

Achieving Better Privacy for the 3GPP AKA Protocol

Abstract: Proposed by the 3rd Generation Partnership Project (3GPP) as a standard for 3G and 4G mobile-network communications, the AKA protocol is meant to provide a mutually-authenticated key-exchange between *clients* and associated network *servers*. As a result AKA must guarantee the indistinguishability from random of the session keys (key-indistinguishability), as well as client- and server-impersonation resistance. A paramount requirement is also that of client privacy, which 3GPP defines in terms of: *user identity confidentiality*, *service untraceability*, and *location untraceability*. Moreover, since servers are sometimes untrusted (in the case of roaming), the AKA protocol must also protect clients with respect to these third parties. Following the description of client-tracking attacks e.g. by using error messages or IMSI catchers, van den Broek *et al.* and respectively Arapinis *et al.* each proposed a new variant of AKA, addressing such problems. In this paper we use the approach of provable security to show that these variants still fail to guarantee the privacy of mobile clients. We propose an improvement of AKA, which retains most of its structure and respects practical necessities such as key-management, but which provably attains security with respect to servers and Man-in-the-Middle (MiM) adversaries. Moreover, it is impossible to link client sessions in the absence of client-corruptions. Finally, we prove that any variant of AKA retaining its mutual authentication specificities *cannot* achieve client-unlinkability in the presence of corruptions. In this sense, our proposed variant is optimal.

Keywords: privacy, security proof, AKA protocol

DOI Editor to enter DOI

Received ..; revised ..; accepted ...

1 Introduction

Authenticated key-exchange (AKE) protocols address a fundamental goal in cryptography, namely that of allowing two parties to communicate securely over an insecure channel (like the Internet, a radio-frequency channel, or a mobile telecommunications network). Such protocols are constructed in two steps. First, the two parties exchange information that allows them to establish several *short-term* session keys. These are

later used to secure and authenticate data exchanged between the parties. AKE schemes ensure that sensitive data or services are securely provided from a server to a legitimate client.

In the context of mobile networks, mobile services, such as calls, SMS privileges, or Internet use, are granted to clients over a secure channel, following an authenticated key-exchange protocol called AKA¹. This protocol was introduced by the Third Generation Partnership Project (3GPP), which wrote and has been maintaining the specifications of 3G telecommunication systems. Note that mobile services must be provided both in domestic, and in roaming scenarios, and only to clients who are entitled to them. Moreover, the nature and destination of such services should remain private.

The AKA protocol. The AKA scheme was developed at the end of the previous century, and is symmetric-key only. The mobile-network context imposes a peculiar architecture to the AKA design. Thus, typical 3G and 4G networks involve three types of participants. Mobile *clients* may subscribe to a single *operator*, thus becoming entitled to mobile services. The latter are provided across a secure channel, not by the operator, but by an intermediate *local* network operator that we call a *server*. In domestic scenarios, the server and operator are affiliated together, and can thus both be trusted. In the case of roaming, the server is associated to a different operator and only trusted to provide services, not to know the client's long-term secret values (known only to the client and the operator). Servers do know the short-term secret values necessary for the secure-channel establishment, in order to then provide the service. The AKA setting requires three participants, unlike the classical two-party AKE setting.

The AKA scheme relies on two long-term symmetric keys. As a peculiarity of the subscriber-operator architecture, clients are associated both with a unique client-key and with their operator's key. Since this key is shared between a potentially very large number of users, the clients store a value derived from both keys denoted Top_C , and the subscriber key. only a (one-way) function of this, and the client-key.

Three crucial constraints influenced the design of AKA in the 1990s. At that time (and even to the present day) SIM cards could not generate (pseudo)random numbers². As

¹ Despite the similarity of the acronyms, the AKA protocol is only an example of an AKE scheme. Other examples are for instance TLS/SSL or PACE/EAC.

² However, the next-generation SIM cards are able – and in fact expected – to be able to perform such computations [8].

a result of this limitation, the freshness of each session had to be guaranteed without user-generated nonces. As a second constraint, note that the communication between servers and operators is financially expensive, thus to be minimized. We note that some communication between these two parties is inevitable, since the server is untrusted, but must still be able to authenticate to the client as legitimate. In the case of AKA, operators minimize costs by generating *batches* of *authentication vectors*, containing short-term authentication and key-establishment values, but hiding the client's long-term secrets. Finally, the protocol is run over notoriously-noisy communication-channels, and it must thus be robust with respect to noise. As a result, AKA was designed to be *stateful*, allowing the two parties to authenticate to each other by means of a *sequence number*, which can be updated at each execution. Resistance to noise translates to a tolerance interval for the authenticated sequence number, which allows some leeway for each execution. Finally, in case the two parties become desynchronized, a procedure is provided to re-synchronize them.

In order to preserve client privacy, clients are associated with several identifiers. Apart from the permanent, unique identifier called an IMSI, the client is also associated (by the servers) with temporary identifiers TMSI, which are unique per server. In an initial identification phase, the server must ascertain the IMSI of the client, which the operator must know to generate the authenticated key-exchange data. However, since IMSI values are permanent and recognizable, clients can use their TMSI values instead. A new TMSI value is sent encrypted, but not authenticated, as part of the secure-channel data. Servers must store the TMSI values for each client that authenticates to them; moreover, the TMSI value that was last associated with a given IMSI is kept for a longer time, since it may be demanded by another server. Finally, as a backup, in case a TMSI value is not found, or not recognized, the server will also accept an identification by IMSI.

AKA Requirements. Apart from several *security requirements*, such as guaranteeing the secure delivery of services, as well as client-to-server and server-to-client authentication, an important additional concern for 3G/4G communication is that of client *privacy*. The 3GPP specifications [3] specifically demand: “user identity confidentiality”, “user untraceability”, and “user location confidentiality”, which are all formulated with respect to passive third-party eavesdroppers. The first notion requires that such an adversary may not learn the *permanent* user identifier IMSI of a client; the second refers to an adversary's ability of learning whether a given client uses the

same, or different services³; and finally the third requires that adversaries cannot learn a client's location (since the servers running the protocol are area-specific).

These explicit requirements ensure a minimal amount of client privacy. The AKA protocol was moreover designed to guarantee a measure of security against corrupted, or malicious servers. This requirement is particularly relevant in the case of roaming, in which case the server may be untrusted or more vulnerable. We consider the three-party architecture with possible server corruptions and formulate security with respect to servers in two properties, namely: (i) *state-confidentiality*: the property that servers cannot learn the client's secret key, the operator's secret key, nor the client's state; (ii) *soundness*: that servers cannot in fact successfully run a key-exchange protocol with the client unless aided by the operator.

Protocol vulnerabilities. In this paper we focus on the (provable) privacy of AKA, but also consider its security. Three attacks in the literature, namely IMSI catcher attacks [10], IMSI paging attacks [14, 25], and impersonation by server-corruption [19], already prove that AKA does not offer the desired degree of client privacy. IMSI catchers allow passive and active adversaries to track clients by exploiting the fact that during the protocol run, the server will require clients to send their permanent identifier IMSI if the TMSI value cannot be traced back to an IMSI. IMSI paging attacks exploit the lack of authentication in the TMSI-reallocation message.

Zhang et Fang [19] note that in the case of roaming, the local server providing mobile services may be poorly secured and corruptible. An immediate consequence is that the data transmitted over the secure channel by this server can be compromised; however, Zhang et Fang also showed that the danger persists even after the client has left the area serviced by the corrupted network, since the latter can impersonate the legitimate server in a new, uncorrupted, strongly-secured network.

