\leftarrow} \{0,1\}^\rho} [C = \text{Enc}(\text{pk}; (r, K))] \leq 2^{-\gamma}.$$

2.4 Security Definitions

We define three security properties for BFKEMs. The two “standard” security notions are indistinguishability under chosen-plaintext (IND-CPA) and chosen-ciphertext (IND-CCA) attacks. In addition, we define one-wayness under chosen-plaintext attacks (OW-CPA). The latter is the weakest notion among the ones considered in this paper, and implied by both IND-CPA and IND-CCA, but sufficient for our generic construction of IND-CCA-secure BFKEMs.

Indistinguishability-based security. Figure 2 defines the IND-CPA and IND-CCA experiments for BFKEMs. The experiments are similar to the security notions for conventional KEMs, but the adversary can arbitrarily puncture the secret key via the Punc oracle and retrieve the punctured secret key via the Corr oracle, once it has been punctured on the challenge ciphertext C^* .

$\mathbf{Exp}_{\mathcal{A}, \text{BFKEM}}^{\text{T}}(\lambda, m, k)$:
 $(\text{sk}, \text{pk}) \leftarrow^{\$} \text{KGen}(1^\lambda, m, k), (C^*, K_0) \leftarrow^{\$} \text{Enc}(\text{pk}), \mathcal{Q} \leftarrow \emptyset$
 $K_1 \leftarrow^{\$} \mathcal{K}, b \leftarrow^{\$} \{0, 1\}$
 $b^* \leftarrow^{\$} \mathcal{A}^{\mathcal{O}, \text{Punc}(\text{sk}, \cdot), \text{Corr}}(\text{pk}, C^*, K_b)$
 where $\mathcal{O} \leftarrow \{\text{Dec}'(\text{sk}, \cdot)\}$ if $\text{T} = \text{IND-CCA}$ and $\mathcal{O} \leftarrow \emptyset$ otherwise.
 $\text{Dec}'(\text{sk}, C)$ behaves as Dec but returns \perp if $C = C^*$
 $\text{Punc}(\text{sk}, C)$ runs $\text{sk} \leftarrow^{\$} \text{Punc}(\text{sk}, C)$ and $\mathcal{Q} \leftarrow \mathcal{Q} \cup \{C\}$
 Corr returns sk if $C^* \in \mathcal{Q}$ and \perp otherwise
 If $b^* = b$ then return 1
 return 0

Fig. 2. Indistinguishability-based security for BFKEMs.

Definition 8 (Indistinguishability-Based Security of BFKEM). For $\text{T} \in \{\text{IND-CPA}, \text{IND-CCA}\}$, we define the advantage of an adversary \mathcal{A} in the T experiment $\mathbf{Exp}_{\mathcal{A}, \text{BFKEM}}^{\text{T}}(\lambda, m, k)$ as

$$\mathbf{Adv}_{\mathcal{A}, \text{BFKEM}}^{\text{T}}(\lambda, m, k) := \left| \Pr \left[\mathbf{Exp}_{\mathcal{A}, \text{BFKEM}}^{\text{T}}(\lambda, m, k) = 1 \right] - \frac{1}{2} \right|.$$

A Bloom Filer key-encapsulation scheme BFKEM is $\text{T} \in \{\text{IND-CPA}, \text{IND-CCA}\}$ secure, if $\mathbf{Adv}_{\mathcal{A}, \text{BFKEM}}^{\text{T}}(\lambda, m, k)$ is a negligible function in λ for all $m, k > 0$ and all PPT adversaries \mathcal{A} .

One-wayness under chosen-plaintext attack. Figure 3 defines the OW-CPA experiment. The experiment is similar to the IND-CPA experiment, except that the goal of the adversary is to recover the encapsulated key, given a random challenge ciphertext.

$\mathbf{Exp}_{\mathcal{A}, \text{BFKEM}}^{\text{OW-CPA}}(\lambda, m, k)$:
 $(\text{sk}, \text{pk}) \leftarrow^{\$} \text{KGen}(1^\lambda, m, k), (C^*, K_0) \leftarrow^{\$} \text{Enc}(\text{pk}), \mathcal{Q} \leftarrow \emptyset$
 $K_0^* \leftarrow^{\$} \mathcal{A}^{\text{Punc}(\text{sk}, \cdot), \text{Corr}}(\text{pk}, C^*)$
 where $\text{Punc}(\text{sk}, C)$ runs $\text{sk} \leftarrow^{\$} \text{Punc}(\text{sk}, C)$ and $\mathcal{Q} \leftarrow \mathcal{Q} \cup \{C\}$
 Corr returns sk if $C^* \in \mathcal{Q}$ and \perp otherwise
 If $K_0^* = K_0$ then return 1
 return 0

Fig. 3. OW-CPA security for BFKEMs.

Definition 9 (One-Wayness Under Chosen-Plaintext Attack). We define the advantage of an adversary \mathcal{A} in experiment $\mathbf{Exp}_{\mathcal{A}, \text{BFKEM}}^{\text{OW-CPA}}(\lambda, m, k)$ as

$$\mathbf{Adv}_{\mathcal{A}, \text{BFKEM}}^{\text{OW-CPA}}(\lambda, m, k) := \Pr \left[\mathbf{Exp}_{\mathcal{A}, \text{BFKEM}}^{\text{OW-CPA}}(\lambda, m, k) = 1 \right].$$

A BFKEM is OW-CPA secure, if $\mathbf{Adv}_{\mathcal{A}, \text{BFKEM}}^{\text{OW-CPA}}(\lambda, m, k)$ is a negligible function in λ for all $m, k > 0$ and all PPT adversaries \mathcal{A} .

Relation to “standard KEM security”. We would like to point out that it is possible to remove both the puncture and corrupt oracle and immediately send the secret key (punctured at the challenge ciphertext) to the adversary. However, this security definition is only equivalent to our security definition if an adversary cannot detect in which order ciphertexts have been punctured. That is, this security definition only provides reasonable security for schemes where the secret key does not reveal the order of puncturings. A formalization of this can be found in [26].

3 BFKEM Constructions

In this section, we present different BFKEM constructions. We start with a CPA secure version inspired by the hashed Boneh-Franklin identity-based encryption (IBE) scheme [12] in Section 3.1 and then show how we can obtain a CCA secure variant via the Fujisaki-Okamoto (FO) transform [25] in the random oracle model (ROM) in Section 3.2. Then, in Section 3.3 we present a construction of a CPA secure BFKEM from ciphertext-policy attribute-based encryption (CP-ABE) schemes and discuss how to obtain CCA security via the FO transform in the ROM. Finally, in Section 3.4 we present a CPA secure BFKEM from identity-based broadcast encryption (IBBE) and discuss how to obtain CCA security via the FO transform in the ROM or the CHK transform [16] without requiring random oracles.

3.1 BFKEM from Hashed IBE

Construction. In the sequel, let $\text{Params} := (q, e, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, g_1, g_2) \leftarrow^{\$} \text{BilGen}(1^\lambda)$, and $g_T = e(g_1, g_2)$. We will always assume that all algorithms described below implicitly receive these parameters as additional input. Let $B = (\text{BFGen}, \text{BFUpdate}, \text{BFCheck})$ be a Bloom filter for set \mathbb{G}_1 . Furthermore, let $G : \mathbb{N} \rightarrow \mathbb{G}_2$ and $E : \mathbb{G}_T \rightarrow \{0, 1\}^\lambda$ be cryptographic hash functions (which will be modeled as random oracles [7] in the security proof).

Let $\text{BFKEM} = (\text{KGen}, \text{Enc}, \text{Punc}, \text{Dec})$ be defined as follows.

$\text{KGen}(1^\lambda, m, k)$: This algorithm first generates a Bloom filter instance by running $(H, T) \leftarrow^{\$} \text{BFGen}(m, k)$. Then it chooses $\alpha \leftarrow^{\$} \mathbb{Z}_q$, and computes and returns

$$\text{sk} := (T, (G(i)^\alpha)_{i \in [m]}) \text{ and } \text{pk} := (g_1^\alpha, H).$$

Remark. The reader familiar with the Boneh-Franklin IBE scheme [12] may note that the secret key contains m elements of \mathbb{G}_2 , each essentially being a secret key of the Boneh-Franklin scheme for “identity” i , $i \in [m]$, with respect to “master public-key” g_1^α .

$\text{Enc}(\text{pk})$: This algorithm takes as input a public key pk of the above form. It samples a uniformly random key $K \leftarrow^{\$} \{0, 1\}^\lambda$ and exponent $r \leftarrow^{\$} \mathbb{Z}_q$. Then it computes $i_j := H_j(g_1^r)$ for $(H_j)_{j \in [k]} := H$, then $y_j = e(g_1^\alpha, G(i_j))^r$ for $j \in [k]$, and finally

$$C := (g_1^r, (E(y_j) \oplus K)_{j \in [k]}).$$

It outputs $(C, K) \in (\mathbb{G}_1 \times \{0, 1\}^{k\lambda}) \times \{0, 1\}^\lambda$.

Remark. Note that for each $j \in [k]$, the tuple $(g_1^r, E(y_j) \oplus K)$ is essentially a “hashed Boneh-Franklin IBE” ciphertext, encrypting K for “identity” $i_j = H_j(g_1^r)$ and with respect to master public key g_1^α , where the identity is derived deterministically from a “unique” (with overwhelming probability) ciphertext component g_1^r . Thus, the ciphertext C essentially consists of k Boneh-Franklin ciphertexts that share the same randomness r , each encrypting the same key K for an “identity” derived deterministically from g_1^r .

Note also that this construction of Enc satisfies the requirement of separable randomness from Definition 5. Furthermore, ciphertexts are γ -spread according to Definition 7 with $\gamma = \log_2 p$, because g_1^r is uniformly distributed over \mathbb{G}_1 .

$\text{Punc}(\text{sk}, C)$: Given a ciphertext $C := (g_1^r, (E(y_j) \oplus K)_{j \in [k]})$ and secret key $\text{sk} = (T, (\text{sk}[i])_{i \in [m]})$, the puncturing algorithm first computes $T' = \text{BFUpdate}(H, T, g_1^r)$. Then, for each $i \in [m]$ it defines

$$\text{sk}'[i] := \begin{cases} \text{sk}[i] & \text{if } T'[i] = 0, \text{ and} \\ \perp & \text{if } T'[i] = 1, \end{cases}$$

where $T'[i]$ denotes the i -th bit of T' . Finally, this algorithm returns

$$\text{sk}' := (T', (\text{sk}'[i])_{i \in [m]}).$$

Remark. Note that the above procedure is correct even if the procedure is applied repeatedly with different ciphertexts C , since the BFUpdate algorithm only changes bits of T from 0 to 1, but never from 1 to 0. So we can delete a secret key element $\text{sk}[i]$ once $T'[i]$ has been set to 1. Furthermore, we have $\text{sk}'[i] = \perp \iff T'[i] = 1$. Intuitively, this will ensure that we can use this key to decrypt a ciphertext $C := (g_1^r, (E(y_j) \oplus K)_{j \in [k]})$ if and only if $\text{BFCheck}(H, T, g_1^r) = 0$, where (H, T) is the Bloom filter instance contained in the public key. Note also that the puncturing algorithm essentially only evaluates k universal hash functions $H = (H_j)_{j \in [k]}$ and then deletes a few secret keys, which makes this procedure extremely efficient. Finally, observe that the filter state T can be efficiently recomputed given only public information, namely the list of hash functions H contained in pk and the sequence of ciphertexts C_1, \dots, C_w on which a secret key has been punctured. This yields the existence of an efficient CheckPunc according to Definition 6.

$\text{Dec}(\text{sk}, C)$: Given a secret key $\text{sk} = (T, (\text{sk}[i])_{i \in [m]})$ and a ciphertext $C := (C[0], C[i_1], \dots, C[i_k])$ it first checks whether $\text{BFCheck}(H, T, C[0]) = 1$, and outputs \perp in this case. Otherwise, note that $\text{BFCheck}(H, T, C[0]) = 0$ implies that there exists at least one index i^* with $\text{sk}[i^*] \neq \perp$. It picks the smallest index $i^* \in \{i_1, \dots, i_k\}$ such that $\text{sk}[i^*] = G(i^*)^\alpha \neq \perp$, computes

$$y_{i^*} := e(g_1^r, G(i^*)^\alpha),$$

and returns $K := C[i^*] \oplus E(y_{i^*})$.

Remark. If $\text{BFCheck}(H, T_n, C[0]) = 0$, then the decryption algorithm performs a “hashed Boneh-Franklin” decryption with a secret key for one of the identities.

Note that $\text{Dec}(\text{sk}_n, C) \neq \perp \iff \text{BFCheck}(H, T, C[0]) = 0$, which guarantees the first extended correctness property required by Definition 4. It is straightforward to verify that the other two extended correctness properties of Definition 4 hold as well.

Design choices. We note that we have chosen to base our BFKEM on *hashed* Boneh-Franklin IBE instead of standard Boneh-Franklin for two reasons. First, it allows us to keep ciphertexts short and independent of the size of the binary representation of elements of \mathbb{G}_T . This is useful, because the recent advances for computing discrete logarithms in finite extension fields [39] apply to the target group of state-of-the-art pairing-friendly elliptic curve groups. Recent assessments of the impact of these advances by Menezes et al. [42] as well as Barbulescu and Duquesne [4] suggest that for currently used efficient curve families such as BN [6] or BLS [5] curves a conservative choice of parameters for the 128 bit security level yields sizes of \mathbb{G}_T elements of $\approx 4600 - 5500$ bits. The hash function allows us to “compress” these group elements in the ciphertext to 128 bits (the size of a symmetric encryption key). Even if future research enables the construction of bilinear maps where elements of \mathbb{G}_T can be represented by 2λ bits for λ -bit security (which is optimal), it is still preferable to hash group elements to λ bits to reduce the ciphertext by a factor of about 2. Second, by modelling E as a random oracle, we can reduce security to a weaker complexity assumption.

Correctness error of this scheme. We will now explain that the correctness error of this scheme is essentially identical to the false-positive probability of the Bloom filter, up to a statistically small distance which corresponds to the probability that two independent ciphertexts share the same randomness r .

For $m, k \in \mathbb{N}$, let $(\text{sk}_0, \text{pk}) \leftarrow^{\$} \text{KGen}(1^\lambda, m, k)$, let $\mathcal{U} := \{C : (C, K) \leftarrow^{\$} \text{Enc}(\text{pk})\}$ denote the set of all valid ciphertexts with respect to pk . Let $\mathcal{S} = (C_1, \dots, C_n)$ be a list of n ciphertexts, where $(C_i, K_i) \leftarrow^{\$} \text{Enc}(\text{pk})$, and run $\text{sk}_i = \text{Punc}(\text{sk}_{i-1}, C_i)$ for $i \in [n]$ to determine the secret key sk_n obtained from puncturing sk_0 iteratively on all ciphertexts $C_i \in \mathcal{S}$.

Now let us consider the probability

$$\Pr[\text{Dec}(\text{sk}_n, C^*) \neq K^* : (C^*, K^*) \leftarrow^{\$} \text{Enc}(\text{pk}), C^* \notin \mathcal{S}]$$

that a newly generated ciphertext $C^* \notin \mathcal{S}$ is not correctly decrypted by sk_n . To this end, let $C^*[0] = g_1^{r^*}$ denote the first component of ciphertext $C^* = (g_1^{r^*}, C_1^*, \dots, C_k^*)$, and likewise let $C_i[0]$ denote the first component of ciphertext C_i for all $C_i \in \mathcal{S}$. Writing $\text{sk}_n = (T_n, (\text{sk}_n[i])_{i \in [m]})$ and $\text{pk} = (g_1^\alpha, H)$, one can now verify that we have $\text{Dec}(\text{sk}_n, C^*) \neq K^* \iff \text{BFCheck}(H, T_n, C^*[0]) = 1$, because $\text{BFCheck}(H, T_n, C^*[0]) = 0$ guarantees that there exists at least one index j such that $\text{sk}_n[H_j(C^*[0])] \neq \perp$, so correctness of decryption follows essentially from correctness of the Boneh-Franklin scheme. Thus, we have to consider the probability that $\text{BFCheck}(H, T_n, C^*[0]) = 1$. We distinguish between two cases:

1. There exists an index $i \in [n]$ such that $C^*[0] = C_i[0]$. Note that this implies immediately that $\text{BFCheck}(H, T_n, C^*[0]) = 1$. However, recall that $C^*[0] = g_1^{r^*}$ is a uniformly random element of \mathbb{G}_1 . Therefore the probability that this happens is upper bounded by n/q , which is negligibly small.

2. $C^*[0] \neq C_i[0]$ for all $i \in [n]$. In this case, as explained in Section 2.1, the soundness of the Bloom filter guarantees that

$$\Pr[\text{BFCheck}(H, T_n, C^*[0]) = 1] \leq \left(1 - e^{-\frac{(n+1/2)k}{m-1}}\right)^k \leq 2^{-k}.$$

In summary, the correctness error of this scheme from the discussion in Section 2.1 is approximately $2^{-k} + n/q$. Since n/q is negligibly small, this essentially amounts to the correctness error of the Bloom filter, which in turn depends on the number of ciphertexts n , and the choice of parameters m, k .

Flexible instantiability of this scheme. Our scheme is highly parameterizable in the sense that we can adjust the size of keys and ciphertexts by adjusting the correctness error (determined by the choice of parameters m, k that in turn determine the false-positive probability of the Bloom filter) of our scheme.

Additional properties. As already explained in the remarks after the description of the individual algorithms of BFKEM, the scheme satisfies the requirements of Definitions 4, 5, 6, and 7.

IND-CPA-security. We base IND-CPA-security on a bilinear computational Diffie-Hellman variant in the bilinear groups generated by BilGen.

Definition 10 (BCDH [12]). We define the advantage of adversary \mathcal{A} in solving the BCDH problem with respect to BilGen as

$$\text{Adv}_{\mathcal{A}, \text{BilGen}}^{\text{BCDH}}(\lambda) := \Pr[e(g_1, h_2)^{r^\alpha} \leftarrow^{\$} \mathcal{A}(\text{Params}, g_1^r, g_1^\alpha, g_2^\alpha, h_2)],$$

where $\text{Params} = (p, e, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, g_1, g_2) \leftarrow^{\$} \text{BilGen}(1^\lambda)$, and $(g_1^r, g_1^\alpha, g_2^\alpha, h_2) \leftarrow^{\$} \mathbb{G}_1^2 \times \mathbb{G}_2^2$.

Theorem 1. From each efficient adversary \mathcal{B} that issues u queries to random oracle E we can construct an efficient adversary \mathcal{A} with

$$\text{Adv}_{\mathcal{B}, \text{BilGen}}^{\text{IND-CPA}}(\lambda, m, k) \leq ku \cdot \text{Adv}_{\mathcal{A}, \text{BilGen}}^{\text{BCDH}}(\lambda).$$

Proof. Algorithm \mathcal{A} receives as input a BCDH-challenge tuple $(g_1^r, g_1^\alpha, g_2^\alpha, h_2)$. It runs adversary \mathcal{B} as a subroutine by simulating the $\text{Exp}_{\mathcal{B}, \text{BilGen}}^{\text{IND-CPA}}(\lambda, m, k)$ experiment, including random oracles G and E , as follows.

First, it defines $\mathcal{Q} := \emptyset$, runs $(H, T) \leftarrow^{\$} \text{BFGen}(m, k)$, and defines the public key as $\text{pk} := (g_1^\alpha, H)$. Note that this public key is identically distributed to a public key output by $\text{KGen}(1^\lambda, m, k)$. In order to simulate the challenge ciphertext, the adversary chooses a random key $K \leftarrow^{\$} \{0, 1\}^\lambda$ and k uniformly random values $Y_j \leftarrow^{\$} \{0, 1\}^\lambda$, $j \in [k]$, and defines the challenge ciphertext as $C^* := (g_1^r, (Y_j)_{j \in [k]})$. Finally, it outputs (pk, C^*, K) to \mathcal{B} .

Whenever \mathcal{B} queries $\text{Punc}(\text{sk}, \cdot)$ on input $C = (C[0], \dots)$, then \mathcal{A} updates T by running $T = \text{BFUpdate}(H, T, C[0])$, and $\mathcal{Q} \leftarrow \mathcal{Q} \cup \{C\}$.

