Determining the Core Primitive for Optimally Secure Ratcheting*

Fatih Balli[1], Paul Rösler[2], Serge Vaudenay[1]

[1] LASEC, École polytechnique fédérale de Lausanne
{firstname.lastname}@epfl.ch
[2] Chair for Network and Data Security, Ruhr University Bochum
paul.roesler@rub.de

Abstract. After ratcheting attracted attention mostly due to practical real-world protocols, recently a line of work studied ratcheting as a primitive from a theoretic point of view. Literature in this line, pursuing the strongest security of ratcheting one can hope for, utilized for constructions strong, yet inefficient key-updatable primitives – based on hierarchical identity based encryption (HIBE). As none of these works formally justified utilizing these building blocks, we answer the yet open question under which conditions their use is actually *necessary*.

We revisit these strong notions of ratcheted key exchange (RKE), and propose a more realistic (slightly stronger) security definition. In this security definition, both exposure of participants' local secrets and attacks against executions' randomness are considered. While these two attacks were partially considered in previous work, we are the first to unify them cleanly in a natural game based notion.

Our definitions are based on the systematic RKE notion by Poettering and Rösler (CRYPTO 2018). Due to slight (but meaningful) changes to regard attacks against randomness, we are ultimately able to show that, in order to fulfill strong security for RKE, public key cryptography with (independently) updatable key pairs is a necessary building block. Surprisingly, this implication already holds for the simplest RKE variant.

Hence, (1) we model optimally secure RKE under randomness manipulation to cover realistic attacks, (2) we (provably) extract the core primitive that is necessary to realize strongly secure RKE, and (3) our results indicate which relaxations in security allow for constructions that only rely on standard public key cryptography.

1 Introduction

The term "ratcheting" as well as the underlying concept of continuously updating session secrets for secure long-term communication settings originates from real-world messaging protocols [13,15,14]. In these protocols, first forward-secrecy [15] and later security after state exposures [14] (also known as future secrecy, backward secrecy, or post-compromise security) were aimed to be achieved

* The full version [2] of this article is available as entry 2020/148 in the IACR eprint archive.

as the exposure of the devices' local states was considered a practical threat. The main motivation behind this consideration is the typical lifetime of sessions in messaging apps. As messaging apps are nowadays usually run on smartphones, the lifetime of messaging sessions is proportional to the ownership duration of a smartphone (typically several years). Due to the long lifetime of sessions and the mobile use of smartphones, scenarios, in which the local storage – containing the messaging apps' secret state – can be exposed to an attacker, are extended in comparison to use cases of other cryptographic protocols.

Practical Relevance of Randomness Manipulation
In addition to exposures of locally stored state secrets, randomness for generating (new) secrets is often considered vulnerable. This is motivated by numerous attacks in practice against randomness sources (e.g., [9]), randomness generators (e.g., [19,5]), or exposures of random coins (e.g., [18]). Most theoretic approaches try to model this threat by allowing an adversary to *reveal* attacked random coins of a protocol execution (as it was also conducted in related work on ratcheting). This, however, assumes that the attacked protocol honestly and uniformly samples its random coins (either from a high-entropy source or using a random oracle) and that these coins are only afterwards leaked to the attacker. In contrast, practically relevant attacks against bad randomness generators or low-entropy sources (e.g., [9,19,5]) change the distribution from which random coins are sampled. Consequently, this threat is only covered by a security model if considered adversaries are also allowed to *influence* the execution's (distribution of) random coins. Thus, it is important to consider randomness *manipulation* (instead of reveal), if attacks against randomness are regarded practically relevant.

The overall goal of ratcheting protocols is to reduce the effect of any such non-permanent and/or non-fatal attack to a minimum. For example, an ongoing communication under a non-fatal attack should become secure as soon as the adversary ends this attack or countermeasures become effective. Examples for countermeasures are replacing bad randomness generators via software updates, eliminating state exposing viruses, etc. Motivated by this, most widely used messaging apps are equipped with mechanisms to regularly update the local secrets such that only a short time frame of communication is compromised if an adversary was successful due to obtaining local secrets and/or attacking random coins.

Real-World Protocols
The most prominent and most widely deployed real-world ratcheting protocol is the Signal protocol (used by WhatsApp, Skype, and others). The analysis of this protocol in a multi-stage key agreement model [6] was the first theoretic treatment of ratcheting in the literature. Cohn-Gordon et al. [6], however, focus on grasping the precise security that Signal offers rather than generically defining ratcheting as an independent primitive. While the security provided by Signal is sufficient in most real-world scenarios, we focus in this work on the theoretic analysis of the (optimally secure) primitive ratcheting.

Generic Treatment of Ratcheting as a Primitive

In the following we shortly introduce and review previous modeling approaches for strongly secure ratcheting. We thereby abstractly highlight modeling choices that crucially affect the constructions, secure according to these models respectively. Specifically, we indicate why some models can be instantiated with only public key cryptography (PKC) – bypassing our implication result – and others cannot. In Table 1 we summarize this overview.

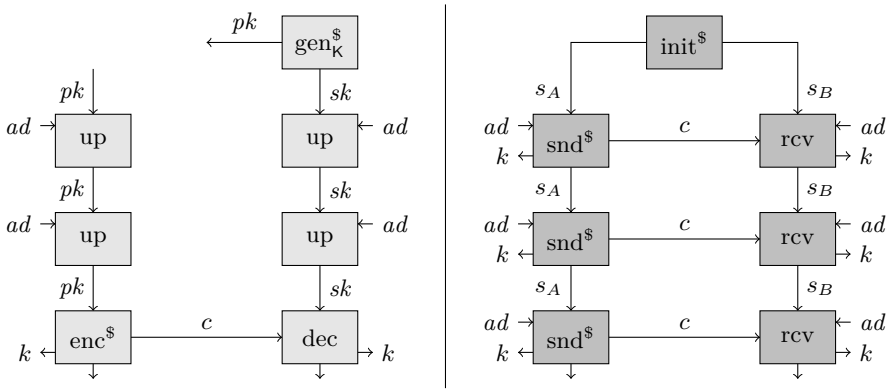

Fig. 1: Conceptual depiction of kuKEM* (on the left) and unidirectional RKE (on the right). '$' in the upper index of an algorithm name denotes that the algorithm runs probabilistically and ad is associated data.

The initial generic work that considers ratcheted key exchange (RKE) as a primitive and defines its syntax, correctness, and security (in a yet impractical variant) is by Bellare et al. [3]. Abstractly, their concept of ratcheted key exchange, depicted in the right part of Figure 1, consist of an initialization that provides two session participants A and B with a state that can then be used by them to repeatedly compute new keys in this session (e.g., for use in higher level protocols). In their restricted communication model, A is allowed to compute new keys with her state and accordingly send ciphertexts to B who can then compute (the same) keys with his state. During these key computations, A's and B's states are updated respectively (to minimize the effect of state exposures). As B can only comprehend key computations from A (on receipt of a ciphertext) but cannot actively initiate the computation of new keys, this variant was later called unidirectional RKE [17]. Beyond this restriction of the communication model, the security definition by Bellare et al. only allows the adversary to expose A's temporary local state secrets, while B's state cannot be exposed (which in turn requires no forward-secrecy with respect to state updates by B). Following Bellare et al., Poettering and Rösler [17,16] propose a revised security definition of unidirectional RKE (URKE: allowing also the exposure of B's state) and extend the communication model to define syntax, correctness, and

security of *sesqui*directional RKE (SRKE: additionally allows B to only send special update ciphertexts to A that do not trigger a new key computation but help him to recover from state exposures) and *bi*directional RKE (BRKE: defines A and B to participate equivalently in the communication). With a similar instantiation as Poettering and Rösler, Jaeger and Stepanovs [10] define security for bidirectional channels under state exposures and randomness reveal.

All of the above mentioned works define security *optimally* with respect to their syntax definition and the adversary's access to the primitive execution (modeled via oracles in the security game). This is reached by declaring secrets insecure *iff* the adversary conducted an unpreventable/trivial attack against them (i.e., a successful attack that no instantiation can prevent). Consequently, fixing syntax and oracle definitions, no stronger security definitions exist.

Relaxed Security Notions

Subsequent to these strongly secure ratcheting notions, multiple weaker formal definitions for ratcheting were proposed that consider special properties such as strong explicit authentication [8], out of order receipt of ciphertexts [1], or primarily target on allowing efficient instantiations [12,4].

	(a) Interaction	**(b) State Exposure**	**(c) Bad Randomness**	**(d) Recovery**
C+ [6]	$\leftrightarrow$	Always allowed	Reveal	Delayed
B+ [3]	$\rightarrow$	Only allowed for A	Reveal	Immediate
PR [17]	$\rightarrow$	Always allowed	Not considered	Immediate
	$\mapsto$	Always allowed	Not considered	Immediate
	$\leftrightarrow$	Always allowed	Not considered	Immediate
JS [10]	$\leftrightarrow$	Always allowed	Reveal	Immediate
DV [8]	$\leftrightarrow$	Always allowed	Not considered	Partial
JMM [12]	$\rightarrow$	Partially restricted	Reveal	(Immediate)
	$\mapsto$	Partially restricted	Reveal	(Immediate)
	$\leftrightarrow$	Partially restricted	Reveal	(Immediate)
ACD [1]	$\leftrightarrow$	Always allowed	Manipulation	Delayed
CDV [4]	$\leftrightarrow$	Always allowed	Not considered	Delayed
This work	$\rightarrow$	Always allowed	Manipulation	Immediate

Table 1: Differences in security notions of ratcheting regarding (a) uni- ($\rightarrow$), sesqui- ($\mapsto$), and bidirectional ($\leftrightarrow$) interaction between A and B, (b) when the adversary is allowed to expose A's and B's state (or when this is *unnecessarily* restricted), (c) the adversary's ability to reveal or manipulate algorithm invocations' random coins, and (d) how soon and how complete recovery from these two attacks into a secure state is required of *secure* constructions (or if *unnecessary* delays or exceptions for recovery are permitted).[1] Recovery from attacks required by Jost et al. [12] is *immediate* in so far as their restrictions of state exposures introduce delays implicitly. Gray marked cells indicate the reason (i.e., relaxations in security) why respective instantiations can rely on standard PKC only (circumventing our implication result). Rows without gray marked cells have no construction based on pure PKC.

[1] '*Unnecessary*' refers to restrictions beyond those that are immediately implied by optimal security definitions (that only restrict the adversary with respect to unpreventable/trivial attacks).

While these works are syntactically similar, we shortly sketch their different relaxations regarding security – making their security notions sub-optimal. Durak and Vaudenay [8] and Caforio et al. [4] forbid the adversary to perform impersonation attacks against the communication between A and B during the establishment of a *secure* key. Thus, they do not require recovery from state exposures – which are a part of impersonation attacks – in all possible cases, which we denote as "partial recovery" (see Table 1). Furthermore, both works neglect bad randomness as an attack vector. In the security experiments by Jost et al. [12] and Alwen et al. [1] constructions can delay the recovery from attacks longer than necessary (Jost et al. therefore temporarily forbid the exposure of the local state). Additionally, they do not require the participants' states to become incompatible (immediately) on active attacks against the communication.

Instantiations of Ratcheting

Interestingly, both mentioned *uni*directional RKE instantiations that were defined to depict optimal security [3,17] as well as bidirectional real-world examples such as the Signal protocol (analyzed in [6]), and instantiations of the above named relaxed security notions [8,12,1,4] only rely on standard PKC (cf. rows in Table 1 with gray cells).