An older work by Ateniese et al. [6] examines the problem of untraceable mobility, in particular noting an informal paradigm: nobody but the client should know both the client's identity and location at the same time. In this context, they provide solutions to achieving better privacy in 3-party settings; however, these solutions are extremely generic, which makes them hard to immediately to AKA. Moreover, note that server-operator communication takes place across a channel that is implemented differently in practice by different operators. The protocols proposed by [6] require this channel to be imple-

³ The exact 3GPP wording allows two interpretations. One is that an adversary must not distinguish between two possible services used by the same user; the other interpretation is that adversaries must not know whether a single, known service is provided to two different clients or to the same client. We formalized this property following the first meaning.

mented in a specific way. Finally, we note that, while highly worthwhile, the goal of preventing operators from learning a client’s operator is incompatible with the way authentication is currently done in the AKA protocol (operators need to prepare session information for servers, across a secure, and mutually authenticated channel which forcibly reveals the identity – and implicitly the location of the server).

Although the AKA protocol was designed to protect client privacy, several attacks can be run on the physical layer or by using side-channel information. Since we focus only on the privacy of AKA at the protocol layer, this type of attacks are out of scope. Another class of attacks that are out of scope for our paper are those that exploit faults in TMSI reallocation choices (for instance the fact that TMSIs are not updated sufficiently fast, or according to protocol), denial-of-service by jamming, or attacks enabled by version-shifting (forcing the user to employ a weaker, more vulnerable version of the protocol on 2G and 3G, rather than 4G) [25]. Such attacks, however, indicate that the infrastructure and protocols existing *around* the AKA handshake should be more closely examined; if possible, backward compatibility should be discouraged in order to guarantee better security and privacy.

We outline further vulnerabilities, e.g. [19, 26, 28], in Appendix ??, and show that as a consequence, the AKA protocol does *not* attain even the basic-most privacy requirements outlined by 3GPP. This suggests the need for stronger security.

Client indistinguishability. Motivated by a heightened concern for client privacy, Arapinis et al. [21] were the first to consider the notion of *untraceability* of clients with respect to external adversaries. They exploited the fact that adversaries can induce faulty behaviour (resynchronization) and then distinguish a client that is forced to perform this operation from a fully synchronized one. They modified the protocol to address this, and gave a formal-security proof using ProVerif. However, this proof idealizes the long-term state used in the protocol runs, making it unclear how far the guarantee holds for the true scheme. As we show in this work, their variant does not in fact guarantee user untraceability. The attack we present breaks client-indistinguishability by exploiting the fact that an adversary can forge the encrypted IMSI message proposed by Arapinis et al. In work orthogonal to ours, Khan et al. [14] also critically examined the modifications proposed by Arapinis et al., pointing out that they fall short of providing adequate security and privacy, whilst being impractical to implement. An important criticism addresses the PKI for clients and servers; as we explain below, in our own variant, we minimize the modifications both in terms of computation and administration costs.

Lee et al. [17] consider the untraceability of the 4G LTE (Long-term Evolution) protocol (similar to AKA, but with a different identifier- and key-management), but do not focus on

the handshake itself. Their main result is that in the absence of server corruptions, LTE is (weakly) untraceable against an active MiM adversary. They focus on the security of TMSI values, and their retransmission, but surprisingly their model cannot capture IMSI-catcher attacks (which directly impact privacy *without* server corruptions). The reason is that Lee et al. assume that the TMSI reallocation and usage is perfect in the sense that a TMSI value will always lead to the legitimate IMSI value, and that once the TMSI allocation process starts, the IMSI will never again be demanded. This is not true for the AKA procedure, for which, if the active attacker replaces the TMSI with a random message, the server will demand the client’s IMSI in clear. Lee et al. do not capture the attack by Arapinis, since they do not model desynchronizations, and they reduce the three-party setting to just two parties, by assuming that the server and the operator are the same entity⁴. Their security proofs rely on an assumption on the underlying cryptographic functions used in the protocol, but they did not study whether TUAK and MILENAGE (the two current instantiations of the cryptographic cipher suites) actually guarantee the required properties.

Finally, the AKA scenario is one use-case of the more generic constructions of anonymous, secure and authenticated channels presented by Alwen et al. [5]. Their security model presents a slightly different scenario than we consider here. For one thing, the constructed resources include one adversarial (eavesdropping) interface, and interfaces to clients and servers; in our setting, this does not capture the 3-party scenario. Furthermore, while the constructive cryptography (CC) framework enables the composability of the obtained resources by design, our model also allows for adaptive client corruptions (in the stronger models), which is problematic in CC. Finally, we note that their constructions rely on key-indistinguishable MACs (KI-MACs), which allow authentication, but prevent tags from leaking information about the keys they were generated with. In our work, we tried to keep as close as possible to the presently-used algorithm suites used in AKA, namely TUAK and MILENAGE. This precludes the use of KI-MACs. However, one of our main sub-results is giving a sufficient security requirement which must be provided by the cryptographic algorithms. It would be interesting to explore whether KI-MACs could be used as an underlying condition to modify existing suites or provide new ones.

Our contributions. We have two main contributions. The first is to show that the AKA protocol and two more promising improvements in the literature do not guarantee client-privacy,

⁴ In the absence of corruption, this treatment is justified, but it is incomplete in terms of real-world implementations.

nor security with respect to the server. As a consequence, our second main contribution is to present a variant of the AKA which provably guarantees the following five properties:

- **Key-indistinguishability:** a BPR-like [15] guarantee that the derived session keys are indistinguishable from random values of equal length. The adversary here is an active Man-in-the-Middle (MiM).
- **Impersonation security:** the guarantee that a MiM attacker cannot impersonate either the server (to the client), or the client (to the server).
- **Wide-weak client-unlinkability:** a property guaranteeing that two sessions run by the same client are not linkable with respect to a MiM attacker which may learn whether the server accepts the client’s authentication or not, but which cannot corrupt clients to learn their keys.
- **State-confidentiality:** the demand that (malicious) servers cannot learn the client’s secret key, the corresponding operator’s secret key, nor the client’s state (in particular, the sequence number). We assume that the server interacts with both operators and with clients, but we only guarantee security for the authenticated key-exchange protocol (we do not address the record-layer primitives).
- **Soundness:** the demand that (malicious) servers cannot make the client accept the server’s authentication (thus completing the key-derivation process), unless they are explicitly given authenticating information by a legitimate operator.

As an additional main result, we also prove that achieving a stronger degree of client-unlinkability (see below) is impossible while the AKA protocol retains its current structure. Consequently, the degree of unlinkability we achieve is in that sense, optimal.

A precise formalization. As an implicit, but significant contribution we formalize the five security and privacy properties outlined above. Consequently, we are able to describe two attacks against AKA, one against client-unlinkability, the other against soundness; then we also proceed to show that the proposed AKA improvements of Arapinis et al. [20] and van den Broek et al. [11] are *not* client-unlinkable.

This precise formalization also allows us to find a gap in an impossibility result regarding client-unlinkability in the presence of corruptions for mutually authenticated protocols [23]. We formulate a different impossibility result, which is more precise, and also more generic in the case of the AKA protocol. Another main contribution of this paper is to extend the narrow-forward-privacy impossibility result for a broader class of protocols including AKA and its variants.