Whenever a random oracle query to $G : \mathbb{N} \rightarrow \mathbb{G}_2$ is made (either by \mathcal{A} or \mathcal{B}), with input $\ell \in \mathbb{N}$, then \mathcal{A} responds with $G(\ell)$, if $G(\ell)$ has already been defined. If not, then \mathcal{A} chooses a random integer $r_\ell \leftarrow^{\$} \mathbb{Z}_q$, and returns $G(\ell)$, where

$$G(\ell) := \begin{cases} h_2 \cdot g_2^{r_\ell} & \text{if } \ell \in \{H_j(g_1^r) : j \in [k]\}, \text{ and} \\ g_2^{r_\ell} & \text{otherwise.} \end{cases}$$

This definition of G allows \mathcal{A} to simulate the Corr oracle as follows. When \mathcal{B} queries Corr, then it first checks whether $C^* \in \mathcal{Q}$, and returns \perp if this does not hold. Otherwise, note that we must have $\forall j \in [k] : T[H_j(g_1^r)] = 0$, where $H = (H_j)_{j \in [k]}$ and $T[\ell]$ denotes the ℓ -th bit of T . Thus, by the simulation of G described above, \mathcal{A} is able to compute and return $G(\ell)^\alpha = (g_2^{r_\ell})^\alpha = (g_2^r)^{r_\ell}$ for all ℓ with $\ell \notin \{H_j(g_1^r) : j \in [k]\}$, and therefore in particular for all ℓ with $T[\ell] = 1$. This enables the perfect simulation of Corr.

Finally, whenever \mathcal{B} queries random oracle $E : \mathbb{G}_T \rightarrow \{0, 1\}^\lambda$ on input y , then \mathcal{A} responds with $E(y)$, if $E(y)$ has already been defined. If not, then \mathcal{A} chooses a random string $Y \leftarrow_{\$} \{0, 1\}^\lambda$, assigns $E(y) := Y$, and returns $E(y)$. Now we have to distinguish between two types of adversaries.

1. A Type-1 adversary \mathcal{B} never queries E on input of a value y , such that there exists $j \in [k]$ such that $y = e(g_1^\alpha, G(H_j(g_1^r)))^r$. Note that in this case the value $Y'_j := E(e(g_1^\alpha, G(H_j(g_1^r))))^r$ remains undefined for all $j \in [k]$ throughout the entire experiment. Thus, information-theoretically, a Type-1 adversary receives no information about the key encrypted in the challenge ciphertext C^* , and thus can only have advantage $\text{Adv}_{\mathcal{B}, \text{BFKEM}}^{\text{IND-CPA}}(\lambda, m, k) = 0$, in which case the theorem holds trivially.
2. A Type-2 adversary queries $E(y)$ such that there exists $j \in [k]$ with $y = e(g_1^\alpha, G(H_j(g_1^r)))^r$. \mathcal{A} uses a Type-2 adversary to solve the BCDH challenge as follows. At the beginning of the game, it picks two indices $(u^*, j^*) \leftarrow_{\$} [u] \times [k]$ uniformly random. When \mathcal{B} outputs y in its u^* -th query to E , then \mathcal{A} computes and outputs $W := y \cdot e(g_1^\alpha, g_2^r)^{-r_\ell}$. Since \mathcal{B} is a Type-2 adversary, we know that at some point it will query $E(y)$ with $y = e(g_1^\alpha, G(H_j(g_1^r)))^r$ for some $j \in [k]$. If this is the u^* -th query and we have $j = j^*$, which happens with probability $1/(uk)$, then we have

$$\begin{aligned} W &= y \cdot e(g_1^r, g_2^\alpha)^{-r_\ell} = e(g_1^\alpha, G(H_j(g_1^r)))^r \cdot e(g_1^\alpha, g_2^r)^{-r_\ell} \\ &= e(g_1^\alpha, h_2 \cdot g_2^{r_\ell})^r \cdot e(g_1^\alpha, g_2^r)^{-r_\ell} = e(g_1^\alpha, h_2)^r \cdot e(g_1^\alpha, g_2^{r_\ell})^r \cdot e(g_1^\alpha, g_2^r)^{-r_\ell} \end{aligned}$$

and thus W is a solution to the given BCDH instance. Note that r_ℓ is chosen in the simulation and therefore known. \square

OW-CPA-Security. The following theorem can either be proven analogous to Theorem 1, or based on the fact that IND-CPA-security implies OW-CPA-security. Therefore we give it without proof.

Theorem 2. *From each efficient adversary \mathcal{B} that issues u queries to random oracle E we can construct an efficient adversary \mathcal{A} with*

$$\text{Adv}_{\mathcal{B}, \text{BFKEM}}^{\text{OW-CPA}}(\lambda, m, k) \leq uq \cdot \text{Adv}_{\mathcal{A}, \text{BilGen}}^{\text{BCDH}}(\lambda).$$

Remark 2. The construction presented above allows to switch the roles of \mathbb{G}_1 and \mathbb{G}_2 , i.e., to switch all elements in \mathbb{G}_1 to \mathbb{G}_2 and vice versa. This might be beneficial regarding the size of the secret key when instantiating our construction using a bilinear group where the representation of elements in \mathbb{G}_2 requires more space than the representation of elements in \mathbb{G}_1 .

3.2 CCA-Security of the BFKEM from Hashed IBE via Fujisaki-Okamoto

We obtain a CCA-secure BFKEM by adopting the Fujisaki-Okamoto (FO) transformation [25] to the BFKEM setting. Since the FO transformation does not work generically for any BFKEM, we have to use the additional requirements on the underlying BFKEM that have been defined in Section 2.3. These additional properties enable us to overcome the difficulty that the original Fujisaki-Okamoto transformation from [25] requires *perfect* correctness. We remark that Hofheinz *et al.* [38] give a new, modular analysis of the FO transformation, which also works for public key *encryption* schemes with *negligible* correctness error, however, it is not applicable to BFKEMs, because, due to their non-negligible correctness error, the bounds given in [38] provide insufficient security in this case.

Construction. Let $\text{BFKEM} = (\text{KGen}, \text{Enc}, \text{Punc}, \text{Dec})$ be a BFKEM with *separable randomness* according to Definition 5. This means that we can write Enc equivalently as $(C, K) \leftarrow^{\$} \text{Enc}(\text{pk}) = \text{Enc}(\text{pk}; (r, K))$ for uniformly random $(r, K) \leftarrow^{\$} \{0, 1\}^{\rho+\lambda}$. In the sequel, let R be a hash function (modeled as a random oracle in the security proof), mapping $R : \{0, 1\}^* \rightarrow \{0, 1\}^{\rho+\lambda}$. We construct a new scheme $\text{BFKEM}' = (\text{KGen}', \text{Enc}', \text{Punc}', \text{Dec}')$ as follows.

$\text{KGen}'(1^\lambda, m, k)$: This algorithm is identical to KGen .

$\text{Enc}'(\text{pk})$: Algorithm Enc' samples $K \leftarrow^{\$} \{0, 1\}^\lambda$. Then it computes $(r, K') := R(K) \in \{0, 1\}^{\rho+\lambda}$, runs $(C, K) \leftarrow^{\$} \text{Enc}(\text{pk}; (r, K))$, and returns (C, K') .

$\text{Punc}'(\text{sk}, C)$: This algorithm is identical to Punc .

$\text{Dec}'(\text{sk}, C)$: This algorithm first runs $K \leftarrow^{\$} \text{Dec}(\text{sk}, C)$, and returns \perp if $K = \perp$. Otherwise, it computes $(r, K') = R(K)$, and checks consistency of the ciphertext by verifying that $(C, K) = \text{Enc}(\text{pk}; (r, K))$. If this does not hold, then it outputs \perp . Otherwise it outputs K' .

Correctness error and extended correctness. Both the correctness error and the extended correctness according to Definition 4 are not affected by the Fujisaki-Okamoto transform. Therefore these properties are inherited from the underlying scheme. The fact that the first property of Definition 4 is satisfied makes the scheme suitable for the application to 0-RTT key establishment.

IND-CCA-security. The security proof reduces security of our modified scheme to the OW-CPA-security of the scheme from Section 3.

Theorem 3. *Let $\text{BFKEM} = (\text{KGen}, \text{Enc}, \text{Punc}, \text{Dec})$ be a BFKEM scheme that satisfies the additional properties of Definitions 4 and 6, and which is γ -spread according to Definition 7. Let $\text{BFKEM}' = (\text{KGen}', \text{Enc}', \text{Punc}', \text{Dec}')$ be the scheme described in Section 3.2. From each efficient adversary \mathcal{A} that issues at most $q_{\mathcal{O}}$ queries to oracle \mathcal{O} and q_R queries to random oracle R , we can construct an efficient adversary \mathcal{B} with*

$$\text{Adv}_{\mathcal{A}, \text{BFKEM}'}^{\text{IND-CCA}}(\lambda, m, k) \leq q_R \cdot \text{Adv}_{\mathcal{B}, \text{BFKEM}}^{\text{OW-CPA}}(\lambda, m, k) + q_{\mathcal{O}}/2^\gamma.$$

Proof. We proceed in a sequence of games. In the sequel, \mathcal{O}_i is the implementation of the decryption oracle in Game i .

Game 0. This is the original IND-CCA security experiment from Definition 8, played with the scheme described above. In particular, the decryption oracle \mathcal{O}_0 is implemented as follows (we omit the check for $C = C^*$):

$\mathcal{O}_0(C)$

```

 $K \leftarrow^{\$} \text{Dec}(\text{sk}, C)$ 
If  $K = \perp$  then return  $\perp$ 
 $(r, K') = R(K)$ 
If  $(C, K) \neq \text{Enc}(\text{pk}; (r, K))$  then return  $\perp$ 
Return  $K'$ 

```

Recall that K_0 denotes the encapsulated key computed by the IND-CCA experiment. K_0 is uniquely defined by the challenge ciphertext C^* via $K_0 := \text{Dec}(\text{sk}_0, C^*)$, where sk_0 is the initial (non-punctured) secret key, since the scheme satisfies extended correctness (Definition 4, second property). Let A_0 denote the event that \mathcal{A} ever queries K_0 to random oracle R . Note that \mathcal{A} has zero advantage in distinguishing K' from random, until A_0 occurs, because R is a random function. Thus, we have $\Pr[A_0] \geq \text{Adv}_{\mathcal{A}, \text{BFKEM}'}^{\text{IND-CCA}}(\lambda, m, k)$. In the sequel, we denote with A_i the event that \mathcal{A} ever queries K_0 to random oracle R in Game i .

Game 1. This game is identical to Game 0, except that after computing $K \leftarrow^{\$} \text{Dec}(\text{sk}, C)$ and checking whether $K \neq \perp$, the experiment additionally checks whether the adversary has ever queried random oracle R on input K , and returns \perp if not. More precisely, the experiment maintains a list

$$L_R = \{(K, (r, K')) : \mathcal{A} \text{ queried } R(K) = (r, K')\}$$

to record all queries K made by the adversary to random oracle R , along with the corresponding response $(r, K') = R(K)$. The decryption oracle \mathcal{O}_1 uses this list as follows (boxed statements highlight changes to \mathcal{O}_0):

$\mathcal{O}_1(C)$

```

 $K \leftarrow^{\$} \text{Dec}(\text{sk}, C)$ 
If  $\nexists (r, K') : (K, (r, K')) \in L_R$  then return  $\perp$ 
 $(r, K') = R(K)$ 
If  $(C, K) \neq \text{Enc}(\text{pk}; (r, K))$  then return  $\perp$ 
Return  $K'$ 

```

Note that Games 0 and 1 are perfectly indistinguishable, unless \mathcal{A} ever outputs a ciphertext C with $\mathcal{O}_1(C) = \perp$, but $\mathcal{O}_0(C) \neq \perp$. Note that this happens if and only if \mathcal{A} outputs C such that $C = \text{Enc}(\text{pk}; (r, K))$, where r is the randomness defined by $(r, K') = R(K)$, but without prior query of $R(K)$.

The random oracle R assigns a uniformly random value $r \in \{0, 1\}^\rho$ to each query, so, by the γ -spreadness of BFKEM, the probability that the ciphertext C output by the adversary “matches” the ciphertext produced by $\text{Enc}(\text{pk}; (r, K))$ is $2^{-\gamma}$. Since \mathcal{A} issues at most $q_{\mathcal{O}}$ queries to \mathcal{O}_1 , this yields $\Pr[A_1] \geq \Pr[A_0] - q_{\mathcal{O}}/2^\gamma$.

Game 2. We make a minor conceptual modification. Instead of computing $(r, K') = R(K)$ by evaluating R , \mathcal{O}_2 reads (r, K') from list L_R . More precisely:

$\mathcal{O}_2(C)$

$K \xleftarrow{\$} \text{Dec}(\text{sk}, C)$

If $\nexists (r, K') : (K, (r, K')) \in L_R$ **then return** \perp

Define (r, K') to be the unique tuple such that $(K, (r, K')) \in L_R$.

If $(C, K) \neq \text{Enc}(\text{pk}; (r, K))$ **then return** \perp

Return K'

By definition of L_R it always holds that $(r, K') = R(K)$ for all $(K, (r, K')) \in L_R$. Indeed (r, K') , is uniquely determined by K , because $(r, K') = R(K)$ is a function. Since R is only evaluated by \mathcal{O}_1 if there exists a corresponding tuple $(K, (r, K')) \in L_R$ anyway, due to the changes introduced in Game 1, oracle \mathcal{O}_2 is equivalent to \mathcal{O}_1 and we have $\Pr[A_2] = \Pr[A_1]$.

Game 3. This game is identical to Game 2, except that whenever \mathcal{A} queries a ciphertext C to oracle \mathcal{O}_3 , then \mathcal{O}_3 first runs the CheckPunct algorithm associated to BFKEM (cf. Definition 6). If $\text{CheckPunct}(\text{pk}, \mathcal{Q}, C) = \perp$, then it immediately returns \perp . Otherwise, it proceeds exactly like \mathcal{O}_2 . More precisely:

$\mathcal{O}_3(C)$

If $\text{CheckPunct}(\text{pk}, \mathcal{Q}, C) = \perp$ **then return** \perp

$K \xleftarrow{\$} \text{Dec}(\text{sk}, C)$

If $\nexists (r, K') : (K, (r, K')) \in L_R$ **then return** \perp

Define (r, K') to be the unique tuple such that $(K, (r, K')) \in L_R$.

If $(C, K) \neq \text{Enc}(\text{pk}; (r, K))$ **then return** \perp

Return K'

Recall that by public checkability (Definition 6) we have $\perp = \text{Dec}(\text{sk}, C) \iff \perp = \text{CheckPunct}(\text{pk}, \mathcal{Q}, C)$. Therefore the introduced changes are conceptual, and $\Pr[A_3] = \Pr[A_2]$.

Game 4. We modify the secret key used to decrypt the ciphertext. Let sk_0 denote the initial secret key generated by the experiment (that is, before any puncturing operation was performed). \mathcal{O}_4 uses sk_0 to compute $K \xleftarrow{\$} \text{Dec}(\text{sk}_0, C)$ instead of $K \xleftarrow{\$} \text{Dec}(\text{sk}, C)$, where sk is a possibly punctured secret key. More precisely:

$\mathcal{O}_4(C)$

If CheckPunct(pk, \mathcal{Q} , C) = \perp **then return** \perp
 $K \leftarrow^{\$} \text{Dec}(\text{sk}_0, C)$
If $\nexists (r, K') : (K, (r, K')) \in L_R$ **then return** \perp
 Define (r, K') to be the unique tuple such that $(K, (r, K')) \in L_R$.
If $(C, K) \neq \text{Enc}(\text{pk}; (r, K))$ **then return** \perp
Return K'

For indistinguishability from Game 3, we show that $\mathcal{O}_4(C) = \mathcal{O}_3(C)$ for all ciphertexts C . Let us first consider the case $\text{Dec}(\text{sk}, C) = \perp$. Then public checkability guarantees that $\mathcal{O}_4(C) = \mathcal{O}_3(C) = \perp$, due to the fact that $\text{Dec}(\text{sk}, C) = \perp \iff \text{CheckPunct}(\text{pk}, \mathcal{Q}, C) = \perp$.

Now let us consider the case $\text{Dec}(\text{sk}, C) \neq \perp$. In this case, the semi-correctness of punctured keys (3rd requirement of Definition 4) guarantees that $\text{Dec}(\text{sk}, C) = \text{Dec}(\text{sk}_0, C) = K \neq \perp$.

After computing $\text{Dec}(\text{sk}_0, C)$, \mathcal{O}_4 performs exactly the same operations as \mathcal{O}_3 after computing $\text{Dec}(\text{sk}, C)$. Thus, in this case both oracles are perfectly indistinguishable, too. This yields that the changes introduced in Game 4 are purely conceptual, and we have $\Pr[A_4] = \Pr[A_3]$.

Remark. Due to the fact that we are now using the initial secret key to decrypt C , we have reached a setting where, due to the perfect correctness of the initial secret key sk_0 , essentially a perfectly-correct encryption scheme is used – except that the decryption oracle implements a few additional abort conditions. Thus, we can now basically apply the standard Fujisaki-Okamoto transformation, but we must show that we are also able to simulate the additional abort imposed by the additional consistency checks properly. To this end, we first replace these checks with equivalent checks before applying the FO transformation.

Game 5. We replace the consistency checks performed by \mathcal{O}_4 with an equivalent check. More precisely, \mathcal{O}_5 works as follows:

$\mathcal{O}_5(C)$

If CheckPunct(pk, \mathcal{Q} , C) = \perp **then return** \perp
 $K \leftarrow^{\$} \text{Dec}(\text{sk}_0, C)$
If $\nexists (r, K') : ((K, (r, K')) \in L_R \wedge (C, K) = \text{Enc}(\text{pk}; (r, K)))$ **then return** \perp
Return K' such that $(K, (r, K')) \in L_R \wedge (C, K) = \text{Enc}(\text{pk}; (r, K))$

This is equivalent, so that we have $\Pr[A_5] = \Pr[A_4]$.

Game 6. Observe that in Game 5 we check whether there exists a tuple (r, K') with $(K, (r, K')) \in L_R$ and $(C, K) = \text{Enc}(\text{pk}; (r, K))$, where K must match the secret key computed by $K \leftarrow^{\$} \text{Dec}(\text{sk}_0, C)$.

In Game 6, we relax this check. We test only whether there exists any tuple $(\tilde{K}, (\tilde{r}, \tilde{K}')) \in L_R$ such that $(C, \tilde{K}) = \text{Enc}(\text{pk}; (\tilde{r}, \tilde{K}'))$ holds. Thus, it is not explicitly checked

whether \tilde{K} matches the value $K \stackrel{\$}{\leftarrow} \text{Dec}(\text{sk}_0, C)$. Furthermore, the corresponding value \tilde{K}' is returned. More precisely:

$\mathcal{O}_6(C)$

If $\text{CheckPunct}(\text{pk}, \mathcal{Q}, C) = \perp$ **then return** \perp
 $K \stackrel{\$}{\leftarrow} \text{Dec}(\text{sk}_0, C)$

If $\nexists(\tilde{r}, \tilde{K}') : ((\tilde{K}, (\tilde{r}, \tilde{K}')) \in L_R \wedge (C, \tilde{K}) = \text{Enc}(\text{pk}; (\tilde{r}, \tilde{K})))$ **then return** \perp

Return \tilde{K}' such that $(\tilde{K}, (\tilde{r}, \tilde{K}')) \in L_R \wedge (C, \tilde{K}) = \text{Enc}(\text{pk}; (\tilde{r}, \tilde{K}))$

By the perfect correctness of the initial secret key sk_0 , we have

$$(C, \tilde{K}) = \text{Enc}(\text{pk}; (\tilde{r}, \tilde{K})) \implies \text{Dec}(\text{sk}_0, C) = \tilde{K},$$

so that we must have $K = \tilde{K}$. \mathcal{O}_6 is equivalent to \mathcal{O}_5 , and $\Pr[A_6] = \Pr[A_5]$.

Game 7. This game is identical to Game 6, except that we change the decryption oracle again. Observe that the value K computed by $K \stackrel{\$}{\leftarrow} \text{Dec}(\text{sk}_0, C)$ is never used by \mathcal{O}_6 . Therefore the computation of $K \stackrel{\$}{\leftarrow} \text{Dec}(\text{sk}_0, C)$ is obsolete, and we can remove it. More precisely, \mathcal{O}_7 works as follows.

$\mathcal{O}_7(C)$

If $\text{CheckPunct}(\text{pk}, \mathcal{Q}, C) = \perp$ **then return** \perp

If $\nexists(\tilde{r}, \tilde{K}') : ((\tilde{K}, (\tilde{r}, \tilde{K}')) \in L_R \wedge (C, \tilde{K}) = \text{Enc}(\text{pk}; (\tilde{r}, \tilde{K})))$ **then return** \perp

Return \tilde{K}' such that $(\tilde{K}, (\tilde{r}, \tilde{K}')) \in L_R \wedge (C, \tilde{K}) = \text{Enc}(\text{pk}; (\tilde{r}, \tilde{K}))$

We have only removed an obsolete instruction, which does not change the output distribution of the decryption oracle. Therefore \mathcal{O}_7 simulates \mathcal{O}_6 perfectly, and we have $\Pr[A_7] = \Pr[A_6]$.