In contrast, both mentioned optimally secure bidirectional ratcheting variants (i.e., sesquidirectional and bidirectional RKE [17], and bidirectional strongly secure channel [10]) are based on a strong cryptographic building block, called *key-updatable public key encryption*, which can be built from hierarchical identity based encryption (HIBE). Intuitively, key-updatable public key encryption is standard public key encryption that additionally allows to update public key and secret key independently with respect to some associated data (a conceptual depiction of this is on the left side of Figure 1). Thereby an updated secret key cannot be used to decrypt ciphertexts that were encrypted to previous (or different) versions of this secret key (where versions are defined over the associated data used for updates).

We emphasize a significant difference between key-updatable public key encryption and HkuPke (introduced in [12]): in HkuPke key updates rely on interactive communication between holders of public key and secret key, and associated data for key updates is not fully adversary-controlled. These differences make it strictly weaker, insufficient for optimal security of RKE (on which we further elaborate in Section 3).

Necessity for Strong Building Blocks

Natural questions that arise from this line of work are, whether and under which conditions such strong (HIBE-like) building blocks are not only sufficient but also necessary to instantiate the strong security of (bidirectional) RKE. In order to answer these questions, we build key-updatable public key cryptography from ratcheted key exchange. Consequently we affirm the necessity and provide (sufficient) conditions for relying on these strong building blocks. We therefore minimally adjust the syntax of key-updatable key encapsulation mechanism (kuKEM) [17] and consider the manipulation of algorithm invocations' random coins in our security definitions of kuKEM and RKE.

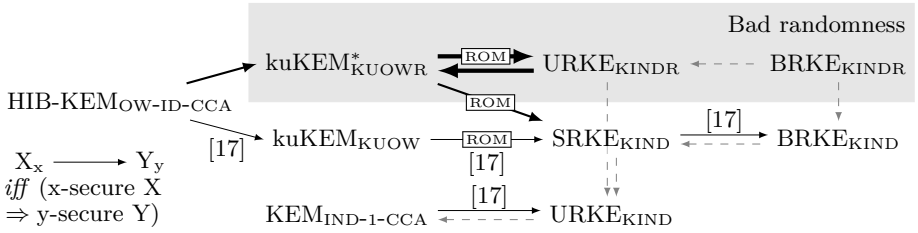

Fig. 2: The contributions of this paper (bold arrows) and their connection to previous work (thin arrows) involving RKE (uni-, sesqui-, and bidirectional) and KEM (standard, hierarchical-identity-based, and key-updatable) primitives. ROM indicates that the proof holds in the random oracle model. $\text{kuKEM}^*_{\text{KUOWR}} \Rightarrow_{\text{ROM}} \text{SRKE}_{\text{KIND}}$ is not formally proven in this paper, but we point out that the proof of $\text{kuKEM}_{\text{KUOW}} \Rightarrow_{\text{ROM}} \text{SRKE}_{\text{KIND}}$ from [17] can be rewound. Gray dashed connections indicate trivial implications (due to strictly weaker syntax or security definitions).

As a result we show that (see Figure 2):

- kuKEM* (with one-way security under manipulation of randomness)[2] $\Rightarrow_{\text{ROM}}$ Unidirectional RKE (with key indistinguishability under manipulation of randomness),
- Unidirectional RKE (with key indistinguishability under manipulation of randomness) $\Rightarrow$ kuKEM* (with one-way security under manipulation of randomness).

Given the security notions established in honest randomness setting and their connections to each other, one would also expect Group RKE $\Rightarrow$ Bidirectional RKE $\Rightarrow$ Sesquidirectional RKE $\Rightarrow$ Unidirectional RKE to follow. Hence, our results indicate that stronger RKE variants also likely require building blocks as hard as kuKEM*. Furthermore, our results show that: One-way security under manipulation of randomness of kuKEM* $\Rightarrow_{\text{ROM}}$ Key indistinguishability of *sesquidirectional* RKE. Interestingly, these results induce that (when considering strong security) ratcheted key exchange requires these strong (HIBE-like) building blocks not only for bidirectional communication settings, but already for the unidirectional case. Both mentioned previous unidirectional RKE schemes can bypass our implication because they forbid exposures of B's state [3] or assume secure randomness [17] (see Table 1). We describe attacks against each of both constructions in our security definition in the full version [2]. Similarly, the discussed relaxed security definitions of ratcheting [6,8,12,1,4] allow for efficient constructions because they restrict the adversary more than necessary (see Table 1). Although our analysis was partially motivated by the use of kuKEM in [17,10], we do not ultimately answer whether these particular constructions

[2] The asterisk at kuKEM* indicates the minimal adjustment to the kuKEM syntax definition from [17]. For the kuKEM* we consider one-way security as it suffices to achieve strong security for RKE.

necessarily relied on it. Rather we provide a clean set of conditions under which RKE and kuKEM clearly imply each other as we do not consider the justification of previous constructions but a clear relation for future work important.

Thus, we show that sufficient conditions for necessarily relying on kuKEM as a building block of RKE are: (a) unrestricted exposure of both parties' local states, (b) consideration of attacks against algorithm invocations' random coins, and (c) required immediate recovery from these two attacks into a secure state by the security definition (i.e., the adversary is only restricted with respect to unpreventable/trivial attacks).[3]

Contributions

The contributions of our work can be summarized as follows:

- We are the first who systematically define optimal security of key-updatable KEM and unidirectional RKE under randomness manipulation (in sections 3 and 4) and thereby consider this practical threat in addition to state exposures in an instantiation-independent notion of RKE. Thereby we substantially enhance the respective models by Poettering and Rösler [17].
- In Section 5, we construct unidirectional RKE generically from a kuKEM* to show that the latter suffices as a building block for the former under manipulation of randomness.
- To show that kuKEM* is not only sufficient but also necessary to build unidirectional RKE (under randomness manipulation), we provide a construction of kuKEM* from a generic unidirectional RKE scheme in Section 6.

With our results we distill the core building block of strongly secure ratcheted key exchange down to its syntax and security definition. This allows further research to be directed towards instantiating kuKEM* schemes that are more familiar and easier in terms of security requirements, rather than attempting to construct seemingly more complex RKE primitives.[4] Simultaneously, our results indicate the cryptographic hardness of ratcheted key exchange and thereby help to systematize and comprehend the security definitions and different dimensions of ratcheting in the literature. As a consequence, our results contribute to a fact-based trade-off between security and efficiency for RKE by providing re-

[3] Note that there may exist further sets of sufficient conditions for relying on kuKEMs since, for example, sesqui- and bidirectional RKE by Poettering and Rösler [17,16] violate condition (b) but base on kuKEMs as well. We refer the reader to Appendix B.2 in [16] for a detailed explanation of why their scheme presumably also must rely on a kuKEM. We leave the identification of further sets of conditions as future work.

[4] For example, the bidirectional channel construction in the proceedings version of [10] is not secure according to the security definition (but a corrected version is published as [11]), in the acknowledgments of [16] it is mentioned that an early submitted version of their construction was also flawed, and for an earlier version of [8] we detected during our work (and informed the authors) that the construction was insecure under bad randomness such that the updated proceedings version (also available as [7]) disregards attacks against randomness entirely. Finally, we detected and reported that the construction of HkuPke in [12] is not even correct.

quirements for relying on heavy building blocks and thereby revealing respective bypasses.

2 Preliminaries

2.1 Notation

By $x \leftarrow y$ we define the assignment of the value of variable y to variable x and thus for a function X, $x \leftarrow \mathrm{X}(y)$ means that x is assigned with the evaluation output of X on input y. We define $\mathtt{T}, \mathtt{F}$ as Boolean values for true and false. The shortcut notion $w \leftarrow x \ \mathtt{?}\ y\ \mathtt{:}\ z$ means that 'if $x = \mathtt{T}$, then $w \leftarrow y$, otherwise $w \leftarrow z$'. For a probabilistic algorithm Y, $x \leftarrow_{\$} \mathrm{Y}(y)$ denotes the probabilistic evaluation of Y on input y with output x and $x \leftarrow \mathrm{Y}(y; r)$ denotes the deterministic evaluation of Y on y with output x where the evaluation's randomness is fixed to r. For a set $\mathcal{X}$, $x \leftarrow_{\$} \mathcal{X}$ is the uniform random sampling of value x from $\mathcal{X}$. We use the shortcut notion $\mathcal{X} \xleftarrow{\cup} \mathcal{Y}$ to denote the union $\mathcal{X} \leftarrow \mathcal{X} \cup \mathcal{Y}$ of sets $\mathcal{X}$ and $\mathcal{Y}$.

Symbol 'ϵ' denotes an empty string and symbol '$\perp$' denotes an undefined element or an output that indicates rejections (thus it is not an element of explicitly defined sets).

By $\mathcal{X}^*$, we denote the set of all lists of arbitrary size whose elements belong to $\mathcal{X}$. We abuse the notation of empty string 'ϵ' by writing $L = \epsilon$ for an empty list L. If an element $x \in \mathcal{X}$ is appended to list L then we denote this by $L \leftarrow L \| x$ (or simply $L \xleftarrow{\shortparallel} x$). Thus, '$\|$' denotes a special concatenation symbol that is not an element of any of the explicitly defined sets. We define relations prefix-or-equal $\preceq$ and strictly-prefix $\prec$ over two lists. For instance, for lists $L, L_0 = L \| x, L_1 = L \| y$ where $x, y \in \mathcal{X}, x \neq y$ we have that $L \preceq L, L \not\prec L, L \prec L_0, L \prec L_1, L_0 \not\preceq L_1, L_1 \not\preceq L_0$ meaning that L is a prefix of L_0 and L_1 but neither of L_0, L_1 is a prefix of the other. By $X[\cdot]$ we denote an associative array.

In our security experiments, that we denote with **Game**, we invoke adversaries via instruction 'Invoke'. These adversaries are denoted by $\mathcal{A}, \mathcal{B}$. Adversaries have access to the security experiment's interface, which is defined by oracles that are denoted by the term **Oracle**. Games are terminated via instructions 'Stop with x' (meaning that x is returned by the game) or 'Reward b' (meaning that the game terminates and returns 1 if $b = \mathtt{T}$). In procedures that we denote by **Proc** and in oracles, we use the shortcut notion 'Require x'. Depending on the procedure's or oracle's number of return values n, that means 'If $x = \mathtt{F}$, then return $\perp^n$'.

2.2 Message Authentication Code

We define a message authentication code to be a set of algorithms $\mathsf{M} = (\mathrm{tag}, \mathrm{vfy}_{\mathsf{M}})$ over a set of symmetric keys $\mathcal{K}$, a message space $\mathcal{M}$, and a tag space $\mathcal{T}$. The syntax is defined as:

$$\mathcal{K} \times \mathcal{M} \rightarrow \mathrm{tag} \rightarrow \mathcal{T}$$

$\mathcal{K} \times \mathcal{M} \times \mathcal{T} \to \text{vfy}_{\mathsf{M}} \to \{\mathtt{T}, \bot\}$

Please note that we define the tag algorithm explicitly deterministic.

For correctness of a MAC we define that for all $k \in \mathcal{K}$ and all $m \in \mathcal{M}$ it is required that $\text{vfy}_{\mathsf{M}}(k, m, \text{tag}(k, m)) = \mathtt{T}$.

We define a one-time multi-instance strong unforgeability notion SUF for MAC security – that is equivalent with standard strong unforgeability – for which the formal security game is depicted in the full version [2]. That is, for a game in which an adversary can generate instances i (with independent uniformly random keys $k_i \leftarrow_{\$} \mathcal{K}$) via an oracle Gen, the adversary can query a Tag oracle on a message m from message space $\mathcal{M}$ for each instance at most once to obtain the respective MAC tag. Additionally, the adversary can verify MAC tags for specified messages and instances via oracle Vfy and obtain an instance's key by querying an Expose oracle for this instance. The adversary wins by providing a forgery (m, τ) for an instance i to the Vfy oracle if there was no $\text{Tag}(i, m)$ query before with output τ and if i's key was not exposed via oracle Expose. We define the advantage of winning the SUF game against a MAC scheme M as $\text{Adv}^{\text{suf}}_{\mathsf{M}}(\mathcal{A}) = \Pr[\text{SUF}_{\mathsf{M}}(\mathcal{A}) \to 1]$.