Our improvement. Our improved variant of the protocol mostly retains the symmetric character of the current version. We allow TMSI values to be backed up by IMSIs; however, we bypass IMSI catcher attacks by sending IMSIs encrypted with a public-key IND-CCA-secure encryption scheme. We assume each operator has a PKE key-tuple, and each client stores the public key of the operator; this minimizes key-management problems. We note that, just as for the AKA protocol, we only use the (encrypted) IMSI value as an alternative for the randomly-chosen temporary identifier values TMSI. However, we choose to update the current TMSI value by using authenticated encryption (AES-GCM) as part of the server’s authentication message. The PKE scheme is also used when moving from area to area: thus, if the client switches from one server (in a given location) to another (in another location), the TMSI is *not* used. This allows us to reveal only the current area that the server is in, rather than the client’s past location (as is the case for the current version of AKA); furthermore, we minimize the duration for which TMSI values must be stored by servers. On the other hand, in order to preserve client-unlinkability, we require that TMSI values are at least as long (large) as the output of the PKE encryption.

We retain the structure of the authenticated key-exchange part of AKA, using the client- and operator-state to authenticate the two parties and the derived session keys. We show that while this feature of the AKA protocol remains in use, no client-untraceability can be achieved in the presence of corruptions. By removing the need for re-synchronization, we also implicitly prevent attacks which link client sessions depending on whether or not the re-synchronization procedure is used.

We note that our improvement follows guidelines by one of the leading mobile service providers in Europe [22]. Table 1 compares our proposal to the AKA protocol and to the two more promising variants we also analyze in this paper.

2 Privacy model

2.1 3GPP Privacy requirements

AKA Infrastructure. The mobile context for which the AKA protocol was designed contains three entities: (1) clients, which register with operators and are allowed to access a subset of services; (2) operators, which know the secret parameters of all their registered clients; and (3) local servers, which are tasked with providing services to mobile clients, but are not trusted to know the clients’ personal information. In the AKA literature, operators are usually called home local registers (HLR), while servers are known as VLR.

	Defeating:			Security:				
	Attack n°1	Attack n°2	Attack n°3	Prop. n°1	Prop. n°2	Prop. n°3	Prop. n°4	Prop. n°5
3G AKA	x [11]	x [21]	x	x	x	x	?	?
Arapinis	✓ [21]	✓ [21]	x	✓	x	x	?	✓* [21] ⁵
Van Den Broek	✓ [11]	x	✓	✓	x	x	?	?
Our variant	✓	✓	✓	✓	✓	x	✓	✓
Attack n°1: IMSI Catcher § 3.3.1.				Prop. n°1: Confidentiality of the previous location §2.				
Attack n°2: Linkability of failure messages § 3.3.1.				Prop. n°2: ww-unlink §2.2.				
Attack n°3 : Our traceability attack § 3.3.2.				Prop. n°3: nf-unlink §2.2.				
Prop. n°4: State-confidentiality & soundness §2.2.				Prop. n°5: Key-indistinguishability & Client- and Server-impersonation §2.2.				

Fig. 1. Comparison between several AKA variants. For attacks, a ✓ denotes the protocol resists the attack, while x denotes a vulnerability; for properties, a x denotes the property is *not* achieved, while a ✓ indicates security with respect to that property.

The security demands of 3G/4G networks are client-centric, revolving around the following parameters related to mobile clients (users) C:

- IMSI : a permanent identifier, unique per customer and highly trackable;
- TMSI : a temporary identifier, unique per server, and updated after each successful protocol run;
- LAI : a unique local-area identifier per server; client store (TMSI, LAI) tuples whenever a server issues a new TMSI;
- sk_C : the client’s unique client key;
- sk_{op} : the key of the operator C subscribes to;
- $Sqn_C, Sqn_{Op,C}$: the client’s state Sqn_C has an equivalent operator state $Sqn_{Op,C}$, which should not be “too far” from the client’s state. The state is updated by the client upon authenticating the server (correct verification of the authentication challenge); the server updates state upon authenticating the challenge (verification of the authentication response).

We refer the reader to Section 3.2 for more details about the protocol description.

The identifier and key-management schemes are as follows. Clients may know the permanent value IMSI, the temporary identifiers TMSI and LAI, the keys sk_C, sk_{op} and a function of sk_C and sk_{op} (they do not store sk_{op} in clear, as discussed in section 3.2). Operators know the tuple: (IMSI, sk_C, sk_{op}). Servers keep track of tuples (TMSI, LAI, IMSI). Furthermore, both servers and operators know the sequence number that the operator associated to each client. The clients update their own sequence numbers, which are highly related to the operator’s sequence number.

Very notably, servers must both authenticate and exchange session keys with mobile clients, despite not knowing their secret material.

Client Privacy. The Third Generation Partnership Project (3GPP), which designed the AKA protocol in the TS.33.102 specification [3], lists the following privacy requirements:

- **user identity confidentiality:** specifically, “*the property that the permanent user identity (IMSI) of a user [...] cannot be eavesdropped on the radio access link.*”
- **user untraceability:** namely, “*the property that an intruder cannot deduce whether different services are delivered to the same user by eavesdropping on the radio access link.*”
- **user location confidentiality:** in particular, “*the property that the presence or the arrival of a user in a certain area cannot be determined by eavesdropping on the radio access link.*”

The requirements quoted above are quite informal; moreover, the nomenclature is confusing, since in the *provable-security* literature, *untraceability* refers to adversaries tracing clients in distinct protocol runs (rather than it being service-related). We discuss the three requirements below, then formalize them into cryptographic requirements.

User identity confidentiality concerns only the client’s *permanent* IMSI value (not, e.g. the client’s sequence number) with respect to *passive* attackers (rather than active ones). However, mobile networks are notoriously prone to Man-in-the-Middle (MiM) active attacks like the IMSI catcher [10], which allows a third party (the MiM) to recover a client’s IMSI. Another highly-trackable client-specific parameter is the sequence number Sqn , whose updating procedure is very simplistic and its output, predictable even without a secrecy key. As a consequence we require the stronger property of *provable unlinkability*, which ensures that even an *active* MiM cannot *link* two AKA protocol runs to the same client.

For user untraceability, no attacker must know whether the same service (i.e. *any* message-exchange over the secure channel) is provided to a client multiple times. From the point of view of provable security, this is equivalent to *key-*

indistinguishability if the authenticated-encryption algorithms are assumed to be secure.

User location confidentiality demands that eavesdroppers \mathcal{A} cannot detect the presence of a client in a given area; however, the definition does not specify what information \mathcal{A} links to each client (e.g. the IMSI, the sequence number, etc.). Attackers are aware of the current LAI; the difficulty lies in learning which clients *enter* the area. Unfortunately the AKA protocol always reveals the *past* location of any arriving client, making unique (or rare) itineraries stand out. We formalize a strong degree of location privacy as a part of client-unlinkability.

Our formalizations of *client unlinkability* and *key-indistinguishability* consequently guarantee 3GPP’s three privacy requirements.

Implicit requirements. As discussed in Section 1, the AKA protocol implicitly addresses security with respect to malicious servers, which are restricted as follows: (1) the servers have no access to the tuple (sk_C, sk_{Op}) ; (2) the (hence necessary) operator-server communication must be minimized in order to minimize costs.

We formulate the following two *implicit* requirements:

- **State-Confidentiality:** Servers must not learn any client-related long-term state.
- **Soundness:** Clients must reject authentication-challenges not explicitly provided by the operator to the server.