Reduction to OW-CPA-security. Now we are ready to describe the OW-CPA-adversary \mathcal{B} . \mathcal{B} receives (pk, C^*) . It samples a uniformly random key $K' \stackrel{\$}{\leftarrow} \{0, 1\}^\lambda$ and runs the IND-CCA-adversary \mathcal{A} as a subroutine on input (pk, C^*, K') . Whenever \mathcal{A} issues a Punc- or Corr-query, then \mathcal{B} forwards this query to the OW-CPA-experiment and returns the response. In order to simulate the decryption oracle \mathcal{O} , adversary \mathcal{B} implements the simulated oracle \mathcal{O}_7 from Game 7 described above. When \mathcal{A} terminates, then \mathcal{B} picks a uniformly random entry $(\hat{K}, (\hat{r}, \hat{K}')) \stackrel{\$}{\leftarrow} L_R$, and outputs \hat{K} .

Analysis of the reduction. Let \hat{Q} denote the event that \mathcal{A} ever queries K_0 to random oracle R . Note that \mathcal{B} simulates Game 7 perfectly until A_7 occurs, thus we have $\Pr[\hat{Q}] \geq \Pr[A_7]$. Summing up, the probability that the value \hat{K} output by \mathcal{B} matches the key encapsulated in C^* is therefore at least

$$\frac{\Pr[\hat{Q}]}{q_R} \geq \frac{\text{Adv}_{\mathcal{A}, \text{BFKEM}'}^{\text{IND-CCA}}(\lambda, m, k) - q_{\mathcal{O}}/2^\gamma}{q_R}.$$

□

Remark on the tightness. Alternatively, we could have based the security of our IND-CCA-secure scheme on the IND-CPA (rather than OW-CPA) security of BFKEM' . In this case, we would have achieved a tighter reduction, as we would have been able to avoid guessing the index $(\hat{K}, (\hat{r}, \hat{K}')) \xleftarrow{\$} L_R$, at the cost of requiring stronger security of the underlying scheme.

From IND-CCA-secure KEMs to IND-CCA-secure encryption. It is well-known that IND-CCA-secure KEMs can be generically transformed into IND-CCA-secure encryption schemes, by combining it with a CCA-secure symmetric encryption scheme [25]. This construction applies to BFKEMs as well.

3.3 BFKEM from CP-ABE

We now present an alternative, generic construction of a BFKEM from ciphertext-policy attribute-based encryption (CP-ABE) [8]. In particular, the construction can be instantiated with any small-universe (i.e., bounded) CP-ABE scheme⁶ that is adaptively secure, supports at least OR-policies, and allows to encrypt messages from an exponentially large space. We note that since the formulation of KEMs in context of ABE is not widely used, we opt to start from a CP-ABE scheme which we implicitly turn into a KEM in the construction via the folklore compiler to obtain KEMs from encryptions schemes.

In contrast to the basic BFKEM construction in Section 3.1, we are able to generically obtain constant-size ciphertexts (independent of the parameters m and k) if the underlying CP-ABE scheme beyond possessing the aforementioned properties, is also compact, i.e., provides constant-size ciphertexts, (as e.g. [18] and [2] which are obtained from static and parameterized assumptions, respectively). Compact-size ciphertexts come at the cost of increased secret key size in existing schemes (at least quadratic in the number of attributes). However, for forward-secret 0-RTT key-exchange storage cost at the server is less expensive than communication bandwidth and thus can be considered a viable trade-off.

CP-ABE. Before we describe our construction let us briefly recall CP-ABE. Therefore, let \mathbb{U} be the universe of attributes and we require only small-universe constructions, i.e., \mathbb{U} is fixed at setup and $|\mathbb{U}|$ is polynomially bounded in the security parameter λ (in our BFKEM construction we will have $|\mathbb{U}| = m$). Intuitively, in a CP-ABE scheme secret keys are issued with respect to attribute sets $\mathbb{U}' \subseteq \mathbb{U}$ and messages are encrypted with respect to access structures (policies) defined over \mathbb{U} . Decryption works iff the attributes in the secret key satisfy the policy used to produce the ciphertext. Let us discuss this a bit more formally.

Definition 11 (Access Structure [8]). Let \mathbb{U} be the attribute universe. A collection $\mathbb{A} \in 2^{\mathbb{U}}$ of non-empty sets is an access structure on \mathbb{U} . The sets in \mathbb{A} are called the authorized sets, and the sets not in \mathbb{A} are called the unauthorized sets. A collection $\mathbb{A} \in 2^{\mathbb{U}}$ is called monotone if $\forall B, C \in \mathbb{A} : \text{if } B \in \mathbb{A} \text{ and } B \subseteq C, \text{ then } C \in \mathbb{A}$.

⁶ Note that any large universe CP-ABE scheme yields a small-universe CP-ABE scheme but not vice-versa.

Subsequently, we do not require arbitrary monotone access structures, but only OR-policies (i.e., threshold policies with threshold 1). In particular, for some attribute set $\mathbb{U}' := (u_1, \dots, u_n) \subseteq \mathbb{U}$ we consider policies of the form $u_1 \text{ OR } \dots \text{ OR } u_n$, representing an access structure $\mathbb{A} := 2^{\mathbb{U}'} \setminus \emptyset$.

Definition 12 (CP-ABE). A ciphertext-policy attribute-based encryption scheme is a tuple $\text{CP-ABE} = (\text{Setup}, \text{KGen}, \text{Enc}, \text{Dec})$ of PPT algorithms:

$\text{Setup}(1^\lambda, \mathbb{U})$: Takes as input a security parameter λ and an attribute universe description \mathbb{U} and outputs a master secret and public key (msk, mpk) . We assume that all subsequent algorithms will implicitly receive the master public key mpk (public parameters) as input which implicitly fixes a message space \mathcal{M} .

$\text{KGen}(\text{msk}, \mathbb{U}')$: Takes as input the master secret key msk and a set of attributes $\mathbb{U}' \subseteq \mathbb{U}$ and outputs a secret key $\text{sk}_{\mathbb{U}'}$.

$\text{Enc}(M, \mathbb{A})$: Takes as input a message $M \in \mathcal{M}$ and an access structure \mathbb{A} and outputs a ciphertext C .

$\text{Dec}(\text{sk}_{\mathbb{U}'}, C)$: Takes as input a secret key $\text{sk}_{\mathbb{U}'}$ and a ciphertext C and outputs a message M or \perp in case of decryption does not work.

Correctness of CP-ABE requires that for all λ , all attribute sets \mathbb{U} , all $(\text{msk}, \text{mpk}) \xleftarrow{\$} \text{Setup}(1^\lambda, \mathbb{U})$, all $M \in \mathcal{M}$, all $\mathbb{A} \in 2^{\mathbb{U}} \setminus \emptyset$, all $\mathbb{U}' \in \mathbb{A}$, all $\text{sk}_{\mathbb{U}'} \xleftarrow{\$} \text{KGen}(\text{msk}, \mathbb{U}')$ we have that $\Pr[\text{Dec}(\text{sk}_{\mathbb{U}'}, \text{Enc}(M, \mathbb{A})) = M] = 1$.

Security of CP-ABE. Figure 4 defines adaptive IND-T with $T \in \{\text{CPA}, \text{CCA}\}$ security for CP-ABE. We stress that we use a formalization for small-universe schemes where the size of \mathbb{U} is polynomially bounded in the security parameter λ (for large universe \mathbb{U} is not required for Setup). We denote this value by n and consider the attribute set to be $\mathbb{U} = \{1, \dots, n\}$.

$\text{Exp}_{\mathcal{A}, \text{CP-ABE}}^{\text{IND-T}}(\lambda, n)$:

$(\text{msk}, \text{mpk}) \xleftarrow{\$} \text{Setup}(1^\lambda, \mathbb{U})$

$b \xleftarrow{\$} \{0, 1\}, \mathcal{Q} \leftarrow \emptyset$

$(M_0, M_1, \mathbb{A}^*) \xleftarrow{\$} \mathcal{A}^{\mathcal{O}, \text{KGen}(\text{msk}, \cdot)}(\text{mpk})$

where $\mathcal{O} \leftarrow \text{Dec}(\cdot, \cdot)$ if $T = \text{CCA2}$ and $\mathcal{O} \leftarrow \emptyset$ otherwise.

$\text{KGen}(\text{msk}, \mathbb{U}')$ returns $\text{sk}_{\mathbb{U}'}$ and sets $\mathcal{Q} \leftarrow \mathcal{Q} \cup \mathbb{U}'$

if $M_0, M_1 \notin \mathcal{M} \vee |M_0| \neq |M_1| \vee \mathbb{A}^* \cap \mathcal{Q} \neq \emptyset$, let $C^* \leftarrow \perp$

else, let $C^* \xleftarrow{\$} \text{Enc}(M_b, \mathbb{A}^*)$

$b^* \xleftarrow{\$} \mathcal{A}^{\mathcal{O}, \text{KGen}(\text{msk}, \cdot)}(C^*)$

where $\mathcal{O} \leftarrow \text{Dec}'(\cdot, \cdot)$ if $T = \text{CCA2}$ and $\mathcal{O} \leftarrow \emptyset$ otherwise.

$\text{Dec}'(\mathbb{U}', C)$ returns $\text{Dec}(\text{KGen}(\text{msk}, \mathbb{U}'), C)$ if $C \neq C^*$

and \perp otherwise.

$\text{KGen}(\text{msk}, \mathbb{U}')$ returns $\text{sk}_{\mathbb{U}'}$ if $\mathbb{U}' \notin \mathbb{A}^*$ and \perp otherwise

return 1, if $b^* = b$

return 0

Fig. 4. IND-T security for small-universe CP-ABE: $T \in \{\text{CPA}, \text{CCA}\}$.

Definition 13 (IND-T Security of CP-ABE). We define the advantage of an adversary A in the IND-T experiment $\text{Exp}_{A, \text{CP-ABE}}^{\text{IND-T}}(\lambda, n)$ as

$$\text{Adv}_{A, \text{CP-ABE}}^{\text{IND-T}}(\lambda, n) := \left| \Pr \left[\text{Exp}_{A, \text{CP-ABE}}^{\text{IND-T}}(\lambda, n) = 1 \right] - \frac{1}{2} \right|.$$

A ciphertext-policy attribute-based encryption scheme CP-ABE is IND-T, $T \in \{\text{CPA}, \text{CCA}\}$, secure, if $\text{Adv}_{A, \text{CP-ABE}}^{\text{IND-T}}(\lambda, n)$ is a negligible function in λ for all $n > 0$ and all PPT adversaries A .

Intuition of the BFKEM construction. The intuition of constructing a CPA-secure BFKEM from CP-ABE is very simple. Basically, we map the indices m in $T \in \{0, 1\}^m$ of a Bloom filter (H, T) to the attribute universe \mathbb{U} . Then we generate for every attribute $i \in [m]$ (we consider $\mathbb{U} = \{1, \dots, m\}$) a secret key $\text{sk}_{\{i\}}$, set our secret key of the BFKEM scheme to be $\text{sk} := (T, (\text{sk}_{\{1\}}, \dots, \text{sk}_{\{m\}}))$ and delete msk . Encryption is with respect to the attributes given by the indices \mathcal{I} obtained from sending a randomly sampled tag r through the hash functions H_j , $j \in [k]$ of the Bloom filter. Decryption works by using one secret key $\text{sk}_{\{i\}}$ indexed by \mathcal{I} . Puncturing a ciphertext simply amounts to discarding all the secret keys $\text{sk}_{\{i\}}$ indexed by \mathcal{I} .

Construction. Subsequently, we describe the generic CPA-secure BFKEM construction from a CP-ABE scheme ABE. We, thereby, require a CP-ABE with exponentially large message space \mathcal{M} , and assume that the key space \mathcal{K} of the BFKEM scheme is equivalent to \mathcal{M} .

KGen $(1^\lambda, m, k)$: Runs $((H_j)_{j \in [k]}, T) \leftarrow^{\$} \text{BFGen}(m, k)$. Then it runs $(\text{msk}, \text{mpk}) \leftarrow^{\$} \text{ABE.Setup}(1^\lambda, [m])$, and for all $i \in [m]$: $\text{sk}_{\{i\}} \leftarrow^{\$} \text{ABE.KGen}(\text{msk}, \{i\})$. Finally it sets and outputs

$$\text{sk} := (T, (\text{sk}_{\{i\}})_{i \in [m]}) \text{ and } \text{pk} := (\text{mpk}, (H_j)_{j \in [k]}).$$

Enc (pk) : Takes as input a public key pk . It samples uniformly at random a key $K \leftarrow^{\$} \mathcal{M}$, as well as a value $r \leftarrow^{\$} \{0, 1\}^\lambda$, computes $\forall j \in [k]$: $i_j = H_j(r)$, sets $\mathbb{U}' = \{i_1, \dots, i_k\}$ and $\mathbb{A} = 2^{\mathbb{U}'} \setminus \emptyset$. Finally, it computes $C' \leftarrow^{\$} \text{ABE.Enc}(K, \mathbb{A})$ and outputs (C, K) where ciphertext $C := (r, C')$.

Remark. We remark that if a CP-ABE is used where \mathcal{K} and \mathcal{M} are different, one can use standard randomness extraction techniques to extract a key $k \in \mathcal{K}$ from a uniformly random message $m \in \mathcal{M}$.

Punc (sk, C) : Takes as input a secret key $\text{sk} := (T, (\text{sk}_{\{i\}})_{i \in [m]})$ and ciphertext $C := (r, C')$. It computes $T' \leftarrow^{\$} \text{BFUpdate}((H_j)_{j \in [k]}, T, r)$ and for each $i \in [m]$ it defines

$$\text{sk}'_{\{i\}} := \begin{cases} \text{sk}_{\{i\}} & \text{if } T'[i] = 0, \text{ and} \\ \perp & \text{if } T'[i] = 1, \end{cases}$$

where $T'[i]$ denotes the i -th bit of T' . Finally, it returns an updated secret key $\text{sk}' = (T', (\text{sk}'_{\{i\}})_{i \in [m]})$.

$\text{Dec}(\text{sk}, C)$: Takes as input a secret key sk and a ciphertext $C := (r, C')$. It computes $\forall j \in [k] : i_j = H_j(r)$ and takes the first element $sk_{\{i_j\}}$ from $(sk_{\{i\}})_{i \in [m]}$ with $sk_{\{i_j\}} \neq \perp$. If such an $sk_{\{i_j\}}$ exists it outputs $K \leftarrow^{\$} \text{ABE.Dec}(sk_{\{i_j\}}, C')$ and \perp otherwise.

Correctness error of this scheme. Under the same argumentation as in the correctness proof in Section 3.1, we obtain that the correctness error is approximately $2^{-k} + n/2^\lambda$.

CPA security. We directly relate the CPA-security of our construction to the hardness of breaking CPA-security for the underlying CP-ABE.

Theorem 4. *From each efficient adversary \mathcal{B} against CPA-security of our BFKEM, we can construct an efficient adversary \mathcal{A} which breaks CPA-security of the underlying CP-ABE, with*

$$\text{Adv}_{\mathcal{A}, \text{CP-ABE}}^{\text{IND-CPA}}(\lambda, n) \geq \text{Adv}_{\mathcal{B}, \text{BFKEM}}^{\text{IND-CPA}}(\lambda, m, k).$$

Proof. We present a reduction which uses an adversary \mathcal{B} against CPA-security of the BFKEM to break CPA-security of the CP-ABE. First, we engage with a CPA challenger for a CP-ABE with respect to universe $[m]$ to obtain mpk . Then we complete the setup by running the following KeyGen' algorithm and obtain pk :

$\text{KeyGen}'(\text{mpk}, m, k)$: Runs $((H_j)_{j \in [k]}, T) \leftarrow^{\$} \text{BFGen}(m, k)$, sets

$$\text{pk} := (\text{mpk}, (H_j)_{j \in [k]}),$$

and outputs pk .

Then, we choose $(K_0, K_1) \leftarrow^{\$} \mathcal{M} \times \mathcal{M}$, $r \leftarrow^{\$} \{0, 1\}^\lambda$, and compute $\forall j \in [k] : i_j = H_j(r)$, set $\mathbb{U}' = \{i_1, \dots, i_k\}$, let $\mathbb{A} = 2^{\mathbb{U}'} \setminus \emptyset$. We output (K_0, K_1, \mathbb{A}) to the challenger to obtain C'^* . We start \mathcal{B} on $(\text{pk}, (r, C'^*), K_0)$ and simulate the oracles as follows:

$\text{Punc}(\text{sk}, C)$: Set $\mathcal{P} \leftarrow \mathcal{P} \cup \{C\}$, and $T \leftarrow \text{BFUpdate}((H_j)_{i \in [k]}, T, r)$.

Corr : If $C^* \notin \mathcal{P}$ return \perp . Otherwise, $\forall j \in [k] : i_j = T[j]$, and, for all $i_j = 0$ obtain $\text{sk}_j \leftarrow \text{KGen}(j)$ using the key generation oracle provided by the challenger and return $\text{sk} \leftarrow (T, \{\text{sk}_j\}_{j \in [k], i_j=0})$.

If \mathcal{B} eventually outputs a bit b^* we output b^* to break CPA-security of the CP-ABE scheme with the same probability as \mathcal{B} breaks the CPA-security of the BFKEM. Note that the Corr oracle can only be called after the challenge ciphertext C^* , and, therefore r , is determined. This ensures that we only request "allowed" keys via the KGen oracle provided by the challenger.

Obtaining CCA-security. The construction satisfies the additional properties of Definitions 4, 5, and 6 with the same arguments as in Section 3.1. Additionally, γ -spreadness (Def. 7) is given by construction: The randomness r is chosen uniformly at random from $\{0, 1\}^\lambda$. Thus, we can apply the Fujisaki-Okamoto [25] transform the same way as done in Section 3.2 to achieve CCA security.

3.4 BFKEM from IBBE

In this section, we present our generic construction of a BFKEM from any identity-based broadcast encryption (IBBE) scheme. We note that taking the path via IBBE allows us to simultaneously obtain small ciphertexts and small public keys.

Identity-Based Broadcast Encryption. We recall the basic definition of IBBE and its security.

Definition 14 (IBBE). *An identity-based broadcast encryption (IBBE) scheme is a tuple $\text{IBBE} = (\text{Setup}, \text{Extract}, \text{Enc}, \text{Dec})$ consisting of four probabilistic polynomial-time algorithms with the following properties:*

$\text{Setup}(1^\lambda, k)$: *Takes as input the security parameter λ and the maximal number of receivers k and outputs a master public key pk and a master secret key msk . We assume that pk implicitly defines the identity space \mathcal{ID} .*

$\text{Extract}(\text{msk}, \text{ID}_i)$: *Takes as input the master secret key msk and an user identity ID_i and outputs and user private key sk_{ID_i} .*

$\text{Enc}(\text{pk}, \mathcal{S})$: *Takes as input the master public key pk and a set of user identities \mathcal{S} and outputs a ciphertext C and a key K .*

$\text{Dec}(\text{sk}_{\text{ID}_i}, \mathcal{S}, C)$: *Takes as input a user secret key sk_{ID_i} , a set of user identities \mathcal{S} and a ciphertext C and outputs the key K .*

Correctness for IBBE requires that for all λ , for all polynomially bounded k in λ , for all $(\text{pk}, \text{msk}) \leftarrow^{\$} \text{Setup}(1^\lambda, k)$, for all $\mathcal{S} = \{\text{ID}_1, \dots, \text{ID}_i\} \in \mathcal{ID}^i$ with $i \leq k$, for all $(C, K) \leftarrow^{\$} \text{Enc}(\text{pk}, \mathcal{S})$, it holds for all $\text{ID}_S \in \mathcal{S}$ that

$$\Pr[\text{Dec}(\text{Extract}(\text{msk}, \text{ID}_S), \mathcal{S}, C) = K] = 1.$$

$\mathbf{Exp}_{\mathcal{A}, \text{IBBE}}^{\text{IND-sID-CPA}}(\lambda, k)$
 $S^* = \{\text{ID}_1^*, \dots, \text{ID}_s^*\} \leftarrow^{\$} \mathcal{A}(1^\lambda)$
 $(\text{pk}, \text{sk}) \leftarrow^{\$} \text{Setup}(1^\lambda, k)$
 $(C^*, K_0) \leftarrow^{\$} \text{Enc}(\text{pk}, S^*)$
 $K_1 \leftarrow^{\$} \mathcal{K}, b \leftarrow^{\$} \{0, 1\}$
 $b^* \leftarrow^{\$} \mathcal{A}^{\text{Extract}(\text{sk}, \cdot)}(\text{pk}, C^*, K_b)$
 $\text{Extract}(\text{sk}, j)$ returns $\text{sk}_{\text{ID}_j} \leftarrow^{\$} \text{Extract}(\text{sk}, j)$ if $j \notin S^*$ and \perp otherwise.
 return 1, if $b^* = b$
 return 0

Fig. 5. IND-sID-CPA security for IBBE.