3 Sufficient Security for Key-Updatable KEM

A key-updatable key encapsulation mechanism (kuKEM) is a key encapsulation mechanism that provides update algorithms for public key and secret key with respect to some associated data respectively. Prior to our work, this primitive was used to instantiate sesquidirectional RKE. In order to allow for our equivalence result, we minimally adjust the original kuKEM notion by Poettering and Rösler [17] and call it kuKEM*. The small, yet crucial changes comprise allowed updates of public and secret key during encapsulation and decapsulation (in our syntax definition) as well as the adversary's ability to manipulate utilized randomness of encapsulations (in our security definition). In Section 6 the rationales behind these changes are clarified. In order to provide a coherent definition, we not only describe alterations towards previous work but define kuKEM* entirely (as we consider our changes to be a significant contribution and believe that this strengthens comprehensibility).

Syntax A kuKEM* is a set of algorithms $\mathsf{K} = (\text{gen}_{\mathsf{K}}, \text{up}, \text{enc}, \text{dec})$ with sets of public keys $\mathcal{PK}$ and secret keys $\mathcal{SK}$, a set of associated data $\mathcal{AD}$ for updating the keys, a set of ciphertexts $\mathcal{C}$ (with $\mathcal{AD} \cap \mathcal{C} = \emptyset$), and a set of encapsulated keys $\mathcal{K}$. Furthermore we define $\mathcal{R}$ as the set of random coins used during the encapsulation:

$\text{gen}_{\mathsf{K}} \to_{\$} \mathcal{PK} \times \mathcal{SK}$
$\mathcal{PK} \times \mathcal{AD} \to \text{up} \to \mathcal{PK}$
$\mathcal{SK} \times \mathcal{AD} \to \text{up} \to \mathcal{SK}$
$\mathcal{PK} \times \mathcal{R} \to \text{enc} \to \mathcal{PK} \times \mathcal{K} \times \mathcal{C}$ or $\mathcal{PK} \to \text{enc} \to_{\$} \mathcal{PK} \times \mathcal{K} \times \mathcal{C}$
$\mathcal{SK} \times \mathcal{C} \to \text{dec} \to (\mathcal{SK} \times \mathcal{K}) \cup \{(\bot, \bot)\}$

Please note that the encapsulation and decapsulation may modify the public key and the secret key respectively – as a result, the kuKEM* is stateful (where the public key is a public state).[5]

Correctness The correctness for kuKEM* is (for simplicity) defined through game $\mathrm{CORR}_{\mathsf{K}}$ (see Figure 3), in which an adversary $\mathcal{A}$ can query encapsulation, decapsulation, and update oracles. The adversary (against correctness) wins if different keys are computed during decapsulation and the corresponding encapsulation even though compatible key updates were conducted and ciphertexts from encapsulations were directly forwarded to the decapsulation oracle.

Definition 1 (kuKEM* correctness). *A* kuKEM* *scheme* K *is correct if for every* $\mathcal{A}$*, the probability of winning game* $\mathrm{CORR}_{\mathsf{K}}$ *from Figure 3 is* $\Pr[\mathrm{CORR}_{\mathsf{K}}(\mathcal{A}) \to 1] = 0$.

Game $\mathrm{CORR}_{\mathsf{K}}(\mathcal{A})$
00 $(pk, sk) \leftarrow_{\$} \mathrm{gen}_{\mathsf{K}}$
01 $key[\cdot] \leftarrow \bot$
02 $trs \leftarrow \epsilon$; $trr \leftarrow \epsilon$
03 Invoke $\mathcal{A}$
04 Stop with 0

Oracle $\mathrm{Up}_S(ad)$
05 Require $ad \in \mathcal{AD}$
06 $pk \leftarrow \mathrm{up}(pk, ad)$
07 $trs \xleftarrow{\shortparallel} ad$
08 Return

Oracle $\mathrm{Up}_R(ad)$
09 Require $ad \in \mathcal{AD}$
10 $sk \leftarrow \mathrm{up}(sk, ad)$
11 $trr \xleftarrow{\shortparallel} ad$
12 Return

Oracle Enc()
13 $(pk, k, c) \leftarrow_{\$} \mathrm{enc}(pk)$
14 $trs \xleftarrow{\shortparallel} c$
15 $key[trs] \leftarrow k$
16 Return (pk, c)

Oracle Dec(c)
17 Require $c \in \mathcal{C}$
18 $(sk, k) \leftarrow \mathrm{dec}(sk, c)$
19 $trr \xleftarrow{\shortparallel} c$
20 If $trr \preceq trs$:
21 Reward $k \neq key[trr]$
22 Return

Fig. 3: The correctness notion of kuKEM* captured through game CORR.

Security Here we describe KUOWR security of kuKEM* as formally depicted in Figure 4. KUOWR defines one-way security of kuKEM* under randomness manipulation in a multi-instance/multi-challenge setting.

Intuitively, the KUOWR game requires that a secret key can only be used for decapsulation of a ciphertext if prior to this decapsulation all updates of this secret key and all decapsulations with this secret key were consistent with the updates of and encapsulations with the respective public key. This is reflected by using the transcript (of public key updates and encapsulations or secret key updates and decapsulations) as a reference to encapsulated "challenge keys" and secret keys.

[5] As kuKEM* naturally provides no security for encapsulated keys if the adversary can manipulate the randomness for $\mathrm{gen}_{\mathsf{K}}$ already, we only consider the manipulation of random coins for enc.

In order to let the adversary play with the kuKEM*'s algorithms, the game provides oracles Gen, Up_S, Up_R, Enc, and Dec. Thereby instances (i.e., key pairs) can be generated via oracle Gen and are referenced in the remaining oracles by a counter that refers to when the respective instance was generated.

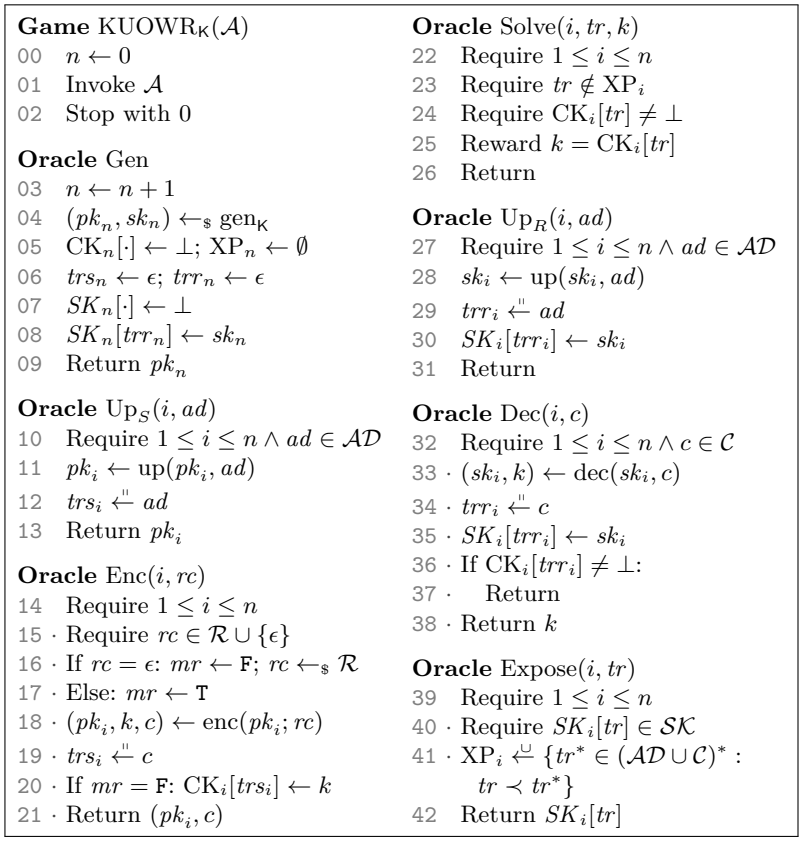

Fig. 4: Security experiment KUOWR, modeling one-way security of key-updatable KEM in a multi-instance/multi-challenge setting under randomness manipulation. Lines of code tagged with '·' are (substantially) modified with respect to KUOW security in [16]. Line 41 is a shortcut notion that can be implemented efficiently. CK: challenge keys, XP: exposed secret keys, trs, trr: transcripts.

For encapsulation via oracle Enc, the adversary can either choose the invocation's random coins by setting rc to some value that is not the empty string ϵ or let the encapsulation be called on fresh randomness by setting $rc = \epsilon$ (line 16). In the former case, the adversary trivially knows the encapsulated key. Thus, only when called with fresh randomness, the encapsulated key is marked as a challenge key in array CK (line 20).

The variables CK, SK, and XP (the latter two are explained below) are indexed via the transcript of operations on the respective key pair part. As public keys and secret keys can uniquely be referenced via the associated data under which they are updated and via ciphertexts that have been encapsulated or decapsulated by them, the concatenation of these values (i.e., sent or received transcripts trs, trr) are used as references to them in the KUOWR game.

On decapsulation of a key that is not marked as a challenge, the respective key is output to the adversary. Challenge keys are of course not provided to the adversary as thereby the challenge would be trivially solved (line 36).

Via oracle Expose, the adversary can obtain a secret key of specified instance i that results from an operation referenced by transcript tr. As described above, the transcript, to which a secret key refers, is built from the associated data of updates to this secret key (via oracle Up_R) and the ciphertexts of decapsulations with this secret key (via oracle Dec) as these two operations may modify the secret key. As all operations, performed with an exposed secret key, can be traced by the adversary (i.e., updates and decapsulations; note that both are deterministic), all secret keys that can be derived from an exposed secret key are also marked exposed via array XP (line 41).

Finally, an adversary can solve a challenge via oracle Solve by providing a guess for the challenge key that was encapsulated for an instance i with the encapsulation that is referenced by transcript tr. Recall that the transcript, to which an encapsulation refers, is built from the associated data of updates to the respective instance's public key (via oracle Up_S) and the ciphertexts of encapsulations with this instance's public key (via oracle Enc) as these two operations may modify the public key for encapsulation. If the secret key for decapsulating the referenced challenge key is not marked exposed (line 23) and the guess for the challenge key is correct (line 24), then game KUOWR stops with '1' (via 'Reward') meaning that the adversary wins.

Definition 2 (KUOWR **Advantage).** *The advantage of an adversary $\mathcal{A}$ against a* kuKEM* *scheme K in game* KUOWR *from Figure 4 is defined as* $\mathrm{Adv}^{\mathrm{kuowr}}_{\mathsf{K}}(\mathcal{A}) = \Pr[\mathrm{KUOWR}_{\mathsf{K}}(\mathcal{A}) \to 1]$.

We chose to consider one-way security as opposed to key indistinguishability for the kuKEM* as it suffices to show equivalence with key indistinguishability of RKE (in the ROM).

Differences compared to previous Security Definition In Figure 4 we denote changes from KUOW security (cf., Figure 1 [16]) by adding '$\cdot$' at the beginning of lines. Below we elaborate on these differences.