2.2 Security models

Due to their orthogonality, it is hard to formalize the notions of key-indistinguishability and client-unlinkability in the same generic framework. One difficulty is the fact that the unlinkability notion requires the adversary to have access to clients without knowing their identities. Following established approaches [13, 24], in the unlinkability model, we associate clients with identifiers, or handles, denoted $VC\%$ (Virtual Client), and this changes the syntax of the oracles we use. Thus, we differentiate between the model for security (including notions of key-indistinguishability, client- and server-impersonation resistance against a MiM adversary, and state-confidentiality and soundness with respect to malicious servers), and that of client-unlinkability. We use similar oracles, with a slightly different syntax, for the two types of definitions, and thus obtain security guarantees based on traditional Bellare, Pointcheval, and Rogaway models [15].

Set up and participants. We consider a set \mathcal{P} of participants, which are either a server S_i or a mobile client C_i of the type respectively VLR or ME/USIM. By contrast operators Op are not modeled as active parties. In all security games apart from state-confidentiality and soundness with respect to the server,

the operators Op are black-box algorithms within the server S ; in those two games, the operators are oracles, which the adversary (i.e. the server) may query. We assume the existence of n_C clients, n_S servers, and n_{Op} operators. If the operators are contained within the servers, we assume that all copies of the same operator Op are synchronized at all times. We associate each client with: a long-term, static secret state consisting of a tuple (sk_C, sk_{Op}) , an ephemeral state st_C consisting of a sequence number Sqn_C , a tuple of a static, permanent identifier IMSI and an ephemeral, temporary identifier TMSI, and finally a tuple of a current, and a past local area identifier, denoted $past.LAI_P$ and $curr.LAI_P$ respectively. Servers are associated with a permanent local area identifier LAI and a unique network identifier ID_{S_i} ; they also keep track of a list of tuples (TMSI, IMSI) associated with clients. Each of the at most n_S servers has black-box access to algorithms (or oracles in the case of state-confidentiality and soundness) $Op_1, \dots, Op_{n_{Op}}$, which are initialized with long-term keys (sk_{Op_i}) and keep track of a list of tuples (IMSI, sk_C , $Sqn_{Op,C}$). In our model, we also assume that the key space of all operators is identical (otherwise it becomes easier to distinguish between clients of different operators).

2.2.1 Client Unlinkability

Client-Unlinkability.

Informally, we call a protocol Π client-unlinkable if no adversary can know whether two executions of Π were run by the same, or by two different clients. Two sessions associated with the same client are called *linked*. Following previous works of Vaudenay [24] and Hermans et al. [13], we give the adversary access to a basic left-or-right oracle, which associates an anonymized handle virtual client VC to one of two possible clients (input by the adversary). We extend this framework to account for client mobility, giving the adversary access to a relocation oracle. Consequently, if an attacker can distinguish between clients based on their location, they will win the unlinkability game, which we detail below.

At the onset of this game, the set of clients is empty and the challenger instantiates two lists \mathcal{L}_{drawn} and \mathcal{L}_{free} . We initialize operators by choosing their secret keys. The adversary can then initialise servers by choosing their locations, and it can create clients to populate the system it attacks. For each newly-created client, the past location $past.LAI_C$ is set to a special symbol \perp and the current location is adversarially-chosen.

The adversary then interacts with clients by means of several oracles. The lists \mathcal{L}_{drawn} and \mathcal{L}_{free} correspond to the two possible states of any one client. Clients can be “drawn” or

“free”; at creation, all clients are “free”, and they may become “drawn” if used as input to a left-or-right Client-Drawing oracle. In particular, we use a left-or-right Client-Drawing oracle (similar to that in [13, 24]), which allows the adversary to interact with one of two clients (the interacting client being chosen depending on a secret bit b). Clients input to the Drawing oracle are moved to the $\mathcal{L}_{\text{drawn}}$ list; further Drawing queries can then be made concurrently as long as the input clients are in the $\mathcal{L}_{\text{free}}$ list⁶. Upon drawing one of two possible clients, the adversary is given a handle on the chosen entity; following the notation of [24], we call this a *virtual* client and we associate it with the handle VC. Virtual clients can then be freed by the adversary (this would remove them from the $\mathcal{L}_{\text{drawn}}$ list and re-add them to the $\mathcal{L}_{\text{free}}$ list). Only free clients can be drawn. This oracle associates a handle to either the left or the right input client, depending on a secret bit b .

The client unlinkability property is defined in terms of the following security experiment $\text{Exp}_{\mathcal{A}}^{\text{c.unlink}}(1^\lambda)$, for a security parameter (in unary) 1^λ .

- The challenger randomly chooses a bit $b \in \{0, 1\}$.
- The adversary may use the oracles below (with restrictions depending on its adversarial class), and the challenger answers the queries.
- The adversary finally outputs a guess d of the bit b .

We say the adversary *wins* if and only if $d = b$, and we define the adversary’s advantage of winning this game against a protocol Π as:

$$\mathcal{A}_{\Pi}^{\text{c.unlink}}(\mathcal{A}) := |\Pr[\mathcal{A} \text{ wins } \text{Exp}_{\mathcal{A}}^{\text{c.unlink}}(1^\lambda)] - \frac{1}{2}|.$$

Recalling the adversarial classes of [13, 24], we call an adversary *narrow* if it may not use the **Result** oracle, permitting it to know whether the server has authenticated the client or not. The opposite of *narrow* are *wide* adversaries. Orthogonal to the use of the **Result** oracle, we also classify adversaries in terms of their use of the **Corrupt** oracle, which gives them access to the client’s data. Thus, adversaries are *weak* if they cannot use the corruption oracle; they are *forward* if any corruption query may only be followed by more corruption queries⁷. Finally, adversaries are classified as *strong* if their access to oracles is unrestricted.

We note that the 3GPP requirements outlined in Section 2.1 restrict their adversaries to “*eavesdroppers on the ra-*

dio link”, which seems to indicate that they target *weak* adversaries that are either *narrow* or *wide*. Moreover, they restrict their adversaries to being passive only; in this paper, we also consider active weak attackers and thus obtain a better privacy guarantee.

Formalization. We quantify adversaries in terms of the following parameters: the adversarial class, which we abbreviate to α -c.unlink, with $\alpha \in \{\text{nw}, \text{ww}, \text{nf}, \text{wf}\}$ (for narrow- and wide-weak, and narrow-, respectively wide-forward adversaries); their execution time is t ; the maximum number q_{exec} of sessions instantiated per client C ; the maximum number q_{id} of user identification per session; and the maximum number q_G of queries to the function G . We formalize the following definitions.

Definition 1. [Weak Unlinkability] A protocol Π is $(t, q_{\text{exec}}, q_{\text{id}}, q_G, \epsilon)$ -nw/ww-client-unlinkable if no narrow/wide-weak-adversary running in time t , creating at most q_{exec} sessions and q_{id} user identification per session, and making at most q_G queries to the function G , has an advantage $\text{Adv}_{\Pi}^{\text{w.c.unlink}}(\mathcal{A}) \geq \epsilon$.

Definition 2. [Forward Unlinkability] A protocol Π is $(t, q_{\text{exec}}, q_{\text{id}}, q_G, \epsilon)$ -nf/wf-client-unlinkable if no narrow/wide-forward-adversary running in time t , creating at most q_{exec} sessions and q_{id} user identification per session, and making at most q_G queries to the function G , has an advantage $\text{Adv}_{\Pi}^{\text{f.c.unlink}}(\mathcal{A}) \geq \epsilon$.