Definition 15 (IND-sID-CPA-security of IBBE). *We define the advantage of an adversary \mathcal{A} in the IND-sID-CPA experiment $\mathbf{Exp}_{\mathcal{A}, \text{IBBE}}^{\text{IND-sID-CPA}}(\lambda, k)$ as*

$$\mathbf{Adv}_{\mathcal{A}, \text{IBBE}}^{\text{IND-sID-CPA}}(\lambda, k) := \left| \Pr \left[\mathbf{Exp}_{\mathcal{A}, \text{IBBE}}^{\text{IND-sID-CPA}}(\lambda, k) = 1 \right] - \frac{1}{2} \right|.$$

We say that an identity-based broadcast encryption scheme IBBE is IND-sID-CPA-secure, if the advantage $\text{Adv}_{\mathcal{A}, \text{IBBE}}^{\text{IND-sID-CPA}}(\lambda, k)$ is a negligible function in λ for all $k > 0$ and all PPT adversaries \mathcal{A} .

Construction. Let $B = (\text{BGen}, \text{BUpdate}, \text{BCheck})$ be a Bloom filter and let IBBE = (Setup, Extract, Enc, Dec) be an identity-based broadcast encryption scheme. We construct a Bloom Filter Key Encapsulation Mechanism BFKEM = (KGen, Enc, Punc, Dec) as follows:

KGen(λ, m, k) : The key generation algorithm generates a Bloom filter instance by running $(H, T) \leftarrow^{\$} \text{BGen}(m, k)$ and generates an IBBE instance by invoking $(\text{pk}_{\text{IBBE}}, \text{msk}) \leftarrow^{\$} \text{IBBE.Setup}(\lambda, k)$. For each $i \in [m]$ it calls

$$\text{sk}_i \leftarrow^{\$} \text{IBBE.Extract}(\text{msk}, i).$$

Finally, it sets

$$\text{pk} := (H, \text{pk}_{\text{IBBE}}) \text{ and } \text{sk} := (T, (\text{sk}_i)_{i \in [m]}).$$

Remark. Observe that the maximum number of recipients is set to the Bloom filter's optimal number of universal hash functions k and the user identity space is bound to the Bloom filter's entries m .

Enc(pk) : Given a public key $\text{pk} = (H, \text{pk}_{\text{IBBE}})$, it samples a random value $r \leftarrow^{\$} \{0, 1\}^\lambda$ and generates indices $i_j := H_j(r)$ for $(H_j)_{j \in [k]} := H$. Then it invokes $(K, C') \leftarrow^{\$} \text{IBBE.Enc}(\text{pk}_{\text{IBBE}}, \mathcal{S})$, where $\mathcal{S} := \{i_j\}_{j \in [k]}$. Finally, it outputs (C, K) , where ciphertext $C := (r, C')$.

Punc(sk, C) : Given a secret key $\text{sk} = (T, (\text{sk}_i)_{i \in [m]})$ and a ciphertext $C = (r, C')$, it invokes $T' = \text{BUpdate}(H, T, r)$ and defines

$$\text{sk}'_i := \begin{cases} \text{sk}_i, & \text{if } T'[i] = 0 \\ \perp, & \text{if } T'[i] = 1. \end{cases}$$

Finally, the algorithm returns $\text{sk}' = (T', (\text{sk}'_i)_{i \in [m]})$.

Remark. From an IBBE's point of view, the puncturing procedure removes participants from the broadcast network by deleting their respective user private keys.

Dec(sk, C) : The input is a secret key $\text{sk} = (T, (\text{sk}_i)_{i \in [m]})$ and ciphertext $C = (r, C')$. Again, let $\mathcal{S} := \{i_j\}_{j \in [k]}$. If $\text{BCheck}(H, T, r) = 0$, then the algorithm returns \perp . Else, there exists at least one index $n \in \mathcal{S}$ such that $\text{sk}_n \neq \perp$. The algorithm picks the smallest index n that meets the previous requirements, computes

$$K := \text{IBBE.Dec}(\text{sk}_n, \mathcal{S}, C')$$

and returns K .

Remark. This algorithm essentially checks, if an user secret key of the user identities in set \mathcal{S} still exists. If so, the ciphertext can be decrypted.

Correctness error. With exactly the same arguments as for the scheme from Section 3.1, one can verify that the correctness error of this scheme is essentially identical to the false positive probability of the Bloom filter, unless a given ciphertext $C = (r, C')$ has a value of r which is identical to the value of r of any previous ciphertext. Since r is uniformly random in $\{0, 1\}^\lambda$, this probability is approximately $2^{-k} + n \cdot 2^{-\lambda}$.

IND-CPA-security. We prove the IND-CPA security of our construction, if the IBBE is IND-sID-CPA-secure.

Theorem 5. *From each efficient adversary \mathcal{B} against IND-CPA-security of our BFKEM, we can construct an efficient algorithm \mathcal{A} against the IND-sID-CPA-security of the underlying IBBE scheme with advantage*

$$\text{Adv}_{\mathcal{A}, \text{IBBE}}^{\text{IND-sID-CPA}}(\lambda, k) \geq \text{Adv}_{\mathcal{B}, \text{BFKEM}}^{\text{IND-CPA}}(\lambda, m, k).$$

Proof. We proceed by presenting a reduction which uses an adversary \mathcal{B} against the IND-CPA-security of the BFKEM to break the IND-sID-CPA-security of the IBBE. The reduction together with \mathcal{B} then forms \mathcal{A} . In order to engage with the IND-CPA Challenger (\mathcal{C} henceforth), we need to commit to a set of recipients \mathcal{S}^* we will attack.

We generate a new Bloom filter instance by invoking $(H, T) \leftarrow^{\$} \text{BFGen}(m, k)$ and sample an additional random value $r^* \leftarrow^{\$} \{0, 1\}^\lambda$. Next, we compute indices $i_j := H_j(r^*)$ where $(H_j)_{j \in [k]} := H$ are the k universal hash functions of the Bloom filter. We define $\mathcal{S}^* := \{i_j\}_{j \in [k]}$ and forward the set to \mathcal{C} . Note that $|\mathcal{S}| = k$.

The challenger \mathcal{C} generates a master public key pk and a master secret key msk by invoking $\text{IBBE.Setup}(\lambda, k)$ and sends us the master public key pk . Additionally, \mathcal{C} prepares a challenge by running $(C', K_0) \leftarrow^{\$} \text{IBBE.Enc}(\text{pk}, \mathcal{S}^*)$ and sampling $K_1 \leftarrow^{\$} \mathcal{K}$, where \mathcal{K} is the symmetric key space. The challenger sends us the challenge (C', K_b) , where b is a bit drawn uniformly at random.

We will initialize the adversary \mathcal{B} with input $(\text{pk}, C^* = (r^*, C'), K_b)$. In the sequel \mathcal{B} has access to several oracles, which we simulate as follows:

- $\text{Punct}(C = (r, C'))$: We invoke $T := \text{BFUpdate}(H, T, r)$ and set $\mathcal{Q} := \mathcal{Q} \cup \{C\}$.
- Corr : If $C^* \notin \mathcal{Q}$, return \perp . Else query $\text{sk}_j := \text{Extract}(j)$ for all $j \in [k]$ such that $T[j] = 0$. Note, that we are allowed to call Extract on all user identities, since puncturing at C^* removes all troublesome secret keys. We return $(T, \{\text{sk}_j\}_{j \in [k] \wedge T[j]=0})$ to \mathcal{A} .

Eventually, \mathcal{B} will output a bit b^* which we will forward to the challenger \mathcal{C} . Since all queries are perfectly simulated, we get

$$\text{Adv}_{\mathcal{A}, \text{IBBE}}^{\text{IND-sID-CPA}}(\lambda, k) \geq \text{Adv}_{\mathcal{B}, \text{BFKEM}}^{\text{IND-CPA}}(\lambda, m, k).$$

This concludes the proof. □

CCA security. IND-CCA security can be achieved with the modified Fujisaki–Okamoto transformation described in Section 3.2. The IBBE-based construction satisfies the additional properties of Definitions 4, 5, and 6 with the same arguments as in Section 3.1. Additionally, γ -spreadness (Def. 7) is given by construction: The randomness r is chosen uniformly at random from $\{0, 1\}^\lambda$.

Separable randomness is not achieved as the symmetric key K algebraically depends on the IBBE (i.e., the symmetric key K is chosen by the IBBE and not by the generic construction).⁷ It is, however, possible to transform any non-separable BFKEM into a separable BFKEM as shown in Section 2.3. Note that this transformation adds an additional component of size λ to the ciphertext.

Thus, we can apply the Fujisaki-Okamoto transform the same way as done in Section 3.2 to achieve CCA security. One notable drawback is that the transformation requires that the encapsulation procedure be run once during each decapsulation. Should the encapsulation procedure be computationally expensive and should the application strive for high efficiency, it might be worth considering a different approach for achieving IND-CCA security.

A different approach to achieve IND-CCA security for our construction would be to directly use an IND-sID-CCA-secure IBBE. This can for example be achieved by using a variant of the CHK transformation [16] sketched in [22]. The basic idea is to derive one of the broadcasted identities from a verification key of a strongly unforgeable one-time signature (sOTS) scheme, which then in turn is used to sign the ciphertext. A reasonable choice for the signature scheme might be the Boneh–Lynn–Shacham signature scheme [13], which is strongly unforgeable due to its unique ciphertexts, or the sOTS due to Groth [32] which avoids pairing evaluations. Drawbacks of the transformation include an expansion of the ciphertext as both the signature verification key and the signature must be included. A formal description and security proof of the transformation can be found in [26].

4 Time-Based Bloom Filter Encryption

For a standard BFKEM scheme, we have to update the public key after the secret key has been punctured n -times, because otherwise the false-positive probability would exceed an acceptable bound. In this section, we describe a construction of a scheme where the lifetime of the public key is split into *time slots*. Ciphertexts are associated with time slots, which assumes loosely synchronized clocks between sender and receiver of a ciphertext. The main advantage is that for a given bound on the correctness error, we are able to handle about the same number of puncturings *per time slot* as the basic scheme during the entire life time of the public key. We call this approach *time-based* Bloom filter encryption. It is inspired by the time-based approach used to construct puncturable encryption in [31,33], which in turn is inspired by the construction of forward-secure public-key encryption by Canetti, Halevi, and Katz [15].

Note that a time-based BFKEM (TB-BFKEM) scheme can trivially be obtained from any BFE scheme, by assigning an individual public/secret key pair for each time slot. However, if we want to split the life time of the public key into, say, 2^t time slots, then this would of course increase the size of keys by a factor 2^t . Since we want

⁷ Depending on the instantiation, it might still be possible to directly achieve separable randomness. For this to work, the IBBE would need *separable keys*, that is, if we can equivalently write $(K, C) \stackrel{\$}{\leftarrow} \text{IBBE.Enc}(\text{mpk}, \mathcal{S}) = \text{IBBE.Enc}'(\text{mpk}, \mathcal{S}; K)$ for uniformly random $K \stackrel{\$}{\leftarrow} \{0, 1\}^\lambda$, where Enc' is a deterministic algorithm. This property is not necessarily given for IBBEs.

to enable a fine-grained use of time slots, to enable a very large number of puncturings over the entire lifetime of the public key without increasing the false positive probability beyond an unacceptable bound, we want to have 2^t as large as possible, but without increasing the size of the public key beyond an acceptable bound. To this end, we give a direct construction which increases the size of secret keys only by an *additive* amount of additional group elements, which is only *logarithmic* in the number of time slots. Thus, for 2^t time slots we have to add merely about t elements to the secret key, while the size of public keys remains even *constant*. Recall also, that due to the time slots, a TB-BFKEM helps to counter message suppression attacks by achieving a form of delayed forward secrecy.

4.1 Formal Model of TB-BFKEM

Likewise to considering our BFKEMs as an instantiation of a puncturable KEM with non-negligible correctness error, we can view the time-based approach analogously as an instantiation of a forward-secret BFKEM [33] with non-negligible correctness error, henceforth referred to as TB-BFKEM. We chose to align our model with the existing formal framework for puncturable forward-secret KEMs. It is essentially our BFKEM Definition 2, augmented by time slots and an additional algorithm Punctnt that allows to puncture a secret key not with respect to a given ciphertext in a given time slot, but with respect to an entire time slot.

Definition 16 (TB-BFKEM). *A puncturable forward-secret key encapsulation (TB-BFKEM) scheme is a tuple of the following PPT algorithms:*

$\text{KGen}(1^\lambda, m, k, t)$: *Takes as input a security parameter λ , parameters m and k for the Bloom filter, and a parameter t specifying the number of time slots. It outputs a secret and public key (sk, pk) , where we assume that the key-space \mathcal{K} is implicit in pk and that pk is implicit in sk .*

$\text{Enc}(\text{pk}, \tau)$: *Takes as input a public key pk and a time slot τ and outputs a ciphertext C and a symmetric key K .*

$\text{PuncCtx}(\text{sk}, \tau, C)$: *Takes as input a secret key sk , a time slot τ , a ciphertext C and outputs an updated secret key sk' .*

$\text{Dec}(\text{sk}, \tau, C)$: *Takes as input a secret key sk , a time slot τ , a ciphertext C and deterministically computes and outputs a symmetric key K or \perp if decapsulation fails.*

$\text{Punctnt}(\text{sk}, \tau)$: *Takes as input a time slot τ and a secret key sk for any time slot $\leq \tau$, and outputs an updated secret key sk' for the time slot $\tau + 1$.*

Correctness. Essentially, the correctness definition is based on that of a BFKEM, but additionally considers time slots (see also [33]).

Definition 17 (Correctness). *We require that the following holds for all $\lambda, m, k, t \in \mathbb{N}$, for all $z \in \mathbb{N}$ with $z \leq t$, and any $(\text{sk}, \text{pk}) \leftarrow^{\$} \text{KGen}(1^\lambda, m, k, t)$.*

- For any ordered sequence τ_1, \dots, τ_z with $1 \leq \tau_1 < \dots < \tau_z \leq t$ and

$$\text{sk}_{i+1} \leftarrow^{\$} \text{Punctnt}(\text{sk}_i, \tau_i)$$

where $i \in \{1, \dots, z\}$ and $\text{sk}_1 := \text{sk}$, and

– for any arbitrary interleaved sequence of invocations of

$$\text{sk}_{z+1,j+1} \leftarrow^{\$} \text{PuncCtx}(\text{sk}_{z+1,j}, \tau_z + 1, C_j)$$

where $j \in \{1, \dots, n\}$, $\text{sk}_{z+1,1} := \text{sk}_{z+1}$, and $(C_j, K_j) \leftarrow^{\$} \text{Enc}(\text{pk}, \tau_z + 1)$

it holds that

$$\Pr[\text{Dec}(\text{sk}_{z+1,n+1}, \tau_z + 1, C^*) \neq K^*] \leq \left(1 - e^{-\frac{(n+1/2)k}{m-1}}\right)^k + \epsilon(\lambda)$$

where $(C^*, K^*) \leftarrow^{\$} \text{Enc}(\text{pk}, \tau_z + 1)$ and $\epsilon(\cdot)$ is a negligible function in λ . The probability is over the random coins of KGen and the random coins of Enc used to compute C_1, \dots, C_n and C^* .

4.2 Additional Properties of a TB-BFKEM

Again, we will require additional properties for the TB-BFKEM, similar to those from Section 2.3.

Definition 18 (Extended Correctness). We require that the following holds for all $\lambda, m, k, t, n \in \mathbb{N}$, for all $z \in \mathbb{N}$ with $z \leq t$, and any $(\text{sk}, \text{pk}) \leftarrow^{\$} \text{KGen}(1^\lambda, m, k, t)$.

– For any ordered sequence τ_1, \dots, τ_z with $1 \leq \tau_1 < \dots < \tau_z \leq t$ and

$$\text{sk}_{i+1} \leftarrow^{\$} \text{PuncInt}(\text{sk}_i, \tau_i)$$

where $i \in \{1, \dots, z\}$ and $\text{sk}_1 := \text{sk}$, and

– for any arbitrary interleaved sequence of invocations of

$$\text{sk}_{z+1,j+1} \leftarrow^{\$} \text{PuncCtx}(\text{sk}_{z+1,j}, \tau_z + 1, C_j)$$

where $j \in \{1, \dots, n\}$, $\text{sk}_{z+1,1} := \text{sk}_{z+1}$, and $(C_j, K_j) \leftarrow^{\$} \text{Enc}(\text{pk}, \tau_z + 1)$

it holds that:

1. **No false-negatives in the current time interval:**

$\text{Dec}(\text{sk}_{z+1,n+1}, \tau_z + 1, C_j) = \perp$ for all $j \in [n]$

2. **No false-negatives with respect to ciphertexts from previous intervals:**

$\text{Dec}(\text{sk}_{z+1,n+1}, \tau^*, C) = \perp$ for all $(C, K) \leftarrow^{\$} \text{Enc}(\text{pk}, \tau^*)$ with $\tau^* < \tau_z + 1$.

3. **Perfect correctness of the initial secret key:**

$\text{Dec}(\text{sk}, \tau, C) = K$ for all $1 \leq \tau \leq t$ and all $(C, K) \leftarrow^{\$} \text{Enc}(\text{pk}, \tau)$.

4. **Semi-correctness of punctured secret keys:**

For all $1 \leq \tau \leq t$ holds: If $\text{Dec}(\text{sk}_{z+1,j+1}, \tau, C) \neq \perp$ then $\text{Dec}(\text{sk}_{z+1,j+1}, \tau, C) = \text{Dec}(\text{sk}, \tau, C)$.

Definition 19 (Separable Randomness). Let $\text{TB-BFKEM} = (\text{KGen}, \text{Enc}, \text{PuncCtx}, \text{Dec}, \text{PuncInt})$ be a TB-BFKEM. We say that TB-BFKEM has separable randomness, if one can equivalently write the encapsulation algorithm Enc as

$$(C, K) \leftarrow^{\$} \text{Enc}(\text{pk}, \tau) = \text{Enc}(\text{pk}, \tau; (r, K)),$$

for uniformly random $(r, K) \in \{0, 1\}^{\rho+\lambda}$, where $\text{Enc}(\cdot, \cdot; \cdot)$ is a deterministic algorithm whose output is uniquely determined by pk, τ and the randomness $(r, K) \in \{0, 1\}^{\rho+\lambda}$.

Definition 20 (Publicly-Checkable Puncturing). Let $\{\mathcal{Q}_{\tau_j}\}_{j=1}^k$ be any list of lists of ciphertexts $\{(C_{\tau_j,1}, \dots, C_{\tau_j,w_j})\}_{j=1}^k$. We say that TB-BFKEM allows publicly-checkable puncturing, if there exists an efficient algorithm `CheckPunct` with the following correctness property.

1. Run $(\text{sk}, \text{pk}) \leftarrow^{\$} \text{KGen}(1^\lambda, m, k, t)$.
2. For $j \in [k]$ do
 - Compute $C_i \leftarrow^{\$} \text{Enc}(\text{pk}, \tau_j)$ and $\text{sk} = \text{PuncCtx}(\text{sk}, \tau_j, C_i)$ for $i \in [w_j]$.
 - Compute $\text{sk} \leftarrow^{\$} \text{PuncInt}(\text{sk}, \tau_j)$
3. Let C and τ be any string. We require that

$$\perp = \text{Dec}(\text{sk}, \tau, C) \iff \perp = \text{CheckPunct}(\text{pk}, \tau, \{\mathcal{Q}_{\tau_j}\}_{j=1}^k, C).$$

Definition 21 (γ -Spreadness). Let TB-BFKEM = (KGen, Enc, PuncCtx, Dec, PuncInt) be a randomness-separable TB-BFKEM with ciphertext space \mathcal{C} . We say that it is γ -spread, if for any honestly generated pk , any key K , any τ and any $C \in \mathcal{C}$

$$\Pr_{r \leftarrow^{\$} \{0,1\}^\rho} [C = \text{Enc}(\text{pk}, \tau; (r, K))] \leq 2^{-\gamma}.$$

4.3 Security Definitions

The security of a TB-BFKEM scheme is defined in a selective-time experiment, where the adversary has to commit to a time slot τ^* to attack before seeing the parameters of the scheme. We present the IND-CPA and IND-CCA experiments in Figure 6.