The main difference in our definition of KUOWR security compared to KUOW security is that we allow the adversary to manipulate the execution's random coins. As we define encapsulation and decapsulation to (potentially) update the used public key or secret key, another conceptual difference is that we only allow the adversary to encapsulate and decapsulate once under each public and secret key. Thus, we assume that public and secret keys are overwritten on

encapsulation and decapsulation respectively. In contrast to our security definition, in the KUOW security game only the current secret key of an instance can be exposed. Even though we assume the secret key to be replaced by its newer versions on updates or decapsulations, there might be, for example, backups that store older secret key versions. As a result we view the restriction of only allowing exposures of the current secret key artificial.[6] An important notational choice is that we index the variables with transcripts trs, trr instead of integer counters. This notation reflects the idea that public key and secret key only stay compatible as long as they are used correspondingly and immediately diverge on different associated data or tampered ciphertexts.

We further highlight the fundamental difference towards HkuPke by Jost et al. [12]. Their notion of HkuPke does not allow (fully adversary-controlled) associated data on public and secret key updates and additionally requires (authenticated) interaction between the holders of the key parts thereby. Looking ahead, this makes this primitive insufficient for diverging the public key from the secret key (in the states) of users A and B during an impersonation of A towards B in (U)RKE (especially under randomness manipulation). This is, however, required in an optimal security definition but explicitly excluded in the sub-optimal RKE notion by Jost et al. [12]. Since the syntax of HkuPke is already inadequate to reflect this security property, we cannot provide a separating attack. Nevertheless, we further expound this weakness in the full version [2].

Instantiation A kuKEM* scheme, secure in the KUOWR game, can be generically constructed from an OW-CCA adaptively secure hierarchical identity based key encapsulation mechanism (HIB-KEM). The construction – the same as in [16] – is provided for completeness in Figure 5. The update of public keys is the concatenation of associated data (interpreted as identities in the HIB-KEM) and the update of secret keys is the delegation to lower level secret keys in the identity hierarchy. The reduction is immediate: After guessing for which public key and after how many updates the challenge encapsulation that is solved by the adversary is queried, the challenge from the OW-CCA game is embedded into the respective KUOWR challenge.

Sufficiency of KUOWR *for SRKE* Before proving equivalence between security of key-updatable KEM and ratcheted key exchange, we shed a light on implications due to the differences between our notion of kuKEM* and its KUOWR security and the notion of kuKEM and its KUOW security in [16].

Remark 1. Even though the KUOWR game provides more power to the adversary in comparison to the KUOW game by allowing manipulation of random coins, exposures of past secret keys, and providing an explicit decapsulation oracle (instead of an oracle that only allows for checks of ciphertext-key pairs; cf.,

[6] It is important to note that the equivalence between KUOWR security of kuKEM* and KINDR security of URKE is independent of this definitional choice – if either both definitions allow or both definitions forbid the exposure of also past secret keys or states respectively, equivalence can be shown.

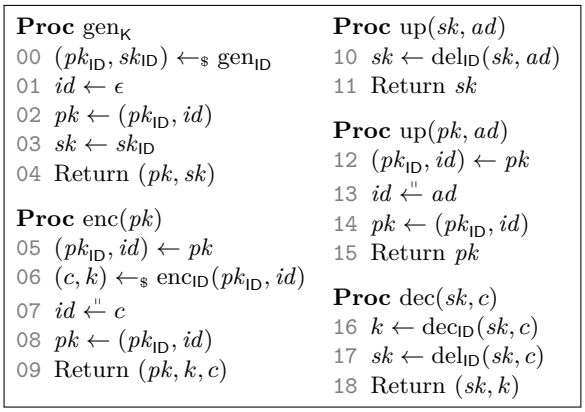

Fig. 5: Generic construction of a kuKEM* from a hierarchical identity based KEM $\mathsf{HK} = (\text{gen}_{\mathsf{ID}}, \text{del}_{\mathsf{ID}}, \text{enc}_{\mathsf{ID}}, \text{dec}_{\mathsf{ID}})$ (slightly differing from construction in [16] Figure 2 by adding an internal key update in encapsulation and decapsulation respectively).

Figure 1 [16]), the game also restricts the adversary's power by only allowing decapsulations under the current secret key of an instance (as opposed to also checking ciphertext-key pairs for past secret keys of an instance as in the KUOW game). One can exploit this and define protocols that are secure with respect to one game definition but allow for attacks in the other game. Consequently, neither of both security definitions implies the other one.

Despite the above described distinction between both security definitions, KUOWR security suffices to build sesquidirectional RKE according to the KIND definition in [17] – which was yet the weakest notion of security of RKE for which a construction was built from a key-updatable public key primitive. The ability to check ciphertext-key pairs under past versions of secret keys of an instance is actually never used in the proof of Poettering and Rösler [16]. The only case in which this Check oracle is used in their proof is B's receipt of a manipulated ciphertext from the adversary. Checking whether a ciphertext-key pair for the current version of a secret key of an instance is valid, can of course be conducted by using the Dec oracle of our KUOWR notion. For full details on their proof we refer the reader to Appendix C in [16].

Consequently, for the construction of KIND secure sesquidirectional RKE (according to [17] Figure 8) from Poettering and Rösler [17], the used kuKEM must either be KUOW secure (see [17] Figure 1) or KUOWR secure (see Figure 4), which is formally depicted in the following observation. Thus, even though these notions are not equivalent, they both suffice for constructing KIND secure sesquidirectional RKE.

Observation 1 *The sesquidirectional RKE protocol* R *from [17] Figure 6 offers* key indistinguishability *according to [17] Figure 8 if function* H *is modeled as a random oracle, the* kuKEM* *provides* KUOWR *security according to Figure 4,*

the one-time signature scheme provides SUF *security according to [16] Figure 22, the* MAC *scheme* M *provides* SUF *security according to Section 2.2, and the symmetric-key space of the* kuKEM* *is sufficiently large.*

4 Unidirectional RKE under Randomness Manipulation

Unidirectional RKE (URKE) is the simplest variant of ratcheted key exchange. After a common initialization of a session between two parties A and B, it enables the continuous establishment of keys within this session. In this unidirectional setting, A can initiate the computation of keys repeatedly. With each computation, a ciphertext is generated that is sent to B, who can then comprehend the computation and output (the same) key. Restricting A and B to this unidirectional communication setting, in which B cannot respond, allows to understand the basic principles of ratcheted key exchange. For the same reasons we provided the whole definition of kuKEM* before (i.e., we see our changes as a significant contribution and aim for a coherent depiction), we fully define URKE under randomness manipulation below.

Syntax We recall that URKE is a set of algorithms $\mathsf{UR} = (\text{init}, \text{snd}, \text{rcv})$ defined over sets of A's and B's states $\mathcal{S}_A$ and $\mathcal{S}_B$ respectively, a set of associated data $\mathcal{AD}$, a set of ciphertexts $\mathcal{C}$, and a set of keys $\mathcal{K}$ established between A and B. We extend the syntax of URKE by explicitly regarding the utilized randomness of the snd algorithm. Consequently we define $\mathcal{R}$ as the set of r̲andom c̲oins $rc \in \mathcal{R}$ used in snd. To highlight that A only sends and B only receives in URKE, we may add 'A' and 'B' as handles to the index of snd, and rcv respectively.

$$\text{init} \rightarrow_\$ \mathcal{S}_A \times \mathcal{S}_B$$
$$\mathcal{S}_A \times \mathcal{AD} \times \mathcal{R} \rightarrow \text{snd} \rightarrow \mathcal{S}_A \times \mathcal{K} \times \mathcal{C} \text{ or } \mathcal{S}_A \times \mathcal{AD} \rightarrow \text{snd} \rightarrow_\$ \mathcal{S}_A \times \mathcal{K} \times \mathcal{C}$$
$$\mathcal{S}_B \times \mathcal{AD} \times \mathcal{C} \rightarrow \text{rcv} \rightarrow \mathcal{S}_B \times \mathcal{K} \cup \{(\bot, \bot)\}$$

Please note that de-randomizing (or explicitly considering the randomness of) the initialization of URKE is of little value since an adversary, when controlling the random coins of init, obtains all information necessary to compute all keys between A and B.

Correctness Below we define correctness for URKE. Intuitively a URKE scheme is correct, if all keys produced with send operations of A can also be obtained with the resulting ciphertext by the respective receive operations of B.

Definition 3 (URKE **Correctness).** *Let* $\{ad_i \in \mathcal{AD}\}_{i\geq 1}$ *be a sequence of associated data. Let* $\{s_{A,i}\}_{i\geq 0}, \{s_{B,i}\}_{i\geq 0}$ *denote the sequences of A's and B's states generated by applying* $\text{snd}(\cdot, ad_i)$ *and* $\text{rcv}(\cdot, ad_i, \cdot)$ *operations iteratively for* $i \geq 1$*, that is,* $(s_{A,i}, k_i, c_i) \leftarrow_\$ \text{snd}(s_{A,i-1}, ad_i)$ *and* $(s_{B,i}, k'_i) \leftarrow \text{rcv}(s_{B,i-1}, ad_i, c_i)$*. We say* URKE *scheme* $\mathsf{UR} = (\text{init}, \text{snd}, \text{rcv})$ *is correct if for all* $s_{A,0}, s_{B,0} \leftarrow_\$ \text{init}$*, for all associated data sequences* $\{ad_i\}_{i\geq 1}$*, and for all random coins used for* snd *calls, the key sequences* $\{k_i\}_{i\geq 1}$ *and* $\{k'_i\}_{i\geq 1}$ *generated as above are equal.*

Security For security, we provide the KINDR game for defining key indistinguishability under randomness manipulation of URKE in Figure 6. In this game, the adversary can let the session participants A and B send and receive ciphertexts via SndA and RcvB oracle queries respectively to establish keys between them. By querying the Reveal or Challenge oracles, the adversary can obtain these established keys or receive a challenge key (that is either the real established key or a randomly sampled element from the key space) respectively. Finally, the adversary can expose A's and B's state as the output of a specified send or receive operation respectively via oracles ExposeA or ExposeB.

When querying the SndA oracle, the adversary can specify the random coins for the invocation of the snd algorithm from the set $\mathcal{R}$ or indicate that it wants the random coins to be sampled uniformly at random by letting $rc = \epsilon$. By allowing the adversary to set the randomness for the invocations of the snd algorithm and exposing past states (which was not permitted in the definition of Poettering and Rösler [17]), new trivial attacks arise.

Below we review and explain the trivial attacks of the original URKE KIND game, map them to our version, and then introduce new trivial attacks that arise due to randomness manipulation.

A conceptual difference between our game definition and the games by Poettering and Rösler [17] is the way variables (especially arrays) are indexed. While the KIND games of [17] make use of counters (of send and receive operations) to index computed keys and adversarial events, we use the communicated transcripts, sent and received by A and B respectively, as indices. We thereby heavily exploit the fact that synchronicity (and divergence) of the communication between A and B are defined over these transcripts, which results in a more comprehensible (but equivalent) game notation. Please note that, due to our indexing scheme, it suffices for our game definition to maintain a common key array $key[\cdot]$ and common sets of known keys KN and challenged keys CH for A and B (as opposed to arrays and sets for each party).[7]

The lines marked with '$\cdot$' in Figure 6 denote the handling of trivial attacks without randomness manipulation (as in [17]). Lines marked with '$\circ$' introduce modifications that become necessary due the new trivial attacks based on manipulation of randomness.

Trivial attacks without randomness manipulations are:

(a) If the adversary reveals a key via oracle Reveal, then challenging this key via oracle Challenge is trivial. In order to prevent reveal and challenge of the same key, sets KN and CH trace which keys have been revealed (line 23) and challenged (line 44). The adversary only wins, if the intersection of both sets is empty (line 08). Additionally, a key must only be challenged once as otherwise bit b can be obtained trivially (line 42).
Example: $c \leftarrow \text{SndA}(\epsilon, \epsilon)$; $k \leftarrow \text{Reveal}((\epsilon, c))$; Return $k = \text{Challenge}((\epsilon, c))$

[7] This is because a key, computed during the sending of A and the corresponding receiving of B, only differs between A and B if the received transcript of B diverged from the sent transcript of A.