Oracles. The adversary interacts with the system by means of the following oracles, in addition to a function G , which we model as a PRF and which “encompasses” all the cryptographic functions of the AKA protocol:

- $\text{CreateCl}(\text{Op}, \text{LAI}) \rightarrow (C_i, \text{IMSI}, \text{st}_{C_i})$: this oracle creates a new, legitimate, free client, labelled C_i at a location LAI for which a server is already defined (else the oracle outputs \perp). The client’s IMSI, its sequence number Sqn_{C_i} , and its secret key sk_{C_i} are chosen uniformly at random from sets ID , ST , and \mathcal{S} respectively; the past location and TMSI are set to a special symbol \perp . The client’s operator key sk_{op} is set to the key of the operator Op. The adversary is given the parameters IMSI, st_{C_i} , and the label C_i (used later to select clients).
- $\text{CreateS}(\text{LAI}) \rightarrow S_i$: this oracle generates a new server S_i at location LAI, if this location is not already defined for another server (else the oracle returns \perp).
- $\text{Launch}(\text{VC}, S_j) \rightarrow (s, m)$: this oracle instantiates a new session (labelled by a unique identifier s) between the client associated with VC and the server S_j , and outputs

⁶ In particular, we want to avoid trivial attacks, in which an adversary can distinguish a client simply because it is not in its original state (having already started the protocol run beforehand).

⁷ In particular, the adversary may no longer free drawn clients, nor interact with servers or clients

an initial protocol message m from S_j to VC. This oracle keeps track of tuples (s, VC, S_j) .

- $\text{DrawCl}(C_i, C_j) \rightarrow VC$: on input a pair of client labels, this oracle generates a handle VC, which is a monotonic counter, if the following conditions are met: (a) both clients were free when the query was made; (b) both clients have the same current location value. If either condition is not met, the oracle outputs \perp . Else, depending on the value of the secret bit b , the challenger associates the handle VC either with C_i (if $b = 0$) or with C_j (if $b = 1$). The challenger stores the triple (VC, C_i, C_j) in a table \mathcal{T} .
- $\text{FreeVC}(VC) \rightarrow \perp$: on input the virtual handle VC, this oracle retrieves the values C_i, C_j associated to VC in the table \mathcal{T} , aborting any ongoing protocol runs.
- $\text{Relocate}(VC, \text{LAI}^*) \rightarrow \perp$: this oracle modifies the *current* location of the two clients C_i, C_j associated with VC in \mathcal{T} , to LAI^* . In particular, the challenger does the following for each of the clients: (1) it sets $\text{past.LAI} := \text{curr.LAI}$; (2) it sets $\text{curr.LAI} := \text{LAI}^*$. Any protocol sessions still running for VC are aborted.
- $\text{Send}(P, s, m) \rightarrow m'$: for the first input, the adversary can input either a handle VC or a server identity S. In the former case, the oracle simulates sending the message m from the adversary to the client associated with VC in session s , returning either the party's message m' or \perp if either s is not associated with VC or if VC does not exist. Parties may also return $m' = \perp$ as an error message. If the first input of this oracle is set to S, the oracle simulates sending the message from the adversary to the server S.
- $\text{Execute}(VC, S, s) \rightarrow \tau$: this oracle simulates a complete protocol run between the client associated with VC and the server S. in the presence of a passive adversary. In particular, by alternating SendToCl and SendToS queries on genuinely-output messages, this oracle generates and outputs the transcript τ of the execution between the server S and the client C for which session s was created.
- $\text{Result}(P, s) \rightarrow \{0, 1\}$: if $P = VC$, this oracle returns a bit indicating whether the client associated with VC has accepted the server that VC ran session s with. If $P = S$, then the bit indicates whether the server accepted the client. For the AKA protocol, an acceptance bit is equivalent to the confirmation of the key-exchange. If the session is incomplete or session s is not associated with P, the oracle returns \perp .
- $\text{Corrupt}(C) \rightarrow \{\text{sk}, \text{st}_C, \text{ID}_C, (\text{past.LAI}, \text{curr.LAI})\}$: For a client C, this oracle returns the full state (static and ephemeral), the identifiers and the location information of client C.

2.2.2 Key-indistinguishability and Impersonation

The notion of *key-indistinguishability* refers to the session keys computed as a result of authenticated key exchange, requiring that they be indistinguishable from random bitstrings of equal length. We use a subset of our previously-defined oracles, this time without anonymizing clients by use of handles. Additionally, the adversary may also know the ephemeral state (in our case, the sequence number) of both clients and servers. Since the state is updated in a probabilistic way, we give the adversary a mean of always learning the updated state of a party without necessarily corrupting it (the latter may rule out certain interactions due to notions of freshness, see below). Corruption results in *adversarially controlled* parties.

In this model, participants (clients and servers) may run concurrent key-agreement executions; we denote the j -th execution of the protocol of party P as P_j . Each client C is associated with a *unique* identifier UID (which in the case of AKA is the IMSI); the identity of the server is of the form S_i . We simplify the key-indistinguishability model by abstracting location data (since it does not affect the security of the session keys).

We associate each instance P_i with a session ID sid , a partner ID pid , and an accept/reject bit accept . The partner ID pid is set to either the server S_i or to a user identifier UID; the session ID sid includes four values: the user ID given by client UID (and implicitly sk_{UID}), the server identifier ID_{S_i} , the randomness generated by the server, and the sequence number used for the authentication. Finally, the accept/reject bit is initialized to 0 and turns to 1 at the successful termination of the key-agreement protocol. We call this "terminating in an accepting state". A successful termination of the protocol yields, for each party, a session key K (which for the AKA protocol consists of two keys), the session identifier sid , and the partner identifier pid of the party identified as the interlocutor. We allow adversaries to learn whether instances have terminated (by sending Send queries) and whether they have accepted or rejected their partners. We also assume that the adversary will learn the session and partner identifiers for any session in which the instance has terminated in an accepting state.

Partners. Each instance of each party keeps track of a session ID string, denoted sid , consisting of the four values listed above. We define partners as party instances that share the same session ID. More formally:

Definition 3. [Partners.] *Two instances P_i and P'_j are partners if the following statements hold:*

- (i) *One of the parties is a user and the other is the server.*
- (ii) *The two instances terminate in an accepting state.*
- (iii) *The instances share the same sid.*

In this case, the partner ID of some party P denotes its (intended) partner.

Formalization. In the key-indistinguishability game, we no longer need to formalize the drawing and freeing of clients; thus we do not use those oracles. We furthermore do not use the Relocation oracle. Instead, we give the adversary access to a Key-Reveal oracle, which returns session keys for an ongoing session terminated in an accepting state. The central oracle in this game is a real-or-random type of Testing oracle, which allow the adversary to know either a tuple of real session keys or a tuple of random keys of equal size. If the adversary can tell whether the keys are real or random, then she is said to win the key-indistinguishability game.

At the outset of the game, the challenger first generates the keys of all the n_{Op} operators and instantiates a number n_S of servers. Each server can then access synchronized copies of the n_{Op} operators internally. Each copy of an operator Op run internally by a server S takes as input the identity of that specific server S (this models the fact that operators and servers communicate via an unspecified secure channel). The adversary is then allowed to query any of the oracles below. We implicitly assume that the $\text{Test}^{K.\text{Ind}}$ oracle keeps state and, once it is queried a first time, it will return \perp on all subsequent queries (we only allow a single query). However, we do allow the adversary to interact with other oracles after the $\text{Test}^{K.\text{Ind}}$ query as well.