Exp $_{\mathcal{A}, \text{TB-BFKEM}}^{\text{s-T}}$ (λ, m, k, t):
 $\tau^* \leftarrow^{\$} \mathcal{A}(1^\lambda)$
 $(\text{sk}, \text{pk}) \leftarrow^{\$} \text{KGen}(1^\lambda, m, k, t), (C^*, K_0) \leftarrow^{\$} \text{Enc}(\text{pk}, \tau^*)$
 $K_1 \leftarrow^{\$} \mathcal{K}, b \leftarrow^{\$} \{0, 1\}, \mathcal{Q}_C \leftarrow \emptyset, \mathcal{Q}_\tau \leftarrow \emptyset$
 $b^* \leftarrow^{\$} \mathcal{A}^{\mathcal{O}, \text{PuncCtx}(\text{sk}, \cdot, \cdot), \text{PuncInt}(\text{sk}, \cdot), \text{Corr}}(\text{pk}, C^*, K_b)$
where $\mathcal{O} \leftarrow \{\text{Dec}'(\text{sk}, \cdot)\}$ if $\text{T} = \text{IND-CCA}$ and $\mathcal{O} \leftarrow \emptyset$ otherwise.
 $\text{Dec}'(\text{sk}, \tau, C)$ behaves as Dec but returns \perp if $C = C^*$ and $\tau = \tau^*$
 $\text{PuncCtx}(\text{sk}, \tau, C)$ runs $\text{sk} \leftarrow^{\$} \text{PuncCtx}(\text{sk}, \tau, C)$ and $\mathcal{Q}_C \leftarrow \mathcal{Q}_C \cup \{(C, \tau)\}$
 $\text{PuncInt}(\text{sk}, \tau)$ runs $\text{sk} \leftarrow^{\$} \text{PuncInt}(\text{sk}, \tau)$ and $\mathcal{Q}_\tau \leftarrow \mathcal{Q}_\tau \cup \{\tau\}$
 Corr returns sk if $(C^*, \tau^*) \in \mathcal{Q}_C$ or $\tau^* \in \mathcal{Q}_\tau$ and \perp otherwise
return 1, if $b^* = b$
return 0

Fig. 6. Security for TB-BFKEM: $\text{T} \in \{\text{IND-CPA}, \text{IND-CCA}\}$.

Definition 22 (s-T-Security of TB-BFKEM). We define the advantage of an adversary \mathcal{A} in the s-T experiment $\text{Exp}_{\mathcal{A}, \text{TB-BFKEM}}^{\text{s-T}}(\lambda, m, k, t)$ as

$$\text{Adv}_{\mathcal{A}, \text{TB-BFKEM}}^{\text{s-T}}(\lambda, m, k, t) := \left| \Pr \left[\text{Exp}_{\mathcal{A}, \text{TB-BFKEM}}^{\text{s-T}}(\lambda, m, k, t) = 1 \right] - \frac{1}{2} \right|.$$

A puncturable forward-secret key-encapsulation scheme TB-BFKEM is s-T, $\text{T} \in \{\text{IND-CPA}, \text{IND-CCA}\}$, secure, if $\text{Adv}_{\mathcal{A}, \text{TB-BFKEM}}^{\text{s-T}}(\lambda, m, k, t)$ is a negligible function in λ for all $m, k, t > 0$ and all PPT adversaries \mathcal{A} .

4.4 A Generic Time-Based BFKEM Construction

Before we can present our construction, we recall hierarchical identity-based key encapsulation schemes (HIB-KEMs). HIB-KEMs represent a building block of our construction.

HIB-KEMs. Below we present the basic definition and the security properties of HIB-KEMs.

Definition 23. A $(t' + 1)$ -level hierarchical identity-based key encapsulation scheme (HIB-KEM) with identity space $\mathcal{D}^{\leq t'+1}$, ciphertext space \mathcal{C} , and key space \mathcal{K} consists of the following four algorithms:

$\text{HIBGen}(1^\lambda)$: Takes as input a security parameter and outputs a key pair $(\text{mpk}, \text{sk}_\varepsilon)$.

We say that mpk is the master public key, and sk_ε is the level-0 secret key.

$\text{HIBDel}(\text{sk}_{d'}, d)$: Takes as input secret key $\text{sk}_{d'}$ and $d \in \mathcal{D}$, and outputs a secret key $\text{sk}_{d'|d}$. (We refer to $|$ as concatenation.)

$\text{HIBEnc}(\text{mpk}, \mathbf{d})$: Takes as input the master public key mpk and an identity $\mathbf{d} \in \mathcal{D}^{\leq t'+1}$ and outputs a ciphertext $C \in \mathcal{C}$ and a key $K \in \mathcal{K}$.

$\text{HIBDec}(\text{sk}_{\mathbf{d}}, C)$: Takes as input a secret key $\text{sk}_{\mathbf{d}}$ and a ciphertext C , and outputs a value $K \in \mathcal{K} \cup \{\perp\}$, where \perp is a distinguished error symbol.

Correctness for HIB-KEM. We require that for all $\lambda \in \mathbb{N}$, for all $(\text{mpk}, \text{sk}_\varepsilon) \xleftarrow{\$} \text{HIBGen}(1^\lambda)$, for all $d \in \mathcal{D}$, for all $\text{sk}_{d'|d} \xleftarrow{\$} \text{HIBDel}(\text{sk}_{d'}, d)$, for all $\mathbf{d} \in \mathcal{D}^{\leq t'+1}$, for all $(C, K) \xleftarrow{\$} \text{HIBEnc}(\text{mpk}, \mathbf{d})$, we have that $\text{HIBDec}(\text{sk}_{\mathbf{d}}, C) = K$ holds.

Security definition for HIB-KEM. As for our generic construction we will essentially follow the proof strategy of the BFKEM construction and thus will rely on the weak notion of one-wayness under selective-ID and chosen-plaintext attacks (OW-sID-CPA) for HIB-KEM. We note that any IND-sID-CPA secure HIB-KEM also satisfies this notion of OW-sID-CPA security.

$$\begin{aligned} & \text{Exp}_{\mathcal{A}, \text{HIB-KEM}}^{\text{OW-sID-CPA}}(\lambda) \\ & \mathbf{d}^* \xleftarrow{\$} \mathcal{A}(1^\lambda) \\ & \text{if } \mathbf{d}^* \notin \mathcal{D} \text{ return } 0 \\ & (\text{mpk}, \text{sk}_\varepsilon) \xleftarrow{\$} \text{HIBGen}(1^\lambda), (C, K) \xleftarrow{\$} \text{HIBEnc}(\text{mpk}, \mathbf{d}^*) \\ & K^* \xleftarrow{\$} \mathcal{A}(\text{mpk}, C) \\ & \text{return } 1, \text{ if } K^* = K \\ & \text{return } 0 \end{aligned}$$

Fig. 7. OW-sID-CPA security.

Definition 24 (OW-sID-CPA Security of HIB-KEM). We define the advantage of an adversary \mathcal{A} in the OW-sID-CPA experiment $\text{Exp}_{\mathcal{A}, \text{HIB-KEM}}^{\text{OW-sID-CPA}}(\lambda)$ as

$$\text{Adv}_{\mathcal{A}, \text{HIB-KEM}}^{\text{OW-sID-CPA}}(\lambda) := \Pr \left[\text{Exp}_{\mathcal{A}, \text{HIB-KEM}}^{\text{OW-sID-CPA}}(\lambda) = 1 \right].$$

We call a HIB-KEM OW-sID-CPA secure, if $\text{Adv}_{\mathcal{A}, \text{HIB-KEM}}^{\text{OW-sID-CPA}}(\lambda)$ is a negligible function in λ for all PPT adversaries \mathcal{A} .

Time slots. We will construct a TB-BFKEM scheme that allows to use $t = 2^{t'}$ time slots. We associate the i -th time slot with the string in $\{0, 1\}^{t'}$ that corresponds to the canonical t' -bit binary representation of integer i .

Following [15,31,33], each time slot forms a leaf of an ordered binary tree of depth t' . The root of the tree is associated with the empty string ϵ . We associate the left-hand descendants of the root with bit string 0, and the right-hand descendant with 1. Continuing this way, we associate the left descendant of node 0 with 00 and the right descendant with 01, and so on. We continue this procedure for all nodes, until we have constructed a complete binary tree of depth t' . Note that two nodes at level $j \leq t'$ of the tree are siblings if and only if their first $j - 1$ bits are equal, and that each bit string in $\{0, 1\}^{t'}$ is associated with a leaf of the tree. Note also that the leaves in the tree are ordered, in the sense that the leftmost leaf is associated with $0^{t'}$, its right neighbor with $0^{t'-1}1$, and so on.

Intuition of the construction. The basic idea behind the construction combines the binary tree approach of [15,31,33] with the BF-KEM construction described in Section 3.1. We use a HIB-KEM with identity space

$$\mathcal{D} = \mathcal{D}_1 \times \cdots \times \mathcal{D}_{t'+1} = \underbrace{\{0, 1\} \times \cdots \times \{0, 1\}}_{t' \text{ times}} \times [m].$$

Each bit vector $\tau \in \mathcal{D}_1 \times \cdots \times \mathcal{D}_{t'} = \{0, 1\}^{t'}$ corresponds to one time slot, and we set $\mathcal{D}_{t'+1} = [m]$, where m is the size of the Bloom filter. The hierarchical key delegation property of the HIB-KEM enables the following features:

First, given a HIB-KEM key sk_τ for some “identity” (= time slot) $\tau \in \{0, 1\}^{t'}$, we can derive keys for all Bloom filter bits from sk_τ by computing

$$\text{sk}_{\tau|d} \stackrel{\$}{\leftarrow} \text{HIBDel}(\text{sk}_\tau, d) \quad \text{for all } d \in [m].$$

Second, in order to advance from time slot τ to $\tau + 1$, we first compute

$$\text{sk}_{\tau|d} \stackrel{\$}{\leftarrow} \text{HIBDel}(\text{sk}_\tau, d) \quad \text{for all } d \in [m].$$

As soon as we have computed all Bloom filter keys for time slot τ , we “puncture” the tree “from left to right”, such that we are able to compute all $\text{sk}_{\tau'}$ with $\tau' > \tau$, but not any $\text{sk}_{\tau'}$ with $\tau' \leq \tau$. Here, we proceed exactly as in [15,31,33]. That is, in order to puncture at time slot τ , we first compute the HIB-KEM secret keys associated to all *right-hand siblings* of nodes that lie on the path from node τ to the root (if existent), and then we delete all secret keys associated to nodes that lie on the path from node τ to the root, including sk_τ itself. This yields a new secret key, which contains m level- $(t' + 1)$ HIB-KEM secret keys plus at most t' HIB-KEM secret keys for levels $\leq t'$, even though we allow for $2^{t'}$ time slots.

Construction. Let $(\text{HIBGen}, \text{HIBDel}, \text{HIBEnc}, \text{HIBDec})$ be a $(t' + 1)$ -level HIB-KEM with key space \mathcal{K} and identity space $\mathcal{D} = \mathcal{D}_1 \times \cdots \times \mathcal{D}_{t'+1}$, where $\mathcal{D}_1 = \cdots = \mathcal{D}_t = \{0, 1\}$, $\mathcal{D}_{t'+1} = [m]$, and m is the size of the Bloom filter. Since we will only need selective security, one can instantiate such a HIB-KEM very efficiently, for example in bilinear groups based on the Boneh-Boyen-Goh [11] scheme, or based

on lattices [1]. In the sequel, we will write $\{0, 1\}^{t'}$ shorthand for $\mathcal{D}_1 \times \dots \times \mathcal{D}_{t'}$, but keep in mind that the HIB-KEM supports more fine-grained key delegation. Let $B = (\text{BFGen}, \text{BFUpdate}, \text{BFCheck})$ be a Bloom filter for set $\{0, 1\}^\lambda$. Furthermore, let $G' : \mathcal{K} \rightarrow \{0, 1\}^\lambda$ be a hash function (which will be modeled as a random oracle [7] in the security proof).

We define TB-BFKEM = (KGen, Enc, PuncCtx, Dec, PuncInt) as follows.

KGen($1^\lambda, m, k, t = 2^{t'}$): This algorithm first runs $((H_j)_{j \in [k]}, T) \leftarrow^{\$} \text{BFGen}(m, k)$ to generate a Bloom filter, and $(\text{mpk}, \text{sk}_\epsilon) \leftarrow^{\$} \text{HIBGen}(1^\lambda)$ to generate a key pair. Finally, the algorithm generates the keys for the first time slot. To this end, it first computes the HIB-KEM key for identity $0^{t'}$ by recursively computing

$$\text{sk}_{0^d} \leftarrow^{\$} \text{HIBDel}(\text{sk}_{0^{d-1}}, 0) \quad \text{for all } d \in [t].^8$$

Then it computes the m Bloom filter keys for time slot $0^{t'}$ by computing

$$\text{sk}_{0^{t'}|d} \leftarrow^{\$} \text{HIBDel}(\text{sk}_{0^{t'}}, d) \quad \text{for all } d \in [m],$$

and setting $\text{sk}_{\text{Bloom}} := (\text{sk}_{0^{t'}|d})_{d \in [m]}$. Finally, it punctures the secret key sk_ϵ at position $0^{t'}$, by computing

$$\text{sk}_{0^{d-1}1} \leftarrow^{\$} \text{HIBDel}(\text{sk}_{0^{d-1}}, 1) \quad \text{for all } d \in [t'],$$

and setting $\text{sk}_{\text{time}} := (\text{sk}_{0^{d-1}1})_{d \in [t']}$. The algorithm outputs

$$\text{sk} := (T, \text{sk}_{\text{Bloom}}, \text{sk}_{\text{time}}) \text{ and } \text{pk} := (\text{mpk}, (H_j)_{j \in [k]}).$$

Enc(mpk, τ): On input mpk and time slot identifier $\tau \in \{0, 1\}^{t'}$, this algorithm first samples a random string $c \leftarrow^{\$} \{0, 1\}^\lambda$ and a random key $K \leftarrow^{\$} \{0, 1\}^\lambda$. Then it defines k HIB-KEM identities as $\mathbf{d}_j := (\tau, H_j(c)) \in \mathcal{D}$ for $j \in [k]$, and generates k HIB-KEM key encapsulations as

$$(C_j, K_j) \leftarrow^{\$} \text{HIBEnc}(\text{mpk}, \mathbf{d}_j) \quad \text{for } j \in [k].$$

Finally, it outputs the ciphertext $C := (c, (C_j, G'(K_j) \oplus K)_{j \in [k]})$.

Note that the ciphertexts essentially consists of $k + 1$ elements of $\{0, 1\}^\lambda$, plus k elements of \mathcal{C} , where k is the Bloom filter parameter.

PuncCtx(sk, C): Given a ciphertext $C := (c, (C_j, G'(K_j) \oplus K)_{j \in [k]})$, and secret key $\text{sk} = (T, \text{sk}_{\text{Bloom}}, \text{sk}_{\text{time}})$ where $\text{sk}_{\text{Bloom}} = (\text{sk}_{\tau|d})_{d \in [m]}$, the puncturing algorithm first computes $T' = \text{BFUpdate}((H_j)_{j \in [k]}, T, c)$. Then, for each $i \in [m]$, it defines

$$\text{sk}'_{\tau|i} := \begin{cases} \text{sk}_{\tau|i} & \text{if } T'[i] = 0, \text{ and} \\ \perp & \text{if } T'[i] = 1, \end{cases}$$

where $T'[i]$ denotes the i -th bit of T' . Finally, this algorithm sets $\text{sk}'_{\text{Bloom}} = (\text{sk}'_{\tau|d})_{d \in [m]}$ and returns $\text{sk}' = (T', \text{sk}'_{\text{Bloom}}, \text{sk}_{\text{time}})$.

⁸ Implicitly, we set $\epsilon := 0^0$.

Remark. We note again that the above procedure is correct even if the procedure is applied repeatedly, with the same arguments as for the construction from Section 3.1. Also, the puncturing algorithm essentially only evaluates k universal hash functions and then deletes a few secret keys, which makes this procedure extremely efficient.

$\text{Dec}(\text{sk}, C)$: Given $\text{sk} = (T, \text{sk}_{\text{Bloom}}, \text{sk}_{\text{time}})$ where $\text{sk}_{\text{Bloom}} = (\text{sk}_{\tau|d})_{d \in [m]}$ and ciphertext $C := (c, (C_j, G_j)_{j \in [k]})$. If $\text{sk}_{\tau|H_j(c)} = \perp$ for all $j \in [k]$, then it outputs \perp . Otherwise, it picks the smallest index j such that $\text{sk}_{\tau|H_j(c)} \neq \perp$, computes

$$K_j = \text{HIBDec}(\text{sk}_{\tau|H_j(c)}, C_j),$$

and returns $K = G_j \oplus G'(K_j)$.

Remark. Again we have $\text{Dec}(\text{sk}, C) \neq \perp \iff \text{BFCheck}(H, T, c) = 0$, which guarantees extended correctness in the sense of Definition 18.

$\text{PunctInt}(\text{sk}, \tau)$: Given a secret key $\text{sk} = (T, \text{sk}_{\text{Bloom}}, \text{sk}_{\text{time}})$ for time interval $\tau' \leq \tau$, the time puncturing algorithm proceeds as follows. First, it resets the Bloom filter by setting $T := 0^m$. Then it uses the key delegation algorithm to first compute sk_{τ} . This key can be computed from the keys contained in sk_{time} , because sk is a key for time interval $\tau' \leq \tau$. Then it computes

$$\text{sk}_{\tau|d} \stackrel{\$}{\leftarrow} \text{HIBDel}(\text{sk}_{\tau}, d) \quad \text{for all } d \in [m],$$

and redefines $\text{sk}_{\text{Bloom}} := (\text{sk}_{\tau|d})_{d \in [m]}$. Finally, it updates sk_{time} by computing the HIB-KEM secret keys associated to all *right-hand siblings* of nodes that lie on the path from node τ to the root and adds the corresponding keys to sk_{time} . Then it deletes all keys from sk_{time} that lie on the path from τ to the root.

Remark. Note that puncturing between time intervals may become relatively expensive. Depending on the choice of Bloom filter parameters, in particular on m , this may range between 2^{15} and 2^{25} HIBKEM key delegations. However, the main advantage of BFKEM over previous constructions of puncturable encryption is that these computations must not be performed “online”, during puncturing, but can actually be computed separately (for instance, parallel on a different computer, or when a server has low workload, etc.).

Correctness error of this scheme. Note that within each time slot we can use the same argumentation as for the scheme from Section 3.1 and one can verify that the correctness-error probability of this scheme is essentially identical to the false-positive probability of the Bloom filter, unless a given ciphertext $C = (c, (C_j, G_j)_{j \in [k]})$ has a value of c which is identical to the value of c of any previous ciphertext. Since c is uniformly random in $\{0, 1\}^\lambda$, this probability is bounded by $\left(1 - e^{-\frac{(n+1/2)k}{m-1}}\right)^k + n \cdot 2^{-\lambda}$. Furthermore, due to the perfect correctness of the underlying HIB-KEM scheme, this yields the error-correctness bound of TB-BFKEM.

More formally, we have that for all $\lambda, m, k, t' \in \mathbb{N}$, for all $t = 2^{t'}$, for all $z \in \mathbb{N}$ with $z \leq t$, and any $(\text{sk}, \text{pk}) \stackrel{\$}{\leftarrow} \text{KGen}(1^\lambda, m, k, t)$:

– For any *ordered* sequence $\tau_1, \dots, \tau_z \in \{0, 1\}^{t'}$ and

$$\text{sk}_{i+1} \stackrel{\$}{\leftarrow} \text{PunctInt}(\text{sk}_i, \tau_i)$$

- where $i \in \{1, \dots, z\}$ and $\text{sk}_1 := \text{sk}$, and
- for any *arbitrary interleaved* sequence of invocations of

$$\text{sk}_{z+1,j+1} \stackrel{\$}{\leftarrow} \text{PuncCtx}(\text{sk}_{z+1,j}, \tau_z + 1, C_j)$$

where $j \in \{1, \dots, n\}$, $\text{sk}_{z+1,1} := \text{sk}_{z+1}$, and $(C_j, K_j) \stackrel{\$}{\leftarrow} \text{Enc}(\text{pk}, \tau_z + 1)$

it holds that

$$\Pr [\text{Dec}(\text{sk}_{z+1,n+1}, \tau_z + 1, C^*) \neq K^*] \leq \left(1 - e^{-\frac{(n+1/2)k}{m-1}}\right)^k + n \cdot 2^{-\lambda}$$

where $(C^*, K^*) \stackrel{\$}{\leftarrow} \text{Enc}(\text{pk}, \tau_z + 1)$, due to the perfect correctness of HIB-KEM, the bounded false-positive probability property of the underlying Bloom filter (depending on m and k , see Section 2.1), and the bounded negligible loss $n \cdot 2^{-\lambda}$ occurred due to collisions of uniform c -values in the ciphertexts.