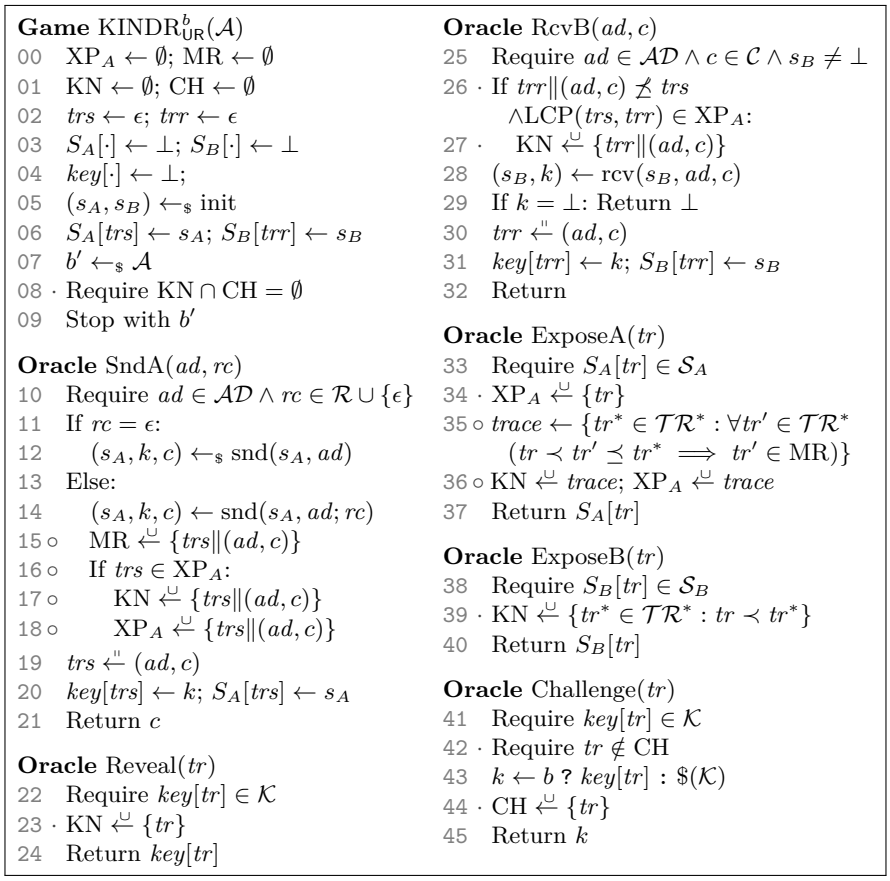

Fig. 6: Games KINDR^b, $b \in \{0, 1\}$, for URKE scheme UR. Lines of code tagged with a '·' denote mechanisms to prevent or detect trivial attacks without randomness manipulation; trivial attacks caused by randomness manipulation are detected and prevented by lines tagged with '○'. We define $\mathrm{LCP}(X, Y)$ to return the longest common prefix between X and Y, which are lists of atomic elements $z_i \in (\mathcal{AD} \times \mathcal{C})$. By longest common prefix we mean the longest list $Z = z_0 \| \ldots \| z_n$ for which $Z \preceq X \wedge Z \preceq Y$. We further define $\mathcal{TR} = \mathcal{AD} \times \mathcal{C}$. Line 39 is a shortcut notion that can be implemented efficiently. XP: exposed states, MR: states and keys affected by manipulated randomness, KN: known keys, CH: challenge keys, trs, trr: transcripts.

(b) As keys, that are computed by both parties (because ciphertexts between them have not been manipulated yet), are stored only once in array *key* (due to the indexing of arrays with transcripts instead of pure counters), the adversary cannot reveal these keys on one side of the communication (e.g., at A) and then challenge them on the other side (e.g., at B). Consequently, this trivial attack (which was explicitly considered in [17]) is implicitly handled by our game definition.

(c) After exposing B's state via oracle ExposeB, the adversary can comprehend all future computations of B. Consequently, all keys that can be received by B in the future are marked known (line 39).
Example: $s_B \leftarrow \text{ExposeB}(\epsilon)$; $c \leftarrow \text{SndA}(\epsilon, \epsilon)$; $\text{RcvB}(\epsilon, c)$; $(s_B, k) \leftarrow \text{rcv}(s_B, \epsilon, c)$; Return $k = \text{Challenge}((\epsilon, c))$

(d) Exposing B's state, as long as the communication between A and B has not yet been manipulated by the adversary, allows the adversary also to compute all future keys established by A (which is also implicitly handled by our indexing of arrays via transcripts).

(e) Exposing A's state via oracle ExposeA allows the adversary to impersonate A towards B by using the exposed state to create and send own valid ciphertexts to B. As creating a forged ciphertext reveals the key that is computed by B on receipt, such keys are marked known (lines 26-27). The detection of this trivial attack works as follows: As soon as B receives a ciphertext that was not sent by A (i.e., B's transcript together with the received ciphertext is not a prefix of A's transcript) and A was exposed after A sent the last ciphertext that was also received by B (i.e., after the last common prefix LCP), the adversary is able to create this ciphertext validly on its own.[8]
Example: $s_A \leftarrow \text{ExposeA}$; $(s_A, k, c) \leftarrow \text{snd}(s_A, \epsilon)$; $\text{RcvB}(\epsilon, c)$; Return $k = \text{Challenge}((\epsilon, c))$

Due to randomness manipulations, the adversary can additionally conduct the following attacks trivially:

(f) If the randomness for sending is set by the adversary (via $\text{SndA}(ad, rc)$, $rc \neq \epsilon$) and the state, used for this sending, is exposed (via ExposeA), then also the next state of A, output by this send operation, will be known (and marked as exposed) as sending is thereby deterministically computed on inputs that are known by the adversary (lines 16,18). Since the adversary can also retrospectively expose A's state, all computations that can be traced, due to continuous manipulated randomness of subsequent SndA oracle queries (unified in set MR) after such an exposure, are also marked as exposed (lines 35-36).
Example: $rc \leftarrow_{\$} \mathcal{R}$; $c' \leftarrow \text{SndA}(\epsilon, rc)$; $\text{RcvB}(\epsilon, c')$; $s_A \leftarrow \text{ExposeA}(\epsilon)$; $(s_A, k', c') \leftarrow \text{snd}(s_A, \epsilon; rc)$; $(s_A, k, c) \leftarrow_{\$} \text{snd}(s_A, \epsilon)$; $\text{RcvB}(\epsilon, c)$; Return $k = \text{Challenge}((\epsilon, c') \| (\epsilon, c))$

(g) Similarly, if the randomness for sending is set by the adversary and the state that A uses during this send operation is exposed, then the key, computed during sending, is known by the adversary since its computation is thereby

[8] Please note that we need to detect this trivial attack this way (in contrast to the game in [17]) because the adversary can forge ciphertexts to B without letting the communication between A and B actually diverge. It can do so by creating an own valid ciphertext which it sends to B (via $s_A \leftarrow \text{ExposeA}(\epsilon)$; $rc \leftarrow_{\$} \mathcal{R}$; $(s_A, k, c) \leftarrow \text{snd}(s_A, \epsilon; rc)$; $\text{RcvB}(\epsilon, c)$) but then it lets A compute the same ciphertext (via $\text{SndA}(\epsilon, rc)$). As a result, A and B are still in sync.

deterministic (lines 16-17,35-36).
Example: $rc \leftarrow_{\$} \mathcal{R}$; $c \leftarrow \text{SndA}(\epsilon, rc)$; $s_A \leftarrow \text{ExposeA}(\epsilon)$; $(s_A, k, c) \leftarrow \text{snd}(s_A, \epsilon; rc)$; Return $k = \text{Challenge}((\epsilon, c))$

Based on this game, we define the advantage of an adversary in breaking the security of an URKE scheme as follows.

Definition 4 (KINDR **Advantage).** *The advantage of an adversary $\mathcal{A}$ against a* URKE *scheme* UR *in game* KINDR *from Figure 6 is defined as* $\text{Adv}^{\text{kindr}}_{\mathsf{UR}}(\mathcal{A}) = \left|\Pr[\text{KINDR}^0_{\mathsf{UR}}(\mathcal{A}) = 1] - \Pr[\text{KINDR}^1_{\mathsf{UR}}(\mathcal{A}) = 1]\right|$.

We say that an URKE scheme UR is secure if the advantage is negligible for all probabilistic polynomial time adversaries $\mathcal{A}$.

Please note that KINDR security of URKE is strictly stronger than both KIND security notions of URKE, defined by Bellare et al. [3] and Poettering and Rösler [17] (which themselves are incomparable among each other).

5 kuKEM* to URKE

Since our ultimate goal is to show that existence of a kuKEM* primitive is a necessary and sufficient condition to construct a URKE primitive – albeit requiring the help of other common cryptographic primitives such as hash functions (modeled as random oracle) and message authentication codes –, we dedicate this section to proving the latter of these implications.

Construction of URKE *from* kuKEM* We give a generic way to construct a URKE scheme UR from a kuKEM* scheme K with the help of random oracle H and MAC scheme M. This transformation $\mathsf{K} \rightarrow \mathsf{UR}$ is fully depicted in Figure 7. Below we briefly describe the algorithms of URKE scheme $\mathsf{UR} = (\text{init}, \text{snd}, \text{rcv})$.

During the state initiation algorithm init, a kuKEM* key pair (sk, pk) is generated such that the encapsulation key pk is embedded into the sender state s_A, and the decapsulation key sk into the receiver state s_B. The remaining state variables are exactly same for A and B. More specifically, two further keys are generated during initialization: the symmetric state key K and a MAC key $k.m$. Furthermore the sent or received transcript (initialized with an empty string ϵ) is stored in each state. For brevity, we assume that K, $k.m$, and the update key $k.u$ (used during sending and receiving; see below) all belong to the same key domain $\mathcal{K}$ that is sufficiently large.

On sending, public key pk in A's state is used by the encapsulation algorithm to generate key k and ciphertext c. Then, MAC key $k.m$, contained in the current state of A, is used to issue a tag τ over the tuple of associated data ad and encapsulation ciphertext c. The finally sent ciphertext, denoted by C, is a concatenation of c and τ. The output key $k.o$, as well as the symmetric keys of the next state of A are obtained from the random oracle, on input of the symmetric state key K, the freshly encapsulated key k, and the history of sent transcript t. Finally, a kuKEM* update is applied on pk under associated data

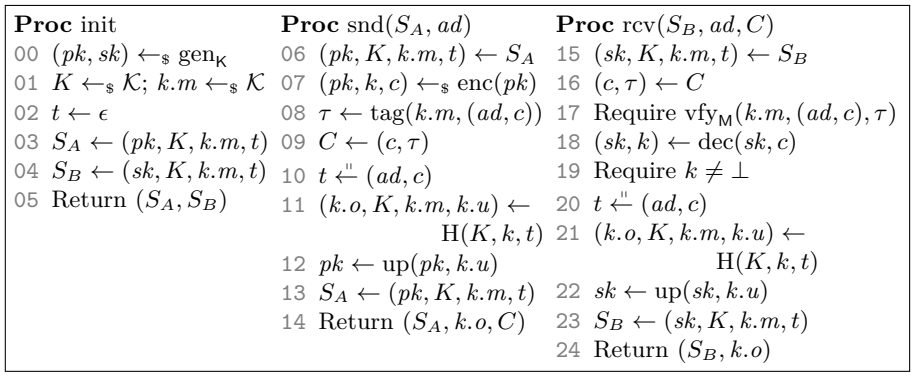

Fig. 7: Construction of a URKE scheme from a kuKEM* scheme $\mathsf{K} = (\text{gen}_\mathsf{K}, \text{up}, \text{enc}, \text{dec})$, a message authentication code $\mathsf{M} = (\text{tag}, \text{vfy}_\mathsf{M})$, and a random oracle H. For simplicity we denote the key space of the MAC and the space of the symmetric key K in s_A with the same symbol $\mathcal{K}$.

that is derived from the random oracle output (denoted by $k.u$). Please note that the encapsulation algorithm is the only randomized operation inside snd. Hence the random coins of the latter are only used by the encapsulation.