Eventually, the adversary \mathcal{A} outputs a bit d , which is a guess for the bit b used internally in the $\text{Test}^{K.\text{Ind}}$ oracle. The adversary *wins* if and only if: $b = d$ and \mathcal{A} has queried a fresh instance to the $\text{Test}^{K.\text{Ind}}$ oracle. We consider the following definition of a fresh instance for the key-indistinguishability. We note that this notion is classical in symmetric-key protocols.

Definition 4. [Freshness : key – indistinguishability.] *An instance P_i is fresh if neither this instance, nor a partner of P_i is adversarially-controlled (its long-term key sk_P has not been corrupted) and the following queries were not previously executed:*

- (i) **Reveal**(.), either on the instance P_i , or on of its partners.
- (ii) **Corrupt**(.) on any instance, either of P , or of their partners.

The advantage of \mathcal{A} in winning the key-indistinguishability game is defined as:

$$\text{Adv}_{\Pi}^{K.\text{Ind}}(\mathcal{A}) := |\Pr[\mathcal{A} \text{ wins}] - 1/2|.$$

Definition 5. [Key – indistinguishability.] *A key-agreement protocol Π has $(t, q_{\text{exec}}, q_{\text{id}}, q_{\text{serv}}, q_{Op}, q_G, \epsilon)$ -key-*

indistinguishability if no adversary running in time t , creating at most q_{exec} party instances, with q_{id} user identification per instance, corrupting at most q_{serv} servers, making at most q_{Op} OpAccess queries per operator per corrupted server, and making at most q_G queries to the function G , has an advantage $\text{Adv}_{\Pi}^{K.\text{Ind}}(\mathcal{A}) \geq \epsilon$.

We also define the related notions of client- and server-impersonation security, which involves a slightly different notion of freshness.

Definition 6. [Freshness : Imp.Sec.] *An instance P_i , with session ID sid and partner ID pid , is fresh if: neither this instance nor a partner of P_i is adversarially-controlled; and if there exists no instance P'_j sharing session sid with the partner $pid = P_i$ (the related transcript is denoted as (m, m', m'')) such that the following events occur:*

- (i) *The message m is sent by the adversary \mathcal{A} to P_i via a $\text{Send}(m)$ query at time $\text{clock} = k$, yielding message m' at time $\text{clock} = k + 1$.*
- (ii) *The message m' is sent by \mathcal{A} to P'_j via a $\text{Send}(m')$ query at time $\text{clock} = k' > k + 1$, yielding message m'' at time $\text{clock} = k' + 1$.*
- (iii) *The message m'' is sent by \mathcal{A} to P_i via a $\text{Send}(m'')$ query at time $\text{clock} = k'' > k' + 1$.*

We note that the messages need not be exactly sequential (i.e. the adversary could query other oracles in different sessions before returning to session sid). Furthermore, the notion of freshness only refers to relays with respect to the partner client pid . We do not restrict the adversary from forwarding received messages to other server or client instances.

The goal of an impersonation adversary is to make a fresh party instance terminate in an accepting state. In this case, the $\text{Test}^{K.\text{Ind}}$ oracle is not used. More formally, the game begins by generating the operator keys and servers as before; then the adversary \mathcal{A} gains access to all the oracles except $\text{Test}^{K.\text{Ind}}$. When \mathcal{A} stops, she *wins* if there exists an instance S_i for client-impersonation (resp. an instance C_i for the server-impersonation) that ends in an accepting state and is fresh as described above. The advantage of the adversary is defined as her success probability, i.e.

$$\text{Adv}_{\Pi}^{C.\text{Imp}}(\mathcal{A}) := \Pr[\mathcal{A} \text{ wins}], \text{Adv}_{\Pi}^{S.\text{Imp}}(\mathcal{A}) := \Pr[\mathcal{A} \text{ wins}].$$

Definition 7. [Impersonation security.] *A key-agreement protocol Π is $(t, q_{\text{exec}}, q_{\text{id}}, q_{\text{serv}}, q_{Op}, q_G, \epsilon)$ -impersonation-secure if no adversary running in time t , creating at most q_{exec} party instances, with q_{id} user identification per instance, corrupting at most q_{serv} servers, making at most q_{Op} OpAccess queries per operator per corrupted server, and making at most q_G*

queries to the function G , has an advantage $\text{Adv}_{\Pi}^{\text{C.Imp}}(\mathcal{A}) \geq \epsilon$ or $\text{Adv}_{\Pi}^{\text{S.Imp}}(\mathcal{A}) \geq \epsilon$.

Oracles: K.Ind and Imp.Sec . The adversary interacts with the system by means of the following oracles, in addition to a function G , which we model as a PRF.

- $\text{CreateCl}(\text{Op}) \rightarrow (\text{UID}, \text{st}_{\text{UID}})$: This oracle creates a client with unique identifier UID . Then the client's secret key sk_{UID} and the sequence number Sq_{UID} . The tuples $(\text{UID}, \text{sk}_{\text{UID}}, \text{sk}_{\text{Op}}, \text{Sq}_{\text{UID}})$ are associated with the client UID and with the corresponding operator Op (i.e. each “copy” of Op in each server does this). The operator sets $\text{st}_{\text{Op,UID}} := \text{Sq}_{\text{UID}}$ and then keeps track of $\text{st}_{\text{Op,UID}}$. The adversary is given UID and st_{UID} .
- $\text{NewInstance}(\text{P}) \rightarrow (\text{P}_j, m)$: this oracle instantiates the new instance P_j , of party P , which is either a client or a server. Furthermore, the oracle also outputs a message m , which is either the first message in an honest protocol session (if P is a server) or \perp (if P is a client). The state st of this party is initiated to be the current state of P , and it is initiated with the current value of TMSI , LAI .
- $\text{Execute}(\text{P}, i, \text{P}', j) \rightarrow \tau$: creates (fresh) instances P_i of a server P and P'_j of a client, then runs the protocol between them. The adversary \mathcal{A} receives the transcript of the protocol.
- $\text{Send}(\text{P}, i, m) \rightarrow m'$: simulates sending message m to instance P_i of P . The output is a response message m' (which is set to \perp in case of an error or an abort).
- $\text{Reveal}(\text{P}, i) \rightarrow \{\text{K}, \perp\}$: if the party has not terminated in an accepting state, this oracle outputs \perp ; else, it outputs the session keys computed by instance P_i .
- $\text{Corrupt}(\text{P}) \rightarrow \text{sk}_{\text{P}}$: if P is a client, this oracle returns the long-term client key sk_{P} , but not sk_{Op} (in this we keep faithful to the implementation of the protocol, which protects the key even from the user himself). If P is corrupted, then this party (and all its instances, past, present, or future), are considered to be adversarially controlled. If P is a server, then this oracle returns the identifier S_i , giving the adversary access to a special oracle OpAccess .
- $\text{OpAccess}(\text{S}, \text{C}) \rightarrow m$: for a corrupted server S , this oracle gives the adversary one access to the server's local copy of all the operators, in particular returning the message that the operator Op would have output to the server on input a client C .
- $\text{StReveal}(\text{P}, i, \text{bits}) \rightarrow x$: for a client P , if $\text{bits} = 0$, then this oracle reveals the current state of P_i ; else, if $\text{bits} = 1$, then the oracle returns the state the operator stores for P .