Concerning the extended correctness, see that

1. $\text{Dec}(\text{sk}_{z+1,n+1}, \tau_z + 1, C_j) = \perp$ for all $j \in [n]$, i.e., no false-negatives in the current time interval due to the perfect-completeness property of the BF B and deleting the respective HIB-KEM secret keys.
2. $\text{Dec}(\text{sk}_{z+1,n+1}, \tau^*, C) = \perp$ for all $(C, K) \stackrel{\$}{\leftarrow} \text{Enc}(\text{pk}, \tau^*)$ with $\tau^* < \tau_z + 1$, i.e., no false-negatives with respect to ciphertexts from previous intervals due to the security properties of HIB-KEM.
3. $\text{Dec}(\text{sk}, \tau, C) = K$ for all $1 \leq \tau \leq t$ and all $(C, K) \stackrel{\$}{\leftarrow} \text{Enc}(\text{pk}, \tau)$, i.e., perfect correctness of the initial secret key due to the perfect correctness property of HIB-KEM and perfect completeness of B .
4. For all $1 \leq \tau \leq t$ holds: If $\text{Dec}(\text{sk}_{z+1,j+1}, \tau, C) \neq \perp$ then $\text{Dec}(\text{sk}_{z+1,j+1}, \tau, C) = \text{Dec}(\text{sk}, \tau, C)$, i.e., semi-correctness of punctured secret keys due to perfect correctness property of HIB-KEM and perfect completeness of B .

CPA Security. Below we state theorem for CPA security of our scheme.

Theorem 6. *From each efficient adversary \mathcal{B} that issues u queries to random oracle G' we can construct an efficient adversary \mathcal{A} with*

$$\text{Adv}_{\mathcal{B}, \text{TB-BFKEM}}^{\text{s-IND-CPA}}(\lambda, m, k) \leq uk \cdot \text{Adv}_{\mathcal{A}, \text{HIB-KEM}}^{\text{OW-sID-CPA}}(\lambda).$$

The proof is almost identical to the proof of Theorem 1 and a straightforward reduction to the security of the underlying HIB-KEM. We sketch it below.

Proof (Sketch). We sketch the proof of Theorem 6. Recall that a ciphertext has the form

$$C := (c, (C_j, G'(K_j) \oplus K)_{j \in [k]}).$$

Essentially, one argues exactly as in Theorem 1 that the adversary receives no information about the key K encapsulated by the Bloom filter encryption scheme, unless it ever queries K_j to random oracle G' for some $j \in [k]$. Therefore assume that \mathcal{B} queries some K_j to G' in its u^* -th query.

At the beginning of the reduction, \mathcal{A} first guesses index $j \leftarrow^{\$} [k]$ and $u^* \leftarrow^{\$} [u]$. It also samples the random string $c \leftarrow^{\$} \{0, 1\}^\lambda$ used for the challenge BFKEM ciphertext at the beginning of the game, generates a Bloom filter

$$((H_j)_{j \in [k]}, T) \leftarrow^{\$} \text{BFGen}(m, k),$$

and requests a challenge ciphertext for identity $d^* = (\tau^* | H_j(c))$, where τ^* is the time slot selected by \mathcal{B} . The challenge ciphertext received back from the HIB-KEM experiment is then embedded in the TB-BFKEM challenge ciphertext. The PuncCtx and PuncInt(sk, \cdot) queries of \mathcal{B} can trivially be simulated by \mathcal{A} . The Corr queries can be answered using the HIBDel oracle provided by the OW-sID-CPA security experiment of the HIB-KEM.

When \mathcal{B} makes its u^* -th query to G' on value K' , then \mathcal{A} terminates and outputs K' . We know that any non-trivial adversary \mathcal{B} queries K_j to G' for some j . If \mathcal{A} has guessed u^* and j correctly, which happens with probability $1/(uk)$, then it holds that $K' = K_j$, which yields the claim. \square

CCA Security. In order to apply the Fujisaki-Okamoto [25] transform in the same way as done in Section 3.2 to achieve CCA security, we need to show that the time based variants of the properties presented in Section 2.3 are satisfied (i.e., Definitions 18, 19, 20, and 21). First, using a full-blown HIBE as a starting point yields a separable HIB-KEM as discussed in Section 2.3. Hence, the separable randomness (Def. 19) is satisfied. Moreover, the publicly-checkable puncturing (Def. 20) is given by construction (as in Section 3.1). Regarding extended correctness (Def. 18), the impossibility of false-negatives is given by construction, the perfect correctness of the non-punctured secret key is given by the perfect correctness of the HIBE and the semi-correctness of punctured secret keys is given by construction. Finally, γ -spreadness (Def. 21) is also given by construction: the ciphertext component c is chosen uniformly at random from $\{0, 1\}^\lambda$. Consequently, all properties are satisfied. We note that one could omit c in the ciphertext if the concretely used HIBE ciphertexts are already sufficiently random. Considering the HIBE of Boneh-Boyen-Goh [11], HIBE ciphertexts are of the form $(g^r, (h_1^{I_1} \cdots h_t^{I_t} \cdot h_0)^r, H(e(g_1, g_2)^r) \oplus K)$, for honestly generated fixed group elements $g, g_1, g_2, h_0, \dots, h_t$, universal hash function H , fixed K and fixed integers I_1, \dots, I_t . Consequently, we have that the ciphertext has at least min-entropy $\log_2 q$ with q being the order of the groups. We want to mention that also many other HIBE construction satisfy the required properties, including, for example [27,48,19].

Remark on CCA Security. Alternatively to applying the FO transform to a TB-BFKEM satisfying the additional properties of extended correctness, separable randomness, publicly checkable puncturing and γ -spreadness to obtain CCA security, we can add another HIBE level to obtain IND-CCA security via the CHK transform [15] in the standard model, and thus to avoid random oracles if required.⁹

⁹ We note, however, that one cannot straightforwardly apply the CHK transform in a black-box way, but needs to take care that all k HIB-KEM ciphertexts C_j , $j \in [k]$ need to use *the same* verification key of the strong one-time signature used to sign the overall ciphertext.

5 Forward-Secret 0-RTT Key Exchange

In [33] (with full version in [34]), Günther, Hale, Jager, and Lauer (GHJL) provide a formal model for forward-secret one-pass key exchange (FSOPKE) by extending the one-pass key exchange [36] by Halevi and Krawczyk. They provide a security model for FSOPKE which requires both forward secrecy and replay protection from the FSOPKE protocol and captures unilateral authentication of the server and mutual authentication simultaneously.

We recap the definition of FSOPKE with a slightly adapted KGen and correctness notion to allow for a non-negligible correctness error, which constitutes the main difference to the work of [33,34]. We remark that we do not change the security model of [34, Sec. 3.2] (i.e., the new input parameters added to KGen do only affect FSOPKE correctness and not the FSOPKE security model).

In particular, we now take into account the maximum number of server and client runs $n \in \mathbb{N}$ and a false-positive probability p of succeeding the runs (in computing the same session key for a particular time step). Looking ahead, these changes are necessary to instantiate FSOPKE with our TB-BFKEM. (See Sec. 2.1 for parameter selections of n and p and corresponding BF parameter m and k if one wants to instantiate the FSOPKE with a BF.)

Definition 25 (FSOPKE). *An FSOPKE scheme FSOPKE providing mutual or unilateral (server-only) authentication consists of the PPT algorithms (FSOPKE.KGen, FSOPKE.RunC, FSOPKE.RunS, FSOPKE.TimeStep):*

$\text{FSOPKE.KGen}(1^\lambda, r, \tau_{max}, n, p)$: *Takes as input a security parameter 1^λ , a role $r \in \{\text{server}, \text{client}\}$, the maximum number of time slots $\tau_{max} \in \mathbb{N}$, and the maximum number of server and client runs $n \in \mathbb{N}$ with false-positive probability p , outputs public and secret keys (pk, sk) for a specific role r (we assume that the key-space \mathcal{K} is implicit in pk).*

$\text{FSOPKE.RunC}(sk, pk)$: *Takes as input a secret key sk , a public key pk , and outputs a (potentially modified) secret key sk' , a session key $K \in \{0, 1\}^* \cup \{\perp\}$, and a message $M \in \{0, 1\}^* \cup \{\perp\}$.*

$\text{FSOPKE.RunS}(sk, pk, M)$: *Takes as input a secret key sk , a public key pk , and a message $M \in \{0, 1\}^*$ and outputs a (potentially modified) secret key sk' and a session key $K \in \{0, 1\}^* \cup \{\perp\}$.*

$\text{FSOPKE.TimeStep}(sk, r)$: *Takes as input a secret key sk and an according role $r \in \{\text{client}, \text{server}\}$ and outputs a (potentially modified) secret key sk' .*

Server and client flow within an FSOPKE scheme. A server and a client are engaging in an FSOPKE scheme as follows. According to their role, a server (i.e., when the role is server) and a client (i.e., when the role is client) execute

$$(pk_s, sk_{s,1,1}) \leftarrow \text{FSOPKE.KGen}(1^\lambda, \text{server}, \tau_{max}, n, p),$$

$$(pk_c, sk_{c,1,1}) \leftarrow \text{FSOPKE.KGen}(1^\lambda, \text{client}, \tau_{max}, n, p)$$

to generate initial public and private keys for a server and client, respectively (where λ , τ_{max} , n , and p are pre-determined). For all $i \in [\tau_{max}]$ and $j \in [n]$, by executing $sk_{s,i+1,1} \leftarrow \text{FSOPKE.TimeStep}(sk_{s,i,j}, \text{server})$ and $sk_{c,i+1,1} \leftarrow \text{FSOPKE}$.

$\text{TimeStep}(\text{sk}_{c,i,j}, \text{client})$, the server and the client progresses from one time slot i to the next slot $i + 1$ to receive (potentially modified) secret keys $\text{sk}_{s,i+1,1}$ and $\text{sk}_{c,i+1,1}$, respectively.

Similarly, within a time step $i \in [\tau_{max}]$ and for all $j \in [n - 1]$, the client proceeds with

$$(\text{sk}_{c,i,j+1}, \text{K}_{c,i,j}, M) \leftarrow \text{FSOPKE.RunC}(\text{sk}_{c,i,j}, \text{pk}_s),$$

for client's private key $\text{sk}_{c,i,j}$ and a server's public key pk_s , to receive a (potentially modified) secret key $\text{sk}_{c,i,j+1}$, a session key $\text{K}_{c,i,j}$, and a message M (e.g., a ciphertext) which is transmitted to the server (associated to pk_s). The server obtains M and executes

$$(\text{sk}_{s,i,j+1}, \text{K}_{s,i,j}) \leftarrow \text{FSOPKE.RunS}(\text{sk}_{s,i,j}, \text{pk}_c, M),$$

for the server's secret key $\text{sk}_{s,i,j}$ and the client's public key pk_c to receive a (potentially modified) secret key $\text{sk}_{s,i,j+1}$ and a session key $\text{K}_{s,i,j}$.

By our adapted correctness notion of the FSOPKE (see Def. 26 below), we have that $\text{K}_{c,i,j} = \text{K}_{s,i,j}$ except with potentially non-negligible probability (bounded by $\mu(n, p) + \varepsilon(\lambda)$), for a potential non-negligible function $\mu(\cdot, \cdot)$ that only depends on n and p , as well as a negligible function $\varepsilon(\cdot)$ that depends on λ .

Definition 26 ((Non-negligible) correctness of FSOPKE). For all $\lambda, \tau_{max}, n \in \mathbb{N}$ and false-positive probability p , for all

$$(\text{pk}_s, \text{sk}_{s,1,1}) \leftarrow \text{FSOPKE.KGen}(1^\lambda, \text{server}, \tau_{max}, n, p),$$

$$(\text{pk}_c, \text{sk}_{c,1,1}) \leftarrow \text{FSOPKE.KGen}(1^\lambda, \text{client}, \tau_{max}, n, p),$$

for all $i \in [\tau_{max} - 1]$ and all $j \in [n]$, for any

$$\text{sk}_{s,i+1,1} \leftarrow \text{FSOPKE.TimeStep}(\text{sk}_{s,i,j}, \text{server}),$$

$$\text{sk}_{c,i+1,1} \leftarrow \text{FSOPKE.TimeStep}(\text{sk}_{c,i,j}, \text{client}),$$

for all $j \in [n - 1]$ and $i \in [\tau_{max}]$, for any

$$(\text{sk}_{c,i',j'+1}, \text{K}_{c,i',j'}, M) \leftarrow \text{FSOPKE.RunC}(\text{sk}_{c,i',j'}, \text{pk}_s),$$

$$(\text{sk}_{s,i',j'+1}, \text{K}_{s,i',j'}) \leftarrow \text{FSOPKE.RunS}(\text{sk}_{s,i',j'}, \text{pk}_c, M) \text{ (if mutual auth.)},$$

$$(\text{sk}_{s,i',j'+1}, \text{K}_{s,i',j'}) \leftarrow \text{FSOPKE.RunS}(\text{sk}_{s,i',j'}, \perp, M) \text{ (if unilateral auth.)},$$

for all $j'' \in [n]$ and $i'' \in [\tau_{max}]$, we have that

$$\Pr[\text{K}_{s,i'',j''} \neq \text{K}_{c,i'',j''}] \leq \mu(n, p) + \varepsilon(\lambda),$$

where $\mu(\cdot, \cdot)$ is some (possibly non-negligible) function that depends on n and p , and $\varepsilon(\cdot)$ is a negligible function that depends on λ .

Security of FSOPKE. The security model of FSOPKE is the same as in defined in [34, Section 3.2] and we omit it here. (See that the additional parameter n and p only affect correctness.) All security guarantees from [34, Theorem 2] directly translate to our FSOPKE construction and we restate the security claim in Corollary 1. (However, we have to argue about the correctness property of our FSOPKE construction where we now allow for a potential non-negligible correctness probability depending on λ, n , and p compared to perfect correctness of the FSOPKE scheme of [34, Def. 12].)

5.1 Construction of an unilateral FSOPKE scheme from a TB-BFKEM

The unilateral-authenticated¹⁰ FSOPKE construction in [33,34] builds on PFSKEM; our TB-BFKEM is defined analogously (cf. Def. 16), but we have to account for the non-negligible correctness errors (in particular, for parameters n and p).

Let TB-BFKEM = (KGen, Enc, PuncCtx, Dec, PuncInt) be a TB-BFKEM. We construct an FSOPKE FSOPKE = (FSOPKE.KGen, FSOPKE.RunC, FSOPKE.RunS, FSOPKE.TimeStep) as follows:

FSOPKE.KGen($1^\lambda, r, \tau_{max}, n, p$) : On input security parameter 1^λ , role $r \in \{\text{server, client}\}$, maximum number of time steps $\tau_{max} \in \mathbb{N}$, maximum number of server and client runs $n \in \mathbb{N}$ with false-positive probability p , proceed as follows:

- if $r = \text{server}$, then compute $(PK, SK_{1,1}) \leftarrow \text{KGen}(1^\lambda, m, k, \tau_{max})$ (for suitable choices of m, k according to Sec. 2.1 depending on n and p) and set $\text{pk}_s := (PK, \tau_{max})$ and $\text{sk}_{s,1,1} := (SK_{1,1}, 1, \tau_{max})$, and output $(\text{pk}_s, \text{sk}_{s,i,1})$.
- if $r = \text{client}$, then set $(\text{pk}_c, \text{sk}_{c,1,1}) := (\perp, 1)$.

FSOPKE.RunC($\text{sk}_{c,i,j}, \text{pk}_s$) : Outputs $(\text{sk}_{c,i,j+1}, K_{c,i,j}, M)$ as follows: for $\text{sk}_{c,i,j} = i$ and $\text{pk}_s = (PK, \tau_{max})$, if $i > \tau_{max}$, then set $(\text{sk}_{c,i,j+1}, K_{c,i,j}, M) := (\text{sk}_{c,i,j}, \perp, \perp)$, otherwise obtain $(C, K_{c,i,j}) \leftarrow \text{Enc}(\text{pk}_s, i)$ and set $(\text{sk}_{c,i,j+1}, K_{c,i,j}, M) := (\tau, K_{c,i,j}, C)$.

FSOPKE.RunS($\text{sk}_{s,i,j}, \text{pk}_c, M$) : Outputs $(\text{sk}_{s,i,j+1}, K_{s,i,j})$ as follows: for $\text{sk}_{s,i,j} = (SK_{i,j}, i, \tau_{max})$ and $\text{pk}_c = \perp$, if $SK_{i,j} = \perp$ or $i > \tau_{max}$, then set $(\text{sk}_{s,i,j+1}, K_{s,i,j}) := (\text{sk}_{s,i,j}, \perp)$ and abort. Obtain $K_{s,i,j} \leftarrow \text{Dec}(SK_{i,j}, M)$. If $K_{s,i,j} = \perp$, then set $(\text{sk}_{s,i,j+1}, K_{s,i,j}) = (\text{sk}_{s,i,j}, \perp)$, otherwise obtain $SK_{i,j+1} \leftarrow \text{PuncCtx}(SK_{i,j}, M)$ and set $(\text{sk}_{s,i,j+1}, K_{s,i,j}) = ((SK_{i,j+1}, i, \tau_{max}), K_{s,i,j})$.

FSOPKE.TimeStep($\text{sk}_{\{s,c\},i,j}, r$) : Outputs $\text{sk}_{\{s,c\},i+1,1}$ as follows:

- if $r = \text{server}$, then for $\text{sk}_{s,i,j} = (SK_{i,j}, i, \tau_{max})$: if $i \geq \tau_{max}$, then set $\text{sk}_{s,i+1,1} := (\perp, i+1, \tau_{max})$ and abort, otherwise compute $SK_{i+1,1} \leftarrow \text{PuncInt}(SK_{i,j}, i)$ and set $\text{sk}_{s,i+1,1} := (SK_{i+1,1}, i+1, \tau_{max})$ and abort.
- if $r = \text{client}$, then for $\text{sk}_{c,i,j} = i$, set $\text{sk}_{c,i+1,1} := i+1$.

Correctness of FSOPKE. Since the underlying TB-BFKEM scheme TB-BFKEM has a potential non-negligible error correctness that is bounded by $\mu'(m, k) + \varepsilon'(\lambda)$, this directly translates to the correctness of the FSOPKE scheme FSOPKE in the sense that the TB-BFKEM-decapsulation Dec fails with potential non-negligible probability $\mu(n, p) + \varepsilon(\lambda) \leq \mu'(m, k) + \varepsilon'(\lambda)$ which in turn translates to FSOPKE.RunS returning $(\text{sk}_{s,i,j+1}, K_{s,i,j} \neq K_{c,i,j})$ with potential non-negligible probability depending on λ, n , and p .

Security guarantees hold due to [34, Theorem 2]. (In particular, see that the parameters n and p we introduce in our FSOPKE definition above only affect the correctness error of the FSOPKE and all FSOPKE security guarantee are required to be independent not only of the time steps (as discussed in [34]), but also independent of n and p .) We state the following corollary:

¹⁰ That is server-side authentication only.

Corollary 1. Let TB-BFKEM be a TB-BFEKEM as defined in Def. 16 with advantage function $\text{Adv}_{\mathcal{A}, \text{TB-BFKEM}}^{\text{s-IND-CCA}}(\lambda, m, k, t)$ (cf. Def. 22). For any PPT adversary \mathcal{B} in the FSOPKE-sec security experiment defined in [34, Def. 11], we have

$$\text{Adv}_{\mathcal{B}, \text{FSOPKE}}^{\text{FSOPKE-sec}}(\lambda) \leq \tau_{\max} \cdot n_{\mathcal{I}} \cdot n_s \cdot \text{Adv}_{\mathcal{A}, \text{TB-BFKEM}}^{\text{s-IND-CCA}}(\lambda, m, k, t),$$

for FSOPKE’s advantage function $\text{Adv}_{\mathcal{B}, \text{FSOPKE}}^{\text{FSOPKE-sec}}(\lambda)$ as defined in [34, Def. 11], for τ_{\max} the maximum number of time slots per session, $n_{\mathcal{I}}$ the maximum number of identities, and n_s the maximum number of sessions (where $n_{\mathcal{I}}$ and n_s are given as in [34, Sec. 3.2]). Furthermore, m and k depend on n and p as described in Sec 2.1.