On receiving, the operations are on par with the sending algorithm. Namely, the received ciphertext C is parsed as the encapsulation ciphertext c and the MAC tag τ. The latter is verified with regards to the MAC key $k.m$, stored in the state of B. After the key k is decapsulated, the same input to the random oracle H is composed. The symmetric components of the next state and $k.o$ are derived from the random oracle's output. Finally, the secret key sk is updated with $k.u$, so that it is in-sync with the update of pk.

We remark that our construction in Figure 7 differs from the unidirectional RKE scheme by Poettering and Rösler [17] only in the output of the random oracle and in the subsequent use of the kuKEM*'s update algorithm (instead they freshly generated a new KEM key pair from the random oracle output). These changes are, nevertheless, significant as their scheme is insecure when the adversary is able to (reveal or) manipulate the random coins for invocations of the snd algorithm. We give a detailed attack description against their scheme in our model in the full version [2].

Theorem 1. *If* kuKEM* *scheme* K *is* KUOWR *secure according to Figure 4,* MAC *scheme* M *is* SUF *secure according to Section 2.2, and* H *is a hash function modeled as random oracle, then* URKE *scheme* UR *from Figure 7 is* KINDR *secure according to Figure 6 with*

$$\text{Adv}^{\text{kindr}}_{\mathsf{UR}}(\mathcal{A}) \leq \text{Adv}^{\text{kuowr}}_{\mathsf{K}}(\mathcal{B}_\mathsf{K}) + \text{Adv}^{\text{suf}}_{\mathsf{M}}(\mathcal{B}_\mathsf{M}) + \frac{q_\text{H} \cdot (q_\text{SndA} + q_\text{RcvB})}{|\mathcal{K}|}$$

where $\mathcal{A}$ *is an adversary against* KINDR *security,* $\mathcal{B}_\mathsf{K}$ *is an adversary against* KUOWR *security,* $\mathcal{B}_\mathsf{M}$ *is an adversary against* SUF *security,* $\mathcal{K}$ *is the key domain*

in the construction UR*,* q_{SndA}*,* q_{RcvB}*, and* q_{H} *are the number of* SndA, RcvB *and* H *queries respectively by* $\mathcal{A}$*, and the running time of* $\mathcal{A}$ *is approximately the running time of* $\mathcal{B}_{\mathsf{K}}$ *and* $\mathcal{B}_{\mathsf{M}}$*.*

Proof (Sketch, Theorem 1). We here give the sketch of the full proof that is in the full version [2]. Our idea is to design a series of games **Game 0-5**, in which differences between subsequent games are only syntactical and the advantage of the adversary $\mathcal{A}$ remains same. From this fifth game we are then ultimately able to reduce either of the following cases, that are explained below, to one of the hardness assumptions.

Consider the following scenarios which lead to a win for the adversary $\mathcal{A}$. Since the challenged keys are derived from the random oracle, we argue that, if $\mathcal{A}$ does not make a random oracle query $\mathrm{H}(K, k, t)$ for any of the challenged keys, then its advantage in guessing the challenge bit correctly remains negligible. We do not consider random oracle queries to keys that are trivially revealed to the adversary, as they do not lead to a win in the KINDR game (e.g., if the exposed state of B helps the adversary to trivially query H). Therefore, we regard the following three events in which $\mathcal{A}$ makes such *special* random oracle queries:

- The random oracle query $\mathrm{H}(K, k, t)$ belongs to one of the keys derived by the sender, in which fresh random coins, unknown to the adversary, are used for sending (and hence for encapsulation). In this case, we can give a reduction to the KUOWR game with respect to kuKEM* scheme K, in which the reduction wins the KUOWR game by using the encapsulated key k as the solution.
- The random oracle query $\mathrm{H}(K, k, t)$ belongs to one of the keys, derived from the sender where the used random coins are chosen by the adversary. We know that $\mathcal{A}$ did not expose the respectively used states of A or B as this leads to a trivial win. Therefore, we can show that the symmetric state keys K in these cases are independent from the view of $\mathcal{A}$. This implies that making such special $\mathrm{H}(K, k, t)$ query requires a collision in the key domain $\mathcal{K}$, whose probability is bounded by $q_{\mathrm{H}} \cdot (q_{\mathrm{SndA}} + q_{\mathrm{RcvB}})/|\mathcal{K}|$.
- The random oracle query $\mathrm{H}(K, k, t)$ belongs to one of the keys, derived by the receiver B, who reaches to an out-of-sync status (if B is still in-sync with A, then one of the two cases above are relevant). Since each received ciphertext contains a MAC tag, we can show that the first received ciphertext by B that is different from the sent ciphertext by A either corresponds to a trivial impersonation or can be used to reduce this event to a forgery in the SUF game with respect to MAC scheme M.

Therefore, by bounding the probability of these three cases, we can deduce the adversary's advantage (which is negligible under the named assumptions). □

6 URKE to kuKEM*

In order to show that public key encryption with independently updatable key pairs (in our case kuKEM*) is a necessary building block for ratcheted key

exchange, we build the former from the latter. The major obstacle is that the updates of public key and secret key of a kuKEM* are conducted independently – consequently no communication between holder of the public key and holder of the secret key can be assumed for updates. In contrast, all actions in ratcheted key exchange are based on communication (i.e., sent or received ciphertexts). Another property that public key updates for kuKEM* must fulfill – in contrast to state updates in KIND secure unidirectional RKE as in [17] – is that they must not leak any information on the according secret key during the update computation. In the following we describe how we solve these two issues and present a reduction of KUOWR security to KINDR security of a generic URKE scheme.

Construction of kuKEM* *from* URKE The weaker KIND security of URKE (as in [17]) already allows that the sender's state s_A can always be exposed without affecting the security of any established keys (as long as this exposed state is not used to impersonate A towards B). Consequently, A's pure state reveals no information on encapsulated keys nor on B's secret key(s). KIND security of URKE further implies that B's state only reveals information on keys that have not yet been computed by B (while earlier computed keys stay secure). One can imagine A's state consequently as the public part of a (stateful) key pair and B's state as the secret counterpart.

The two above mentioned crucial properties of KUOW(R) security are, however, not implied by KIND security when using s_A as the public key and s_B as the secret key of a kuKEM. Firstly, updating s_B (as part of receiving a ciphertext) requires that the ciphertext, generated during sending of A (and updating of s_A), is known by B but the syntax of kuKEM does not allow an interaction between public key holder and secret key holder. This issue can be solved by de-randomizing the snd algorithm. If A's state as part of the public key is updated via a de-randomized invocation of snd, the secret key holder can also obtain the ciphertext that A would produce for the same update (by invoking the de-randomized/deterministic snd) and then update s_B with this ciphertext via rcv. A conceptional depiction of this is in Figure 8. Thereby the secret key is defined to contain s_A in addition to s_B.

Secondly, in the URKE construction of Poettering and Rösler [17] A temporarily computes secrets of B that match A's updated values during sending. As a result, normal KIND security allows that a de-randomized snd invocation reveals the secrets of B to an adversary if s_A is known (see the full version [2] for a detailed description of this attack). In order to solve this issue, the security definition of URKE must ensure that future encapsulated keys' security is not compromised if snd is invoked under a known state s_A and with random coins that are chosen by an adversary (i.e., KINDR security).

Our generic construction of a KUOWR secure kuKEM* from a generic KINDR secure URKE scheme is depicted in Figure 9. As described before, the public key contains state s_A and the secret key contains both states (s_A, s_B) that are derived from the init algorithm. In order to update the public key, the snd algorithm is invoked on state s_A, with the update associated data, and fixed randomness.

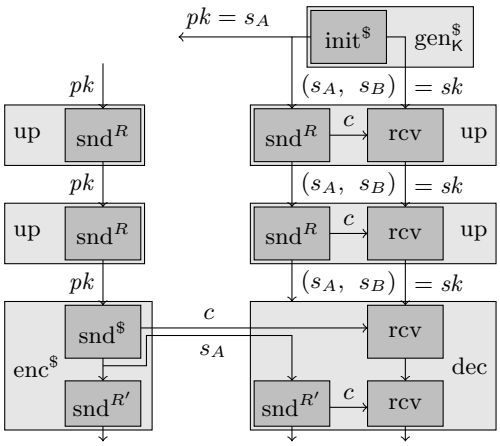

Fig. 8: Conceptual depiction of kuKEM* construction from generic URKE scheme. The symbol in the upper index of an algorithm name denotes the source of random coins ('$' indicates uniformly sampled). R is a fixed value. For clarity we omit *ad* inputs and k outputs (cf. Figure 1).

The output key and ciphertext are thereby ignored. Accordingly, the secret key is updated by first invoking the snd algorithm on state s_A with the same fixed randomness and the update associated data. This time the respective ciphertext from A to B is not omitted but used as input to rcv algorithm with the same associated data under s_B.

Encapsulation and decapsulation are conducted by invoking snd probabilistically and rcv respectively. In order to separate updates from en-/decapsulation, a '0' or '1' is prepended to the associated data input of snd and rcv respectively. For bounding the probability of a ciphertext collision in the proof, a randomly sampled 'collision key' *ck* is attached to the associated data of the snd invocation in encapsulation. In order to accordingly add *ck* to the associated data of rcv as part of the decapsulation, *ck* is appended to the ciphertext. Since state s_A, output by the snd algorithm during the encapsulation, is computed probabilistically, it is also attached to the encapsulation ciphertext, so that (the other) s_A, embedded in the secret key, can be kept compatible with the public key holder's. To bind *ck* and s_A to the ciphertext, both are integrity protected by a message authentication code (MAC) that takes one part of the key from the snd invocation as MAC key (only the remaining key bytes are output as the encapsulated kuKEM* key). Additionally the whole ciphertext (i.e., URKE ciphertext, collision key, state s_A, and MAC tag) is used as associated data for an additional 'internal update' of the public key and the secret key in encapsulation and decapsulation respectively. This is done to escalate manipulations of collision key, state s_A, or MAC tag (as part of the ciphertext) back into the URKE states s_A and s_B (as part of public key and secret key). For full details on the rationales behind these two *binding* steps we refer the reader to the proof.

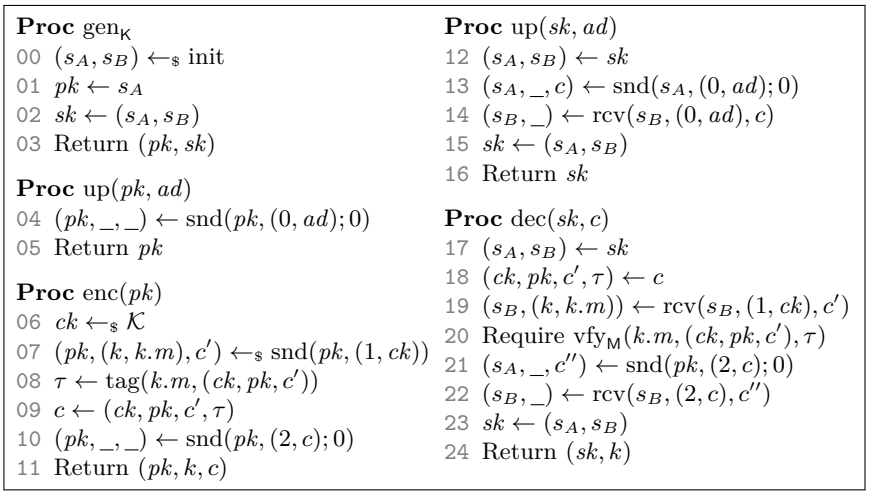

Fig. 9: Construction of a key-updatable KEM from a generic URKE scheme $\mathsf{UR} = (\mathrm{init}, \mathrm{snd}, \mathrm{rcv})$ and one-time message authentication code $\mathsf{M} = (\mathrm{tag}, \mathrm{vfy}_{\mathsf{M}})$.