- $\text{Test}^{\text{K.Ind}}(\text{P}, i) \rightarrow \hat{\text{K}}$: this oracle is initialized with a secret random bit b . It returns \perp if the instance P_i is unfresh or if it has not terminated in an accepting state (with a session key K). If $b = 0$, then the oracle returns $\hat{\text{K}} := \text{K}$, else it returns $\hat{\text{K}} := \text{K}'$, which is a value drawn uniformly at random from the same space as K . We assume that the adversary makes a single $\text{Test}^{\text{K.Ind}}$ query (a standard hybrid argument can extend the notion to multiple queries).

2.2.3 Further security notions: Security w.r.t. servers

In this section, we define two further notions, namely soundness and state-confidentiality with respect to servers. In these games the adversary is a malicious, but legitimate server S , and this is the only server we consider. We note that 3GPP specifications allow servers to communicate with each other, but *how* they do this is not apparent. We use the OpAccess oracle to give the server access to operators on specific clients, but change the syntax so that the oracle takes a single input, namely a client identifier C . We also demand that the output of this oracle represents material only for a single protocol session sid . The adversary uses the Send , CreateCl , NewInstance , Execute , and StReveal oracles as described in the key-indistinguishability model. We additionally modify the corruption oracle, as noted below:

- $\text{Corrupt}(\text{P}) \rightarrow \text{sk}_{\text{P}}$: if P is a client, this oracle returns the long-term client key sk_{P} , but not sk_{Op} (in this we keep faithful to the implementation of the protocol, which protects the key even from the user himself). If P is an operator, then this oracle returns sk_{Op} and the list of tuples $(\text{UID}, \text{sk}_{\text{UID}}, \text{st}_{\text{Op,C}})$ for all clients C subscribing with that operator.

State-Confidentiality. Unlike key-indistinguishability, which guarantees that *session* keys are indistinguishable from random with respect to MiM adversaries, the property of state confidentiality demands that *long-term* client keys remain confidential with respect to *malicious servers*

This game begins by generating the material for n_{Op} operators and n_{C} clients. The adversary can then interact arbitrarily with these entities by using the oracles above. At the end of the game, the adversary must output a tuple: $(\text{P}_i, \text{sk}_{\text{UID}}^*, \text{sk}_{\text{Op}}^*, \text{st}_{\text{UID}}^*, \text{st}_{\text{Op,UID}}^*)$ such that UID is the long-term identifier of P and P_i is a fresh instance of P in the sense formalized below. The adversary wins if at least one of the values: $\text{sk}_{\text{UID}}^*, \text{sk}_{\text{Op}}^*, \text{st}_{\text{UID}}^*, \text{st}_{\text{Op,UID}}^*$ is respectively equal to $\text{sk}_{\text{UID}}, \text{sk}_{\text{Op}}, \text{st}_{\text{UID}}, \text{st}_{\text{Op,UID}}$, the real secret values of the fresh instance P_i .

Definition 8. [Freshness : St.conf] An instance P_i is fresh if neither this instance, nor a partner of P_i is adversarially-controlled (its long-term key sk_P has not been corrupted) and the following queries were not previously executed:

- (i) $\text{StReveal}(\cdot)$ on any instance of P .
- (ii) $\text{Corrupt}(\cdot)$ on any instance of P or on the operator Op to which P subscribes.

The advantage of the adversary is defined as:

$$\text{Adv}_{\Pi}^{\text{St.conf}}(\mathcal{A}) := \Pr[\mathcal{A} \text{ wins}].$$

Definition 9. [State – Confidentiality.] A key-agreement protocol Π is $(t, q_{\text{exec}}, q_{\text{id}}, q_{\text{Op}}, q_G, \epsilon)$ -state-confidential if no adversary running in time t , creating at most q_{exec} party instances, with at most q_{id} user identification per instance, making at most q_{Op} queries to any operator Op , and making at most q_G queries to the function G , has an advantage $\text{Adv}_{\Pi}^{\text{St.conf}}(\mathcal{A}) \geq \epsilon$.

In the *Soundness* game, we demand that no server is able to make a fresh client instance terminate in an accepting state without help from the operator. This game resembles impersonation-security; however, this time the adversary is a legitimate server (not a MiM) and it has access to operators. The adversary may interact with oracles in the soundness game arbitrarily, but we only allow a maximum number of q_{Op} queries to the OpAccess oracle per client. The adversary wins if there exist $(q_{\text{Op}} + 1)$ fresh client instances of a given client which terminated in an accepting state. Freshness is defined as in the impersonation game. The advantage of the adversary is defined as:

$$\text{Adv}_{\Pi}^{\text{S.sound}}(\mathcal{A}) := \Pr[\mathcal{A} \text{ wins}].$$

Definition 10. [Soundness w.r.t. server.] A key-agreement protocol Π is $(t, q_{\text{exec}}, q_{\text{id}}, q_{\text{Op}}, q_G, \epsilon)$ -server-sound if no adversary running in time t , creating at most q_{exec} protocol instances, with at most q_{id} user identification per instance, making at most q_{Op} queries to any operator Op , and making q_G queries to the function G , has an advantage $\text{Adv}_{\Pi}^{\text{S.sound}}(\mathcal{A}) \geq \epsilon$.

3 The AKA protocol

3.1 Notations

Notations. Throughout the rest of the document, we denote the length of a bitstring x by $|x|$, while $[x]_{i..j}$ denotes bits i

through j of x . If f is a function, then $y \leftarrow f(x)$ denotes that y was the output of f on input x . Similarly, $y \xleftarrow{\$} \{0, 1\}^n$ indicates that y is chosen uniformly at random from the set of bitstrings of length n , i.e. $\{0, 1\}^n$. For strings x, y , the string $x||y$ is the concatenation of x and y , while $x \oplus y$ denotes the exclusive OR (XOR) of x and y . For a bit b , \mathbf{b}^n denotes the string $b||b||\dots||b$ of length n . We denote by \perp the error message and by 1^λ the security parameter (in unary). The indistinguishability property is considered against the chosen plaintext attacks (*ind – cpa*). We denoted this property *ind – cpa* to lighten the expressions.

3.2 Description of the AKA protocol

3G (and 4G) mobile networks use a variant of the AKA protocol, which is fully depicted in Appendix 3.2, in order to establish secure channels between mobile clients and servers. Ultimately, the server uses the secure channel to transmit a specific service to the mobile client.




Fig. 2. The five phases of the AKA Procedure.

This protocol is actively run by clients and servers in the (selectively-active) presence of an operator. Servers and operators communicate over a secure and private channel; however, the server is considered only partially trusted. Section 2 describes in detail the setup of the three parties.

The AKA protocol consists of five phases. The first phase, *user identification*, is run by a client C and a server S on the *insecure channel* and it allows S to associate C to an IMSI value. A user ID request is first sent from the server to the client. The client's response is a tuple $(\text{TMSI}, \text{LAI})$, consisting of a temporary identifier and the local area identifier in which C received TMSI. If the LAI value corresponds to the LAI of S , then the latter searches for a tuple $(\text{TMSI}, \text{IMSI})$ in its own database; else S requests this tuple from the server S^* associated with LAI, over an unspecified channel. If no IMSI can be found, then the server demands the IMSI in clear. This procedure is the essential vulnerability leading to IMSI catcher attacks 3.3.