6 Analysis

BFKEM. Finally, we compare our different BFKEM instantiations as presented in Section 3 regarding their time and space complexity (also see Table 2). Regarding computational efficiency of Dec and Punc, all schemes are roughly the same. The space complexity in the direct construction is optimal with respect to the size of public and secret keys, and we achieve ciphertexts of size $O(k)$. Our construction based on CP-ABE can achieve constant size ciphertexts when instantiated with an ABE scheme that achieves constant size ciphertexts. We, however, note that all ABE schemes achieving constant size ciphertext we are aware of (i.e., [18,2]) come at the cost of large public and secret keys. Those key sizes also carry over to our BFKEM construction. Finally, our construction from IBBE can be viewed as the dual to our direct construction in terms of space complexity. That is, the scheme based on IBBE is optimal regarding the size of ciphertexts and secret keys, while it requires $O(k)$ sized public keys.

Scheme	pk	sk	C	Dec	Punc
direct	$O(1)$	$O(m)$	$O(k)$	$O(k)$	$O(k)$
ABE [18,2]	$O(m)$	$O(m^2)$	$O(1)$	$O(k)$	$O(k)$
IBBE [22]	$O(k)$	$O(m)$	$O(1)$	$O(k)$	$O(k)$

Table 2. Performance Comparison

When taking concrete values regarding space complexity into account, our IBBE based construction is the favorable one. In particular, when we use the IBBE by Delerablée [22] (for convenience we recall it in Appendix D), we obtain ciphertexts $C \in \{0, 1\}^\lambda \times \mathbb{G}_1 \times \mathbb{G}_2$ and secret keys $\text{sk} \in \{0, 1\}^m \times \mathbb{G}_1^m$. That is, ciphertexts are shorter and secret key entries are only half the size of the ones in our direct construction. It is, however, important to note that those efficiency gains come at the cost of a stronger assumption (whose validity was analyzed in the generic bilinear group model in [22]).

TB-BFKEM. In Table 3, we provide an overview of all existing practically instantiable approaches to construct a PFSKEM and compare them to the TB-BFKEM proposed in this paper.¹¹ We compare all schemes for an arbitrary number ℓ of time slots, where for

¹¹ We consider all but the PE schemes from indistinguishability obfuscation [20,17].

sake of simplicity we assume $\ell = 2^t$ for some integer t , (corresponding to our time-based BFKEM) and only count the expensive cryptographic operations, i.e., such as group exponentiations and pairings.

Scheme	$ pk $	$ sk $	$ C $	Dec	PuncCtx	PuncInt
$\ell = 2^t$ time slots (PFSKEM/TB-BFKEM)						
GM	$(t + 5) \mathbb{G}_1 $	$(2t + 3p + 5) \mathbb{G}_2 $	$3 \mathbb{G}_1 + \mathbb{G}_T $	$O(p)$	$O(1)$	$O(t^2)$
GHJL	$(t + 35) \mathbb{G}_2 $	$\leq 3(p \cdot 2\lambda + t) \mathbb{G}_2 $	$6 \mathbb{G}_1 + 2 \mathbb{Z}_q $	$O(\lambda^2)$	$O(\lambda^2)$	$O(t^2)$
Ours	$(t + 4) \mathbb{G}_2 $	$(2me^{-kp/m} + t(2 + t)) \mathbb{G}_2 $	$2 \mathbb{G}_1 + (4k + 2)\lambda$	$O(k)$	$O(k)$	$O(t^2 + m)$

Table 3. Overview of the existing approaches to PFSKEM (resp. TB-BFKEM). By p we denote the number of times a secret key has already been punctured, and by ℓ we denote the maximum number of time slots. We consider the GHJL [33] instantiation with the BKP-HIBE of [9], the GM [31] and our instantiations with the BBG-HIBE [11], though other HIBE schemes may lead to different parameters. Finally, note that $p \leq 2^{20}$, k and m refer to the parameters in the Bloom filter, where k is some orders of magnitude smaller than λ , i.e., $k = 10$ vs. $\lambda = 128$, and $|\mathbb{G}_i|$ denotes the bitlength of an element from \mathbb{G}_i .

To quickly summarize the schemes: The most interesting characteristic of our approach compared to previous approaches is that our scheme allows to offload all expensive operations to an offline phase, i.e., to the puncturing of time intervals. Here, in addition to the $O(w^2)$ operations which are common to all existing approaches, we have to generate a number of keys, linear in the size m of the Bloom filter. We believe that accepting this additional overhead in favor of blazing fast online puncturing and decryption operations is a viable tradeoff. For the online phase, our approach has a ciphertext size depending on k (where $k = 10$ is a reasonable choice), decryption depends on k , the secret key shrinks with increasing amount of puncturings and one does only require to securely delete secret keys during puncturing.¹² In contrast, decryption and puncturing in GHJL is highly inefficient and takes several seconds to minutes on decent hardware for reasonable deployment parameters as it involves a large amount of $O(\lambda^2)$ HIBE delegations and consequently expensive group operations. In the GM scheme¹³, puncturing is efficient, but the size of the secret key and thus cost of decryption grows in the number of puncturings p . Hence, it gets impractical very soon. More precisely,

¹² First, note that all constructions have to implement a secure-delete functionality for secret keys within puncturing anyways. Second, note that the question regarding which data structures to choose so that implementations can actually benefit from the shrinking keys is out of scope here. Note that low-level optimizations for sparse files are typically implemented by modern operating systems and file systems [28]. This way one would even benefit when the memory for the secret keys is allocated as a single monolithic block and the deleted keys are simply zeroed out.

¹³ Although GM supports an arbitrary number d of tags in a ciphertext, we consider the scheme with only using a single tag (which is actually favourable for the scheme) to be comparable to GHJL as well as our approach.

cost of decryption requires a number of pairing evaluations that depends on the number of puncturings, and can be in the order of 2^{20} for realistic deployment parameters.

7 Conclusion

In this paper we introduced the new notion of Bloom filter encryption as a variant of puncturable encryption which tolerates a non-negligible correctness error. We presented various BFKEM constructions. The first one is a simple and very efficient construction which builds upon ideas known from the Boneh-Franklin IBE. It achieves constant size public keys. The second one is a generic construction from CP-ABEs, where a suitable choice of the CP-ABE achieves constant size ciphertexts are available. Those constant size ciphertexts, however, come at the cost of larger keys in existing schemes. The third one is a generic construction from IBBEs, which can be instantiated with the IBBE by Delerablée [22]. This instantiation simultaneously yields constant size ciphertexts and compact public keys. Furthermore, we extended the notion of BFKEM to the forward-secrecy setting and also presented a construction of what we call a time-based BFKEM (TB-BFKEM). This construction is based on HIBEs and in particular can be instantiated very efficiently using the Boneh-Boyen-Goh HIBE [11]. Our time-based BFKEM can directly be used to instantiate forward-secret 0-RTT key exchange (fs 0-RTT KE) as in [33].

From a practical viewpoint, our motivation stems from the observation that forward-secret 0-RTT KE requires very efficient decryption and puncturing. Our framework—for the first time—allows to realize practical forward-secret 0-RTT KE, even for larger server loads: while we only require to delete secret keys upon puncturing, puncturing in [33] requires, besides deleting secret-key components, additional computations in the order of seconds to minutes on decent hardware. Likewise, when using [31] in the forward-secret 0-RTT KE protocol given in [33], one requires computations in the order of the current number of puncturings upon decryption, while we achieve decryption independent of this number. Finally, we believe that BFE will find applications beyond forward-secret 0-RTT KE protocols.

Acknowledgments. We thank the anonymous reviewers of the Journal of Cryptology for their valuable comments. This research was supported by the European commission through H2020 under grant agreement n°644962 (PRISMACLOUD), grant agreement n°653454 (CREDENTIAL), n°802823 (REWOCRYPT), ECSEL-JU 2018 program under grant agreement n°826610 (COMP4DRONES), project IOT4CPS funded by the Austrian “ICT of the future” program of the Austrian Research Promotion Agency (FFG) and the Federal Ministry of Austria for Climate Action, Environment, Energy, Mobility, Innovation and Technology (BMK), by the Austrian Science Fund (FWF) and netidee SCIENCE under grant agreement P31621-N38 (PROFET), and the German Research Foundation (DFG), project JA 2445/2-1.

References

1. Agrawal, S., Boneh, D., Boyen, X.: Efficient lattice (H)IBE in the standard model. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 553–572. Springer, Heidelberg, Germany, French Riviera (May 30 – Jun 3, 2010)

2. Attrapadung, N., Hanaoka, G., Yamada, S.: Conversions among several classes of predicate encryption and applications to ABE with various compactness tradeoffs. In: Iwata, T., Cheon, J.H. (eds.) ASIACRYPT 2015, Part I. LNCS, vol. 9452, pp. 575–601. Springer, Heidelberg, Germany, Auckland, New Zealand (Nov 30 – Dec 3, 2015)
3. Aviram, N., Gellert, K., Jager, T.: Session resumption protocols and efficient forward security for TLS 1.3 0-RTT. In: Ishai, Y., Rijmen, V. (eds.) EUROCRYPT 2019, Part II. LNCS, vol. 11477, pp. 117–150. Springer, Heidelberg, Germany, Darmstadt, Germany (May 19–23, 2019)
4. Barbulescu, R., Duquesne, S.: Updating key size estimations for pairings. *Journal of Cryptology* 32(4), 1298–1336 (Oct 2019)
5. Barreto, P.S.L.M., Lynn, B., Scott, M.: Constructing elliptic curves with prescribed embedding degrees. In: Cimato, S., Galdi, C., Persiano, G. (eds.) SCN 02. LNCS, vol. 2576, pp. 257–267. Springer, Heidelberg, Germany, Amalfi, Italy (Sep 12–13, 2003)
6. Barreto, P.S.L.M., Naehrig, M.: Pairing-friendly elliptic curves of prime order. In: Preneel, B., Tavares, S. (eds.) SAC 2005. LNCS, vol. 3897, pp. 319–331. Springer, Heidelberg, Germany, Kingston, Ontario, Canada (Aug 11–12, 2006)
7. Bellare, M., Rogaway, P.: Random oracles are practical: A paradigm for designing efficient protocols. In: Denning, D.E., Pyle, R., Ganesan, R., Sandhu, R.S., Ashby, V. (eds.) ACM CCS 93. pp. 62–73. ACM Press, Fairfax, Virginia, USA (Nov 3–5, 1993)
8. Bethencourt, J., Sahai, A., Waters, B.: Ciphertext-policy attribute-based encryption. In: 2007 IEEE Symposium on Security and Privacy. pp. 321–334. IEEE Computer Society Press, Oakland, CA, USA (May 20–23, 2007)
9. Blazy, O., Kiltz, E., Pan, J.: (Hierarchical) identity-based encryption from affine message authentication. In: Garay, J.A., Gennaro, R. (eds.) CRYPTO 2014, Part I. LNCS, vol. 8616, pp. 408–425. Springer, Heidelberg, Germany, Santa Barbara, CA, USA (Aug 17–21, 2014)
10. Bloom, B.H.: Space/time trade-offs in hash coding with allowable errors. *Commun. ACM* 13(7), 422–426 (1970)
11. Boneh, D., Boyen, X., Goh, E.J.: Hierarchical identity based encryption with constant size ciphertext. In: Cramer, R. (ed.) EUROCRYPT 2005. LNCS, vol. 3494, pp. 440–456. Springer, Heidelberg, Germany, Aarhus, Denmark (May 22–26, 2005)
12. Boneh, D., Franklin, M.K.: Identity-based encryption from the Weil pairing. In: Kilian, J. (ed.) CRYPTO 2001. LNCS, vol. 2139, pp. 213–229. Springer, Heidelberg, Germany, Santa Barbara, CA, USA (Aug 19–23, 2001)
13. Boneh, D., Lynn, B., Shacham, H.: Short signatures from the Weil pairing. In: Boyd, C. (ed.) ASIACRYPT 2001. LNCS, vol. 2248, pp. 514–532. Springer, Heidelberg, Germany, Gold Coast, Australia (Dec 9–13, 2001)
14. Boyd, C., Gellert, K.: A modern view on forward security. *The Computer Journal* (2020), to appear
15. Canetti, R., Halevi, S., Katz, J.: A forward-secure public-key encryption scheme. In: Biham, E. (ed.) EUROCRYPT 2003. LNCS, vol. 2656, pp. 255–271. Springer, Heidelberg, Germany, Warsaw, Poland (May 4–8, 2003)
16. Canetti, R., Halevi, S., Katz, J.: Chosen-ciphertext security from identity-based encryption. In: Cachin, C., Camenisch, J. (eds.) EUROCRYPT 2004. LNCS, vol. 3027, pp. 207–222. Springer, Heidelberg, Germany, Interlaken, Switzerland (May 2–6, 2004)
17. Canetti, R., Raghuraman, S., Richelson, S., Vaikuntanathan, V.: Chosen-ciphertext secure fully homomorphic encryption. In: *Public-Key Cryptography - PKC 2017*. pp. 213–240 (2017)
18. Chen, C., Chen, J., Lim, H.W., Zhang, Z., Feng, D., Ling, S., Wang, H.: Fully secure attribute-based systems with short ciphertexts/signatures and threshold access structures. In: Dawson, E. (ed.) CT-RSA 2013. LNCS, vol. 7779, pp. 50–67. Springer, Heidelberg, Germany, San Francisco, CA, USA (Feb 25 – Mar 1, 2013)

19. Chen, J., Wee, H.: Fully, (almost) tightly secure IBE and dual system groups. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013, Part II. LNCS, vol. 8043, pp. 435–460. Springer, Heidelberg, Germany, Santa Barbara, CA, USA (Aug 18–22, 2013)
20. Cohen, A., Holmgren, J., Nishimaki, R., Vaikuntanathan, V., Wichs, D.: Watermarking cryptographic capabilities. In: Wichs, D., Mansour, Y. (eds.) 48th ACM STOC. pp. 1115–1127. ACM Press, Cambridge, MA, USA (Jun 18–21, 2016)
21. Dallmeier, F., Drees, J.P., Gellert, K., Handirk, T., Jager, T., Klauke, J., Nachtigall, S., Renzelmann, T., Wolf, R.: Forward-secure 0-RTT goes live: Implementation and performance analysis in QUIC. Cryptology ePrint Archive, Report 2020/824 (2020), <https://eprint.iacr.org/2020/824>
22. Delerablée, C.: Identity-based broadcast encryption with constant size ciphertexts and private keys. In: Kurosawa, K. (ed.) ASIACRYPT 2007. LNCS, vol. 4833, pp. 200–215. Springer, Heidelberg, Germany, Kuching, Malaysia (Dec 2–6, 2007)
23. Derler, D., Jager, T., Slamanig, D., Striecks, C.: Bloom filter encryption and applications to efficient forward-secret 0-RTT key exchange. In: Nielsen, J.B., Rijmen, V. (eds.) EUROCRYPT 2018, Part III. LNCS, vol. 10822, pp. 425–455. Springer, Heidelberg, Germany, Tel Aviv, Israel (Apr 29 – May 3, 2018)
24. Derler, D., Krenn, S., Lorünser, T., Ramacher, S., Slamanig, D., Striecks, C.: Revisiting proxy re-encryption: Forward secrecy, improved security, and applications. In: Abdalla, M. (ed.) PKC 2018. LNCS, Springer (2018)
25. Fujisaki, E., Okamoto, T.: Secure integration of asymmetric and symmetric encryption schemes. In: Wiener, M.J. (ed.) CRYPTO'99. LNCS, vol. 1666, pp. 537–554. Springer, Heidelberg, Germany, Santa Barbara, CA, USA (Aug 15–19, 1999)
26. Gellert, K.: Construction and security analysis of 0-RTT protocols. PhD thesis (2020)
27. Gentry, C., Silverberg, A.: Hierarchical ID-based cryptography. In: Zheng, Y. (ed.) ASIACRYPT 2002. LNCS, vol. 2501, pp. 548–566. Springer, Heidelberg, Germany, Queenstown, New Zealand (Dec 1–5, 2002)
28. Giampaolo, D.: Practical File System Design with the Be File System. Morgan Kaufmann Publishers Inc., 1st edn. (1998)
29. Goel, A., Gupta, P.: Small subset queries and bloom filters using ternary associative memories, with applications. In: Misra, V., Barford, P., Squillante, M.S. (eds.) SIGMETRICS 2010, Proceedings of the 2010 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems, New York, New York, USA, 14-18 June 2010. pp. 143–154. ACM (2010), <https://doi.org/10.1145/1811039.1811056>
30. Goyal, V., Pandey, O., Sahai, A., Waters, B.: Attribute-based encryption for fine-grained access control of encrypted data. In: Juels, A., Wright, R.N., De Capitani di Vimercati, S. (eds.) ACM CCS 2006. pp. 89–98. ACM Press, Alexandria, Virginia, USA (Oct 30 – Nov 3, 2006), available as Cryptology ePrint Archive Report 2006/309
31. Green, M.D., Miers, I.: Forward secure asynchronous messaging from puncturable encryption. In: 2015 IEEE Symposium on Security and Privacy. pp. 305–320. IEEE Computer Society Press, San Jose, CA, USA (May 17–21, 2015)
32. Groth, J.: Simulation-sound NIZK proofs for a practical language and constant size group signatures. In: Lai, X., Chen, K. (eds.) ASIACRYPT 2006. LNCS, vol. 4284, pp. 444–459. Springer, Heidelberg, Germany, Shanghai, China (Dec 3–7, 2006)
33. Günther, F., Hale, B., Jager, T., Lauer, S.: 0-RTT key exchange with full forward secrecy. In: Coron, J., Nielsen, J.B. (eds.) EUROCRYPT 2017, Part III. LNCS, vol. 10212, pp. 519–548. Springer, Heidelberg, Germany, Paris, France (Apr 30 – May 4, 2017)
34. Günther, F., Hale, B., Jager, T., Lauer, S.: 0-RTT key exchange with full forward secrecy. Cryptology ePrint Archive, Report 2017/223 (2017), <http://eprint.iacr.org/2017/223>

35. Hale, B., Jager, T., Lauer, S., Schwenk, J.: Simple security definitions for and constructions of 0-RTT key exchange. In: Gollmann, D., Miyaji, A., Kikuchi, H. (eds.) ACNS 17. LNCS, vol. 10355, pp. 20–38. Springer, Heidelberg, Germany, Kanazawa, Japan (Jul 10–12, 2017)
36. Halevi, S., Krawczyk, H.: One-pass HMQV and asymmetric key-wrapping. In: Catalano, D., Fazio, N., Gennaro, R., Nicolosi, A. (eds.) PKC 2011. LNCS, vol. 6571, pp. 317–334. Springer, Heidelberg, Germany, Taormina, Italy (Mar 6–9, 2011)
37. Hofheinz, D., Hövelmanns, K., Kiltz, E.: A modular analysis of the Fujisaki-Okamoto transformation. In: Kalai, Y., Reyzin, L. (eds.) TCC 2017, Part I. LNCS, vol. 10677, pp. 341–371. Springer, Heidelberg, Germany, Baltimore, MD, USA (Nov 12–15, 2017)
38. Hofheinz, D., Hövelmanns, K., Kiltz, E.: A modular analysis of the Fujisaki-Okamoto transformation. Cryptology ePrint Archive, Report 2017/604 (2017), <http://eprint.iacr.org/2017/604>
39. Kim, T., Barbulescu, R.: Extended tower number field sieve: A new complexity for the medium prime case. In: Robshaw, M., Katz, J. (eds.) CRYPTO 2016, Part I. LNCS, vol. 9814, pp. 543–571. Springer, Heidelberg, Germany, Santa Barbara, CA, USA (Aug 14–18, 2016)
40. Lauer, S., Gellert, K., Merget, R., Handirk, T., Schwenk, J.: T0RTT: Non-interactive immediate forward-secret single-pass circuit construction. Proceedings on Privacy Enhancing Technologies 2020(2), 336–357 (2020)
41. Lovett, S., Porat, E.: A space lower bound for dynamic approximate membership data structures. SIAM Journal on Computing 42(6), 2182–2196 (2013)
42. Menezes, A., Sarkar, P., Singh, S.: Challenges with assessing the impact of NFS advances on the security of pairing-based cryptography. IACR Cryptology ePrint Archive 2016, 1102 (2016), <http://eprint.iacr.org/2016/1102>
43. Naor, M., Yogev, E.: Bloom filters in adversarial environments. In: Gennaro, R., Robshaw, M.J.B. (eds.) CRYPTO 2015, Part II. LNCS, vol. 9216, pp. 565–584. Springer, Heidelberg, Germany, Santa Barbara, CA, USA (Aug 16–20, 2015)
44. Ostrovsky, R., Sahai, A., Waters, B.: Attribute-based encryption with non-monotonic access structures. In: Ning, P., De Capitani di Vimercati, S., Syverson, P.F. (eds.) ACM CCS 2007. pp. 195–203. ACM Press, Alexandria, Virginia, USA (Oct 28–31, 2007)
45. Rescorla, E.: The Transport Layer Security (TLS) Protocol Version 1.3. RFC 8446 (Aug 2018), <https://rfc-editor.org/rfc/rfc8446.txt>
46. Sun, S.F., Sakzad, A., Steinfeld, R., Liu, J., Gu, D.: Public-key puncturable encryption: Modular and compact constructions. Cryptology ePrint Archive, Report 2020/126 (2020), <https://eprint.iacr.org/2020/126>
47. Thomson, M., Iyengar, J.: QUIC: A UDP-Based Multiplexed and Secure Transport. Internet-Draft draft-ietf-quic-transport-02, Internet Engineering Task Force (Mar 2017), <https://datatracker.ietf.org/doc/html/draft-ietf-quic-transport-02>, work in Progress
48. Waters, B.: Dual system encryption: Realizing fully secure IBE and HIBE under simple assumptions. In: Halevi, S. (ed.) CRYPTO 2009. LNCS, vol. 5677, pp. 619–636. Springer, Heidelberg, Germany, Santa Barbara, CA, USA (Aug 16–20, 2009)
49. Wu, D.J., Taly, A., Shankar, A., Boneh, D.: Privacy, discovery, and authentication for the internet of things. In: Askoxylakis, I.G., Ioannidis, S., Katsikas, S.K., Meadows, C.A. (eds.) ESORICS 2016, Part II. LNCS, vol. 9879, pp. 301–319. Springer, Heidelberg, Germany, Heraklion, Greece (Sep 26–30, 2016)

A Formal Definitions for Bloom Filter Encryption

Definition 27 (Bloom Filter Encryption). A Bloom Filter encryption (BFE) scheme is a tuple $(\text{KGen}, \text{Enc}, \text{Punc}, \text{Dec})$ of PPT algorithms:

$\text{KGen}(1^\lambda, m, k)$: Takes as input a security parameter λ , parameters m and k and outputs a secret and public key (sk, pk) .