Interestingly, the public key holder can postpone the de-randomized snd invocation for public key updates until encapsulation and instead only remember the updates' associated data without compromising security. However, the updates of the secret key must be performed immediately as otherwise an exposure of the current secret key reveals also information on its past versions. Thereby the computation of snd in the up algorithm must be conducted during the secret key update without interaction between public key holder and secret key holder.

Theorem 2. *If* URKE *scheme* UR *is* KINDR *secure according to Figure 6, one-time* MAC M *is* SUF *secure according to Section 2.2, and for all* $(k, k.m) \in \mathcal{K}_{\mathsf{UR}}$ *it holds that* $k \in \mathcal{K}_{\mathsf{K}}$ *and* $k.m \in \mathcal{K}_{\mathsf{M}}$*, then* kuKEM* *scheme* K *from Figure 9 is* KUOWR *secure according to Figure 4 with*

$$\mathrm{Adv}^{\mathrm{kuowr}}_{\mathsf{K}}(\mathcal{A}) \leq q_{\mathrm{Gen}} q_{\mathrm{Enc}} \cdot \left(\mathrm{Adv}^{\mathrm{kindr}}_{\mathsf{UR}}(\mathcal{B}_{\mathsf{UR}}) + \mathrm{Adv}^{\mathrm{suf}}_{\mathsf{M}}(\mathcal{B}_{\mathsf{M}}) + \frac{1}{|\mathcal{K}|} \right),$$
$$\textit{with } \mathrm{Adv}^{\mathrm{suf}}_{\mathsf{M}}(\mathcal{B}_{\mathsf{M}}) \leq \mathrm{Adv}^{\mathrm{kindr}}_{\mathsf{UR}}(\mathcal{B}_{\mathsf{UR}})$$

where $\mathcal{A}$ *is an adversary against* KUOWR *security,* $\mathcal{B}_{\mathsf{UR}}$ *is an adversary against* KINDR *security of* UR*,* $\mathcal{B}_{\mathsf{M}}$ *is an adversary against* SUF *security of* M*,* q_{Gen} *and* q_{Enc} *are the number of* Gen *and* Enc *queries by* $\mathcal{A}$ *respectively,* $\mathcal{K}$ *is the space from which ck is sampled, and the running time of* $\mathcal{A}$ *is approximately the running time of* $\mathcal{B}_{\mathsf{UR}}$ *and* $\mathcal{B}_{\mathsf{M}}$*.*

In the full version [2] we show how to construct an SUF secure one-time MAC from a generic KINDR secure URKE scheme, which implies the second term in Theorem 2. We prove Theorem 2 below and provide a formal pseudo-code version of the simulation's game hops in the full version [2].

Proof (Theorem 2).

We conduct the proof in four game hops: In the first game hop we guess for which instance the first valid Solve oracle query is provided by the adversary; in the second game hop, we guess for which Enc oracle query of the previously guessed instance the first valid Solve oracle query is provided; additionally the simulation aborts in this game hop if the adversary crafts this first valid ciphertext and provides it to the Dec oracle before it is output by the Enc oracle; in the third game hop, we replace the key, output by the first snd invocation in this guessed Enc oracle query by a randomly sampled key (which is reduced to KINDR security of UR); in the final game hop, we abort on a MAC forgery, provided to the Dec oracle, that belongs to the ciphertext that is output by the guessed Enc oracle query (which is reduced to the SUF security of M).

Game 0 This game is equivalent to the original KUOWR game.

Game 1 The simulation guesses for which instance n_{Gen} the first key k^* is provided to the Solve oracle such that the secret key for decapsulation is not marked exposed (i.e., $tr^* \notin \text{XP}_{n_{\text{Gen}}}$) and the provided key equals the indicated challenge key (i.e., $k^* = \text{CK}_{n_{\text{Gen}}}[tr^*]$). Therefore n_{Gen} is randomly sampled from $[q_{\text{Gen}}]$, where q_{Gen} is the number of Gen oracle queries by the adversary. The reduction aborts if n_{Gen} is not the instance for which the first valid Solve oracle query is provided.

Consequently we have $\text{Adv}^{G_0} = q_{\text{Gen}} \cdot \text{Adv}^{G_1}$.

Game 2 The simulation guesses in which of n_{Gen}'s Enc queries the challenge is created, that is the first valid query to the Solve oracle by the adversary. Therefore n_{Enc} is randomly sampled from $[q_{\text{Enc}}]$ and the simulation aborts if either the randomness for the n_{Enc}'s Enc query is manipulated as thereby no challenge would be created, or the first valid query to the Solve oracle is for another challenge than the one created by n_{Gen}'s n_{Enc}th Enc query, or a secret key that helps to trivially solve the challenge from n_{Gen}'s n_{Enc}th Enc query is exposed.

In addition, the simulation aborts if, before the n_{Gen}'s n_{Enc}th Enc query was made, Dec was queried on a ciphertext (with the same preceding transcript) that contains the same URKE ciphertext and 'collision key' ck as n_{Gen}'s n_{Enc}th Enc query. As the probability of a collision in the URKE transcript (i.e., associated data and ciphertext of the first snd invocation of n_{Gen}'s n_{Enc}th Enc query were previously already provided to n_{Gen}'s n_{Enc}th Dec query under the same preceding transcript) is bounded by a collision in the the key space $\mathcal{K}$ (as thereby ck as associated data must collide), we have $\text{Adv}^{G_1} = q_{\text{Enc}} \cdot \left(\text{Adv}^{G_2} + \frac{1}{|\mathcal{K}|}\right)$.

Game 3 The simulation replaces the output $(k, k.m)$ from the first snd invocation of n_{Gen}'s n_{Enc}th Enc query by values randomly sampled.

An adversary that can distinguish between **Game 2** and **Game 3** can be turned into an adversary that breaks KINDR security of URKE scheme UR. We describe the reduction below: The reduction obtains n_{Gen}'s public key in oracle Gen via oracle ExposeA from the KINDR game. Invocations of snd in Up_S to n_{Gen} are replaced by SndA and ExposeA queries. Invocations of snd in Up_R to n_{Gen} are processed by the reduction itself and the subsequent rcv invocations

are replaced by RcvB queries. The state s_B in queries to Expose for n_{Gen} is obtained via ExposeB queries to the KINDR game. For all queries to Enc of n_{Gen} the snd invocations are replaced by SndA and ExposeA queries. kuKEM* key and MAC key $(k, k.m)$ for n_{Gen}'s Enc oracle queries are obtained via Reveal – except for n_{Gen}'s n_{Enc}th Enc query, in which these two keys are obtained from the Challenge oracle in the KINDR game. Invocations of rcv in the Dec oracle for n_{Gen} are replaced by RcvB queries and Reveal queries (in case the respective key was not already computed in the Enc oracle). The snd invocation in oracle Dec is directly computed by the reduction.

In order to show that manipulations of transcripts in the KUOWR game manipulate equivalently the transcripts in the KINDR game (such that the state s_A in the public key diverges from state s_B in the secret key iff the transcripts $trs_{n_{\text{Gen}}}$ and $trr_{n_{\text{Gen}}}$ diverge), we define the translation array TR[·] that maps the transcript of n_{Gen} in the KUOWR game to the according transcripts in the KINDR game (see the pseudo-code game-hop in the full version [2]).

As **Game 2** aborts if n_{Gen}'s n_{Enc}th Enc query entails no valid KINDR challenge, or if the respective ciphertext was already crafted by the adversary (and provided to the Dec oracle), an adversary, distinguishing the real key pair $(k, k.m)$ from the randomly sampled one, breaks KINDR security. Formally, the solution for n_{Gen}'s n_{Enc}th Enc query to the Solve oracle is compared with the challenge key k from the KINDR Challenge oracle (which is obtained during n_{Gen}'s n_{Enc}th Enc query): If the keys equal, the reduction terminates with $b' = 0$ (as thereby the KINDR game's challenge entailed the real key), otherwise it terminates with $b' = 1$.

Consequently we have $\text{Adv}^{G_2} \leq \text{Adv}^{G_3} + \text{Adv}^{\text{kindr}}_{\mathsf{UR}}(\mathcal{B}_{\mathsf{UR}})$.

Game 4 The only way, the adversary can win in **Game 3**, is to keep secret key and public key of n_{Gen} compatible (by updating them equivalently and forwarding all Enc queries to the Dec oracle) and then forwarding only the URKE ciphertext c' of n_{Gen}'s n_{Enc}th Enc query to the Dec oracle while manipulating parts of the remaining challenge ciphertext. Thereby the Dec oracle outputs the correct challenge key such that the adversary trivially wins.[9]

We therefore define **Game 4** to let the simulation abort if a forgery of the MAC tag for the challenge ciphertext is provided to the Dec oracle. Distinguishing between **Game 3** and **Game 4** can hence be reduced to the SUF security of the one-time MAC M. We describe the reduction below: Instead of sampling $k.m$ randomly, the MAC tag for n_{Gen}'s n_{Enc}th Enc query is derived from the Tag oracle of the SUF game. Since an abort requires that the URKE challenge ciphertext c' is indeed received in oracle Dec (and also the transcripts prior to this ciphertext equal for $trs_{n_{\text{Gen}}}$ and $trr_{n_{\text{Gen}}}$), the URKE key (containing $k.m$)

[9] Please note that after this manipulation, the states s_A and s_B in the public key and secret key respectively diverge, but the key, output by the Dec oracle, still equals the challenge key. In case, the URKE ciphertext c' from the challenge ciphertext is already provided manipulately to the Dec oracle, the challenge key is already independent from the key, computed in the Dec oracle.

equals. As a consequence, a crafted ciphertext (pk, c', τ), provided to the Dec oracle, is a forgery τ for message (pk, c') in the SUF game.

Consequently we have $\mathrm{Adv}^{G_3} \leq \mathrm{Adv}^{G_4} + \mathrm{Adv}^{\mathrm{suf}}_{\mathsf{M}}(\mathcal{B}_{\mathsf{M}})$.

As the challenge key from n_{Gen}'s n_{Enc}th Enc query is randomly sampled and cannot be derived from any other oracle, the advantage of winning in **Game 4** is $\mathrm{Adv}^{G_4} = 0$.

Summing up the advantages above, we have:

$$\mathrm{Adv}^{\mathrm{kuowr}}_{\mathsf{K}}(\mathcal{A}) \leq q_{\mathrm{Gen}} q_{\mathrm{Enc}} \cdot \left(\mathrm{Adv}^{\mathrm{kindr}}_{\mathsf{UR}}(\mathcal{B}_{\mathsf{UR}}) + \mathrm{Adv}^{\mathrm{suf}}_{\mathsf{M}}(\mathcal{B}_{\mathsf{M}}) + \frac{1}{|\mathcal{K}|} \right)$$

where $\mathrm{Adv}^{\mathrm{suf}}_{\mathsf{M}}(\mathcal{B}_{\mathsf{M}}) \leq \mathrm{Adv}^{\mathrm{kindr}}_{\mathsf{UR}}(\mathcal{B}_{\mathsf{UR}})$ follows from an SUF secure one-time MAC construction from a generic KINDR secure URKE scheme UR (which is described in the full version [2]). □

7 Discussion

Our results clearly show that key-updatable key encapsulation is a necessary building block for optimally secure ratcheted key exchange, if the security definition of the latter regards manipulation of the algorithm invocations' random coins. As unidirectional RKE can naturally be built from sesquidirectional RKE, which in turn can be built from bidirectional RKE (which can be derived from optimally secure group RKE), our results are expected to hold also for the according security definitions under these extended communication settings. In contrast, security definitions of ratcheting that restrict the adversary more than necessary in exposing the local state or in solving embedded game challenges (i.e., by excluding more than unpreventable attacks) allow for instantiations that can dispense with these inefficient building blocks.

However, the two previous security definitions fulfilled by constructions that use kuKEM as a building block (cf. Table 1) consider only *randomness reveal* [10] or even *secure randomness* [17]. This raises the question whether using kuKEM in these cases was indeed necessary (or not). The resulting gap between the notions of ratcheting that can be built from only standard PKC and our optimally secure URKE definition with *randomness manipulation*, implying kuKEM, will be discussed in the following.

Implications under Randomness Reveal The core of our proof (showing that URKE implies kuKEM under randomness manipulation) is to utilize URKE's state update in algorithms snd and rcv for realizing public key and secret key updates in kuKEM's up algorithm. In order to remove the otherwise necessary communication between snd and rcv algorithms of RKE, snd is de-randomized by fixing its random coins to a static value. While this de-randomization *trick* is not immediately possible if the reduction to URKE KIND security cannot manipulate the randomness of snd invocations, one can utilize a programmable random oracle to emulate it: instead of fixing the (input) random coins of snd invocations to a static value, one could derive these coins from the output of a

random oracle on input of the respective update's associated data (i.e., *ad* input of algorithm up). Additionally, instead of directly forwarding the update's associated data to the associated data input of snd, another random oracle could be interposed between them. The reduction then simply pre-computes all kuKEM up invocations independent of associated data inputs by querying the SndA oracle in the URKE KIND game on random associated data strings. Then the reduction reveals all used random coins in the URKE KIND game and programs them as output into the random oracle lazily (i.e., as soon as the adversary queries the random oracle on update associated data strings). By correctly guessing, which of the adversary's random oracle queries fit its queried kuKEM update invocations, the reduction can perform the same de-randomization trick as in our proof. The probability of guessing correctly is, however, exponential in the number of queried kuKEM updates such that a useful implication may only be derivable for a constant number of queried updates.

In conclusion, we conjecture that URKE under randomness reveal already requires the use of a kuKEM-like building block with a constantly bounded number of public key and secret key updates. Thereby we argue that our proof approach partially carries over to the case of randomness reveal. This would indicate that the use of a kuKEM-like building block in the construction of Jaeger and Stepanovs [10] is indeed necessary. The formal analysis of this conjecture is left as an open question for future work.

Implications under Secure Randomness For optimal security under secure randomness, Poettering and Rösler [17] show that URKE can be instantiated from standard PKC only (cf. Table 1). In contrast, their construction for sesquidirectional RKE (SRKE: a restricted interactive RKE variant) uses kuKEM for satisfying optimal security under secure randomness. Since a reduction towards SRKE (under KIND security with secure randomness) has no access to random coins respectively used in the RKE algorithms, our de-randomization trick seems inapplicable. Furthermore, while the RKE algorithms snd and rcv can use exchanged ciphertexts for their state updates, generically transforming this state update to realize a 'silent', non-interactive key update needed for kuKEM without our trick appears (at least) problematic.

Nevertheless, it is likely that SRKE KIND security under secure randomness requires kuKEM-like building blocks. This intuition is based on an example attack by Poettering and Rösler [16, Appendix B.2]. It illustrates that a key k^*, computed by any secure SRKE construction under the following attack, needs to be indistinguishable from a random key according to this security notion. The attack proceeds as follows: 1. Alice's and Bob's states are exposed ($s_A \leftarrow \mathrm{ExposeA}(\epsilon)$; $s_B \leftarrow \mathrm{ExposeB}(\epsilon)$), 2. Bob sends update information to Alice (which is possible in SRKE) to recover from his exposure ($c \leftarrow \mathrm{SndB}(\epsilon, \epsilon)$; $\mathrm{RcvA}(\epsilon, c)$). Keys established by Alice after receiving the update information are required to be secure again. Translated to the kuKEM setting, this step corresponds to Bob generating a new key pair and publishing the respective public key. 3. Simultaneously Alice is impersonated towards Bob ($(s'_A, k', c') \leftarrow_{\$} \mathrm{snd}_A(s_A, \epsilon)$; $\mathrm{RcvB}(\epsilon, c')$). This requires Bob's state to become incompatible with Alice's state. In the kuKEM

setting, this corresponds to the secret key being updated with c' as associated data. Note that c' can be independent of Bob's state update from step 2 via c, and the computation of c' is controlled by the adversary. 4. Afterwards Bob's state is again exposed ($s'_B \leftarrow \text{ExposeB}((\epsilon, c)\|(\epsilon, c'))$). 5. Finally, Alice sends and establishes key k^* which is required to be secure ($c'' \leftarrow \text{SndA}(\epsilon, \epsilon)$). 6. Exposing Alice's state thereafter should not harm security of k^* ($s''_A \leftarrow \text{ExposeA}((\epsilon, c''))$).

We observe that, as with a kuKEM public key, Alice's state is publicly known during the entire attack. Only Alice's random coins when establishing k^* and updating her state, and Bob's random coins when sending, as well as his resulting state until he receives c' are hidden towards the adversary. We furthermore note that, by computing ciphertext c', the adversary controls Bob's state update. As a consequence, Bob's state update must reach forward-secrecy for key k^* with respect to adversarially chosen associated update data c' and Bob's resulting (diverged) state s'_B.

All in all, the security requirements highlighted by this attack emphasize the similarity of kuKEM's and SRKE's security. Nevertheless, we note that all our attempts to apply our proof technique for this case failed due to the above mentioned problems. Therefore, formally substantiating or disproving the intuition conveyed by this attack remains an open question for future work.

Open Questions and Impact With our work we aim to motivate research on another remaining open problem: can key-updatable KEM be instantiated more efficiently than generically from HIBE? It is, in contrast, evident that equivalence between HIBE and RKE is unlikely as constructions of the latter only utilize "one identity path" of the whole "identity tree" of the former.

Conclusively, we note that defining security for, and constructing schemes of interactive ratcheted key exchange variants (i.e., under bidirectional communication) is highly complicated and consequently error-prone.[4] By providing generic constructions (instead of ad-hoc designs) and grasping core components and concepts of ratcheted key exchange, complexity is reduced and sources of errors are eliminated. Additionally, our equivalence result serves as a benchmark for current and future designs of ratcheted key exchange – especially group RKE. For future constructions that only rely on standard public key cryptography either of the following questions may arise: how far is the adversary restricted such that our implication is circumvented, or how far is the construction secure under the respective security definition?

References

1. Alwen, J., Coretti, S., Dodis, Y.: The double ratchet: Security notions, proofs, and modularization for the Signal protocol. In: Ishai, Y., Rijmen, V. (eds.) EUROCRYPT 2019, Part I. LNCS, vol. 11476, pp. 129–158. Springer, Heidelberg (May 2019)
2. Balli, F., Rösler, P., Vaudenay, S.: Determining the core primitive for optimally secure ratcheting. Cryptology ePrint Archive, Report 2020/148 (2020), full version of this article. Available at `https://eprint.iacr.org/2020/148`

3. Bellare, M., Singh, A.C., Jaeger, J., Nyayapati, M., Stepanovs, I.: Ratcheted encryption and key exchange: The security of messaging. In: Katz, J., Shacham, H. (eds.) CRYPTO 2017, Part III. LNCS, vol. 10403, pp. 619–650. Springer, Heidelberg (Aug 2017)
4. Caforio, A., Durak, F.B., Vaudenay, S.: On-demand ratcheting with security awareness. Cryptology ePrint Archive, Report 2019/965 (2019), https://eprint.iacr.org/2019/965
5. Checkoway, S., Niederhagen, R., Everspaugh, A., Green, M., Lange, T., Ristenpart, T., Bernstein, D.J., Maskiewicz, J., Shacham, H., Fredrikson, M.: On the practical exploitability of dual EC in TLS implementations. In: Fu, K., Jung, J. (eds.) USENIX Security 2014. pp. 319–335. USENIX Association (Aug 2014)
6. Cohn-Gordon, K., Cremers, C.J.F., Dowling, B., Garratt, L., Stebila, D.: A formal security analysis of the signal messaging protocol. In: 2017 IEEE European Symposium on Security and Privacy, EuroS&P 2017, Paris, France, April 26-28, 2017. pp. 451–466 (2017)
7. Durak, F.B., Vaudenay, S.: Bidirectional asynchronous ratcheted key agreement with linear complexity. Cryptology ePrint Archive, Report 2018/889 (2018), https://eprint.iacr.org/2018/889
8. Durak, F.B., Vaudenay, S.: Bidirectional asynchronous ratcheted key agreement with linear complexity. In: Attrapadung, N., Yagi, T. (eds.) IWSEC 19. LNCS, vol. 11689, pp. 343–362. Springer, Heidelberg (Aug 2019)
9. Heninger, N., Durumeric, Z., Wustrow, E., Halderman, J.A.: Mining your ps and qs: Detection of widespread weak keys in network devices. In: Kohno, T. (ed.) USENIX Security 2012. pp. 205–220. USENIX Association (Aug 2012)
10. Jaeger, J., Stepanovs, I.: Optimal channel security against fine-grained state compromise: The safety of messaging. In: Shacham, H., Boldyreva, A. (eds.) CRYPTO 2018, Part I. LNCS, vol. 10991, pp. 33–62. Springer, Heidelberg (Aug 2018)
11. Jaeger, J., Stepanovs, I.: Optimal channel security against fine-grained state compromise: The safety of messaging. Cryptology ePrint Archive, Report 2018/553 (2018), https://eprint.iacr.org/2018/553
12. Jost, D., Maurer, U., Mularczyk, M.: Efficient ratcheting: Almost-optimal guarantees for secure messaging. In: Ishai, Y., Rijmen, V. (eds.) EUROCRYPT 2019, Part I. LNCS, vol. 11476, pp. 159–188. Springer, Heidelberg (May 2019)
13. Langley, A.: Source code of Pond (05 2016), https://github.com/agl/pond
14. Marlinspike, M., Perrin, T.: The double ratchet algorithm (11 2016), https://whispersystems.org/docs/specifications/doubleratchet/doubleratchet.pdf
15. Off-the-Record Messaging. http://otr.cypherpunks.ca (2016)
16. Poettering, B., Rösler, P.: Asynchronous ratcheted key exchange. Cryptology ePrint Archive, Report 2018/296 (2018), https://eprint.iacr.org/2018/296
17. Poettering, B., Rösler, P.: Towards bidirectional ratcheted key exchange. In: Shacham, H., Boldyreva, A. (eds.) CRYPTO 2018, Part I. LNCS, vol. 10991, pp. 3–32. Springer, Heidelberg (Aug 2018)
18. Rescorla, E., Salter, M.: Extended random values for tls (2009), https://tools.ietf.org/html/draft-rescorla-tls-extended-random-02
19. Yilek, S., Rescorla, E., Shacham, H., Enright, B., Savage, S.: When private keys are public: results from the 2008 debian openssl vulnerability. In: Proceedings of the 9th ACM SIGCOMM Internet Measurement Conference, IMC 2009, Chicago, Illinois, USA, November 4-6, 2009. pp. 15–27 (2009)