$\text{Enc}(\text{pk}, M)$: Takes as input a public key pk , a message $M \in \mathcal{M}$ and outputs a ciphertext C .

$\text{Punc}(\text{sk}, C)$: Takes as input a secret key sk , a ciphertext C and outputs an updated secret key sk' .

$\text{Dec}(\text{sk}, C)$: Takes as input a secret key sk , a ciphertext C and outputs a message $M \in \mathcal{M}$ or \perp if decryption fails.

Definition 28 (Correctness). We require that the following holds for all $\lambda, m, k \in \mathbb{N}$, all arbitrary sequences of messages $(M_1, \dots, M_n, M^*) \in \mathcal{M}^{n+1}$ and any $(\text{sk}, \text{pk}) \stackrel{\$}{\leftarrow} \text{KGen}(1^\lambda, m, k)$.

For any (arbitrary interleaved) sequence of invocations of

$$\text{sk}_{j+1} \stackrel{\$}{\leftarrow} \text{Punc}(\text{sk}_j, C_j),$$

where $j \in \{1, \dots, n\}$, $\text{sk}_1 := \text{sk}$, and $C_j \stackrel{\$}{\leftarrow} \text{Enc}(\text{pk}, M_j)$, it holds that

$$\Pr[\text{Dec}(\text{sk}_{n+1}, C^*) \neq M^*] \leq \left(1 - e^{-\frac{(n+1/2)k}{m-1}}\right)^k + \epsilon(\lambda),$$

where $C^* \stackrel{\$}{\leftarrow} \text{Enc}(\text{pk}, M^*)$ and $\epsilon(\cdot)$ is a negligible function in λ . The probability is over the random coins of KGen , Punc , and Enc .

Note that the correctness is essentially our BFKEM definition ported to the encryption setting. As in Section 2.3, we can define the extended correctness, separable randomness, publicly checkable puncturing and γ -spreadness. As this is straightforward, we do not explicitly repeat the definitions here.

Security notions. Subsequently, in Figure 8, we define the IND-CPA/IND-CCA2-experiment for BFE. The experiment is identical to IND-CPA/IND-CCA2-security for conventional public-key encryption, but in addition, the adversary in the second phase can arbitrarily puncture the secret key and retrieve the punctured secret key as long as the key has been punctured on the challenge ciphertext C^* . This still should not help the adversary to obtain any information about the message hidden in C^* .

Definition 29 (IND-T Security of BFE). We define the advantage of an adversary \mathcal{A} in the IND-T-experiment $\text{Exp}_{\mathcal{A}, \text{BFE}}^{\text{IND-T}}(\lambda, m, k)$ as

$$\text{Adv}_{\mathcal{A}, \text{BFE}}^{\text{IND-T}}(\lambda, m, k) := \left| \Pr \left[\text{Exp}_{\mathcal{A}, \text{BFE}}^{\text{IND-T}}(\lambda, m, k) = 1 \right] - \frac{1}{2} \right|.$$

A Bloom Filter encryption scheme BFE is IND-T-secure, $\text{T} \in \{\text{IND-CPA}, \text{IND-CCA2}\}$, if $\text{Adv}_{\mathcal{A}, \text{BFE}}^{\text{IND-T}}(\lambda, m, k)$ is a negligible function in λ for all $m, k > 0$ and all PPT adversaries \mathcal{A} .

$\text{Exp}_{\mathcal{A}, \text{BFE}}^{\text{IND-T}}(\lambda, m, k)$:
 $(\text{sk}, \text{pk}) \leftarrow^{\$} \text{KGen}(1^\lambda, m, k)$
 $b \leftarrow^{\$} \{0, 1\}, \mathcal{Q} \leftarrow \emptyset$
 $(M_0, M_1) \leftarrow^{\$} \mathcal{A}^{\mathcal{O}}(\text{pk})$
 where $\mathcal{O} \leftarrow \text{Dec}(\text{sk}, \cdot)$ if $T = \text{IND-CCA2}$ and $\mathcal{O} \leftarrow \emptyset$ otherwise.
 if $M_0, M_1 \notin \mathcal{M} \vee |M_0| \neq |M_1|$, then let $C^* \leftarrow \perp$
 else, let $C^* \leftarrow^{\$} \text{Enc}(\text{pk}, M_b)$
 $b^* \leftarrow^{\$} \mathcal{A}^{\mathcal{O}, \text{Punc}(\text{sk}, \cdot), \text{Corr}}(C^*)$
 where $\mathcal{O} \leftarrow \text{Dec}'(\text{sk}, \cdot)$ if $T = \text{IND-CCA2}$ and $\mathcal{O} \leftarrow \emptyset$ otherwise.
 $\text{Dec}'(\text{sk}, C)$ behaves as Dec but returns \perp if $C = C^*$
 $\text{Punc}(\text{sk}, C)$ runs $\text{sk} \leftarrow^{\$} \text{Punc}(\text{sk}, C)$ and $\mathcal{Q} \leftarrow \mathcal{Q} \cup \{C\}$
 Corr returns sk if $C^* \in \mathcal{Q}$ and \perp otherwise
 return 1, if $b^* = b$
 return 0

Fig. 8. IND-T-security for BFE: $T \in \{\text{IND-CPA}, \text{IND-CCA2}\}$.

B Proof of Lemma 1

We begin by computing the expected number of bits set to one in the BF after α random elements have been added. Let $T_0 = b_0 b_1 \dots b_m$ be the sequence of bits of a BF with size m . After initialization, we have $b_i := 0$ for all $i \in [m]$. With each added element we set up to k bits to one, that is we sample k elements $a_1, \dots, a_k \in [m]$ with replacement and set $b_{a_i} := 1$. At the end of α time steps we have at most αk bits equal to one and at least 1 bit equal to one. Let X_i^t be the event that b_i is set at time t , that is, $X_i^t = 1 \implies b_i = 1$ has already been set to one at time t . Thus the number bits set to one after α time steps is

$$\sum_{i=1}^m X_i^{\alpha k}.$$

To compute the expected number of bits set to one at time t , we need to compute the expected value of each X_i . Then the expected number of bits set to one after α time steps is

$$\begin{aligned} X^\alpha &= \sum_{i=1}^m X_i^\alpha = \sum_{i=1}^m \Pr[X_i^\alpha] = \sum_{i=1}^m (1 - \Pr[\overline{X_i^\alpha}]) = \sum_{i=1}^m 1 - \left(1 - \frac{1}{m}\right)^{\alpha k} \\ &= m \cdot \left(1 - \left(1 - \frac{1}{m}\right)^{\alpha k}\right). \end{aligned}$$

We can now bound the probability by applying a simple combinatorial argument. When choosing a random bit b_i , we have a probability of X^α/m to choose an index i with $b_i = 1$. Independently repeating this process k times, leads us to the expected false-positive probability of a random element $u \in \mathcal{U}$ being recognized by the Bloom filter:

$$\Pr[\text{BFCheck}(H, T_n, u)] = \left(\frac{m \cdot \left(1 - \left(1 - \frac{1}{m}\right)^{\alpha k}\right)}{m}\right)^k = \left(1 - \left(1 - \frac{1}{m}\right)^{\alpha k}\right)^k.$$

C Growth of the False-Positive Probability

Figures 9, 10, and 11 show the growth of the false-positive probability after α elements have been added to a Bloom filter with $n \in 2^{16}, 2^{24}, 2^{30}$ for $k \in \{8, 11, 17, 21\}$. Note that both axes are scaled logarithmically.

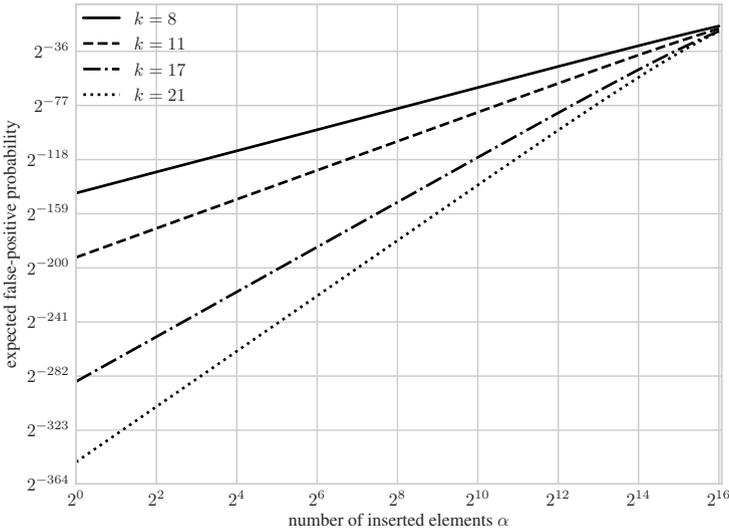


Fig. 9. The false-positive probability of a random element after α elements have been added to a Bloom filter with $n = 2^{16}$ for $k \in \{8, 11, 17, 21\}$.

D Identity-based Broadcast Encryption with Constant Size Ciphertexts and Private Keys

The subsequent construction is the identity-based broadcast encryption scheme by Delerablée [22]. The main advantages of her scheme are the constant size ciphertexts and private keys.

Let $(p, e, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T) \leftarrow^{\$} \text{BilGen}(1^\lambda)$ be public parameters of a bilinear map $e : \mathbb{G}_1 \times \mathbb{G}_2 \rightarrow \mathbb{G}_T$ with prime orders p and $|p| = \lambda$. Let $\mathcal{H} : \mathbb{Z}_q^* \rightarrow \mathbb{Z}_q^*$ be a cryptographic hash function. We construct an identity-based broadcast encryption scheme $\text{IBBE} = (\text{Setup}, \text{Extract}, \text{Enc}, \text{Dec})$ as follows:

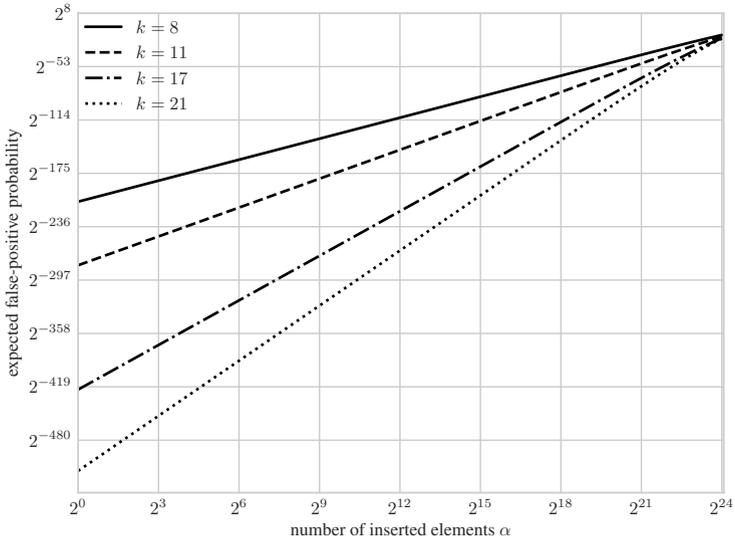


Fig. 10. The false-positive probability of a random element after α elements have been added to a Bloom filter with $n = 2^{24}$ for $k \in \{8, 11, 17, 21\}$.

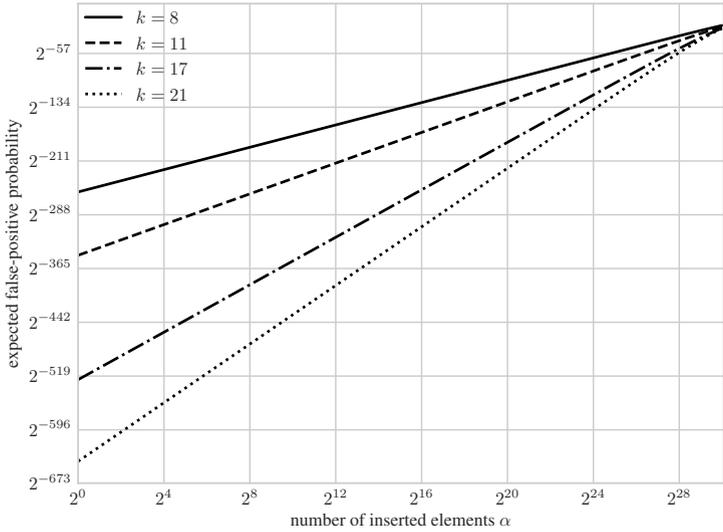


Fig. 11. The false-positive probability of a random element after α elements have been added to a Bloom filter with $n = 2^{30}$ for $k \in \{8, 11, 17, 21\}$.

Setup(λ, k) : The key generation algorithm chooses two generators $g_1 \in \mathbb{G}_1$ and $g_2 \in \mathbb{G}_2$ and a secret value $\gamma \leftarrow_{\mathbb{S}} \mathbb{Z}_q^*$. Finally, we set and output the public key pk and master secret key msk as

$$\text{pk} := \left(w = g_1^\gamma, v = e(g_1, g_2), g_2^\gamma, \dots, g_2^{\gamma^k} \right) \text{ and } \text{msk} := \gamma.$$

Extract(msk, ID) : The key extraction algorithm takes as input the master secret key $\text{msk} = \gamma$ and an identity ID . Output is an extracted secret key

$$\text{sk}_{\text{ID}} = g_1^{\frac{1}{\gamma + \mathcal{H}(\text{ID})}}.$$

Enc(pk, \mathcal{S}) : Given a public key $\text{pk} = (w, v, g_2^\gamma, \dots, g_2^{\gamma^k})$ and a set of identities $\mathcal{S} = \{\text{ID}_j\}_{j \in [s]}$ with $s \leq k$, it samples a symmetric key K by choosing a secret value $\rho \leftarrow_{\mathbb{S}} \mathbb{Z}_q$ and computing $K := v^\rho = e(g_1, g_2)^\rho$. Finally, the algorithm computes a ciphertext $C = (c_1, c_2)$ with

$$c_1 := w^{-\rho} \text{ and } c_2 := g_2^{\rho \cdot \prod_{j=1}^s (\gamma + \mathcal{H}(\text{ID}_j))}.$$

It outputs (K, C) .

Dec($\text{sk}_{\text{ID}_j}, \mathcal{S}, C$) : Given a ciphertext $C = (c_1, c_2)$, it computes

$$K = \left(e \left(c_1, g_2^{p_{i,\mathcal{S}}(\gamma)} \right) \cdot e(\text{sk}_{\text{ID}_j}, c_2) \right)^{\frac{1}{\prod_{j=1, j \neq i}^s (\gamma + \mathcal{H}(\text{ID}_j))}}, \text{ where}$$

$$p_{i,\mathcal{S}}(\gamma) = \frac{1}{\gamma} \left(\prod_{j=1, j \neq i}^s (\gamma + \mathcal{H}(\text{ID}_j)) - \prod_{j=1, j \neq i}^s \mathcal{H}(\text{ID}_j) \right)$$

and returns K .

Note that indeed the ciphertext is $C \in \mathbb{G}_1 \times \mathbb{G}_2$ and an extracted secret key is $\text{sk}_{\text{ID}_j} \in \mathbb{G}_1$.

Remark on computation of $p_{i,\mathcal{S}}$. The decapsulation algorithm uses a function p whose description is dependent of γ . However, neither γ nor any other secret value is needed to compute it. Instead we can compute $g_2^{p_{i,\mathcal{S}}}$ by only using public values.

Let $c_v(a_1, \dots, a_n)$ be a function that on input of n values returns the sum of all possible pairwise distinct v -combinations of the input values, i.e. $c_2(a, b, c) = ab + ac + bc$, and let $\mathcal{S}_i = \{\mathcal{H}(\text{ID}_j) | j \in \mathcal{S} \setminus \{i\}\}$. Then we can rewrite

$$\begin{aligned} p_{i,\mathcal{S}}(\gamma) &= \frac{1}{\gamma} \left(\prod_{j=1, j \neq i}^s (\gamma + \mathcal{H}(\text{ID}_j)) - \prod_{j=1, j \neq i}^s \mathcal{H}(\text{ID}_j) \right) \\ &= \frac{1}{\gamma} \left(\gamma^{s-1} + \gamma^{s-2} c_1(\mathcal{S}_i) + \dots + \gamma c_{s-2}(\mathcal{S}_i) + c_{s-1}(\mathcal{S}_i) - \prod_{j=1, j \neq i}^s \mathcal{H}(\text{ID}_j) \right) \\ &= \gamma^{s-2} + \gamma^{s-3} c_1(\mathcal{S}_i) + \dots + c_{s-2}(\mathcal{S}_i). \end{aligned}$$

In our case it suffices to compute

$$\begin{aligned} g_2^{p_{i,S}(\gamma)} &= g_2^{\gamma^{s-2} + \gamma^{s-3}c_1(S_i) + \dots + c_{s-2}(S_i)} \\ &= g_2^{\gamma^{s-2}} \cdot \left(g_2^{\gamma^{s-3}}\right)^{c_1(S_i)} \cdot \dots \cdot g_2^{c_{s-2}(S_i)}. \end{aligned}$$

As $s \leq k$, all g_2^γ -like values are publicly known and thus, $g_2^{p_{i,S}}$ is computable without any secret knowledge. The given argument also holds for the computation of ciphertext c_2 in the key encapsulation.

Security of the IBBE. In [22] Delerablée also analyzes the security of the above scheme under the so called (g, f, F) -GDDHE assumption. This is a variant of a generalization of the Diffie-Hellman Exponent Assumption introduced in [11] and analyzed in the generic bilinear group model in [22]. For the sake of completeness, we restate the theorem from [22].

Theorem 7. *From each efficient adversary \mathcal{B} against IND-sID-CPA-security of the IBBE scheme, we can construct an efficient algorithm \mathcal{A} against the (g, f, F) -GDDHE assumption with advantage*

$$\mathbf{Adv}_{\mathcal{A}}^{\text{GDDHE}}(g, f, F) \geq \frac{1}{2} \cdot \mathbf{Adv}_{\mathcal{B}, \text{IBBE}}^{\text{IND-sID-CPA}}(\lambda, k).$$