The second phase is run only optionally by the server S and the client's operator Op over a *secure channel*; its purpose is to enable S to then run a batch of AKE sessions with C . The server sends the client's IMSI to Op , which generates a batch of vectors AV each providing material for one out of a maximum of n sessions. For each vector, the operator's state with respect to the client $Sqn_{S,C}$ is augmented, and then the following values are generated: a fresh random value R ; an server-authentication value Mac_S (for the values $Sqn_{S,C}$ and R); a client-authentication value Mac_C (for R only); the session keys CK and IK ; and an anonymity key AK . Of these six values, the last five are computed by using each time a different cryptographic algorithm, denoted $\mathcal{F}_1, \dots, \mathcal{F}_5$. In fact, the AKA protocol uses seven such algorithms, but two of them, denoted $\mathcal{F}_1^*, \mathcal{F}_5^*$, are only used in the *re-synchronization procedure*. The seven algorithms are generic, and can currently be instantiated in one of two ways, one using AES (called MILENAGE), the other using Keccak (called TUAK). Both \mathcal{F}_1 and \mathcal{F}_1^* take as input the keys sk_C and sk_{Op} , the random value R , and a sequence number $Sqn_{S,C}$. The other algorithms use the secret keys and the random value, but not the sequence number. At the end of this phase, the following values are sent to the server for each of the n sessions: $AV = (R, CK, IK, Mac_S, Mac_C, AMF, AK \oplus Sqn_{S,C})$, in which AMF is the Authentication Management Field is a 16-bit value used only in radio access specifications (for example E-UTRAN or non-3GPP access to EPS).

The sequence number $Sqn_{S,C}$ is notably *not* sent in clear to the server, but rather blinded by the value AK .

The third phase of the protocol, *authenticated-key-exchange*, is a mutual authentication and key agreement between the server and the client over the insecure channel. The server chooses the next vector available (if phase 2 was run, then this is the first tuple; else, for returning clients phase 3 is run directly, with the next authentication vector), and sends an *authentication challenge* consisting of the random value R and an authentication string $Autn = (Sqn_{S,C} \oplus AK) \parallel AMF \parallel Mac_S$. The client uses R to compute AK , then it recovers $Sqn_{S,C}$ and verifies Mac_S . If the verification succeeds, and if the recovered Sqn is within a distance of Δ (a pre-defined constant) of the client's own state Sqn_C , then the client computes CK, IK , and the response Mac_C , sending this latter value to S ; else, if the two sequence numbers are too far apart, then the client forces a *re-synchronization procedure*, which is the fourth phase of the protocol. If no re-synchronization is needed, then the client updates $Sqn_C := Sqn_{Op,C}$, and S verifies the received authentication value with respect to the Mac_C sent by Op . If the verification succeeds, then the server sends an acknowledgement to Op and goes directly to phase five. If the verification fails, then the protocol is aborted.

The fourth phase of the protocol, *resynchronization*, is run by all three parties. The client essentially retraces the operator's steps, using its own sequence number Sqn_C and computing the values Mac_S^* and $AK^* \oplus Sqn_C$ by using algorithms \mathcal{F}_1^* and \mathcal{F}_5^* (rather than \mathcal{F}_1 and \mathcal{F}_5), but keeping the same random value R . If the authentication string Mac_S^* verifies for the Sqn value Op recovers, then Op resets its sequence number to Sqn_C and sends to the server another batch of authentication sessions. The protocol restarts. We note that this phase is executed only optionally.

Finally, the fifth phase of AKA, *TMSI re-allocation*, is run by the server and client. As the first message of the record layer, the server sends an (unauthenticated) encryption of a new TMSI value to the client C , using the session key CK computed in phases 3 or 4. The encryption is done by means of the A5/3 algorithm detailed in TS 43.020 [4], run in cipher mode. The new TMSI value, called $TMSI_{new}$, is only permanently saved by the server if the client acknowledges the receipt; else, both values $TMSI_{new}$ and the old $TMSI_{old}$ are retained and can be used in the next authentication procedure. In appendix, the entire protocol is depicted in Figure 11.

3.3 Some Privacy breaches in AKA

To the best of our knowledge, the specifications do not detail how the permanent user identifiers are generated, we consider that all the user IDs are different.

We do not need to have a very formal/theoric analysis to point out the pure AKA protocol cannot guarantee the user unlinkability as defined previously. Indeed, the user identification based on temporary identifiers independent to the permanent identifier is not sufficient. As related to our introduction, we clearly analyze the two main issues which restrict the AKA protocol to guarantee the user unlinkability. The first one is focused on the linkability of the failure messages. The next one is focused on the operational difficulties related to the TMSI. These weaknesses have been partially described by Arapinis and al. in [20, 21].

3.3.1 Linkability of failure messages

Arapinis and al [21] notably provides a novel practical attack establishing the traceability of a user based on the study of the failure messages. Considering a victim mobile client C_v , an adversary can detect its presence in a specific area, only considering the features of the failure messages, and replaying one old legitimate authentication vector including $(R, Autn)$.

This latter can be replayed to any client each time it wants to check the presence or not of the victim client C_v . In fact

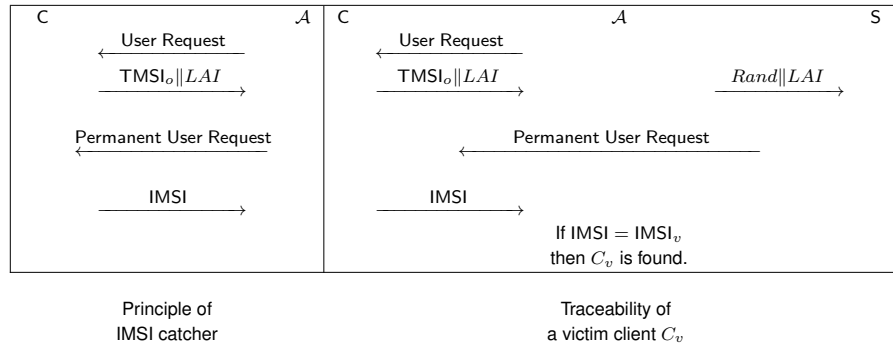


Fig. 3. Attacks based on TMSI.

this replay always implies a failure message. If this replayed authentication vector had been sent to C_v , the answer will be a "de-synchronization message". Indeed, the received MAC will be successfully verified, but the sequence number will be not in the correct range. Otherwise, the answer will be an "Mac failure message", because the the Mac could not be correct, except to a negligible probability of success ($1/2^{(128)}$) (due to the collision's probability). This difference can permit to trace a mobile client in a specific area just replaying a legitimate authentication vector. That represents a breach of the user's untraceability.

3.3.2 Operational difficulties with TMSI

As we explained previously, 3GPP has standardized the use of temporary values TMSI instead of the basic permanent identifier IMSI. These temporary values are generated independently to the permanent identifier. At first glance, this option could guarantee at least the user identity confidentiality. However, the permanent identity IMSI is not as private as the sequence number. In fact, this value can be easily obtained by a weak active adversary. In the identification of the user, when the VLR cannot recognize the TMSI, the procedure is not aborted but the permanent identifier of the user is requested to the client. Thus, by a basic attack 3 an adversary can obtain the permanent identifier of a user during a fresh session. Obviously that represents a breach in the confidentiality of the user, that can also imply a breach in the intraceability. This weakness is exploited by the well-known "IMSI catcher" [10], which is the best known attack to mobile telephony users' privacy. This attack consists in forcing a mobile phone to reveal its permanent identity.

The Temporary identifier of each user must be used, as its name suggests, temporary. The specifications consider the uniqueness of the TMSI's use. As specified by the protocol, a TMSI re-allocation is provided. However, a basic active ad-

versary can corrupt this uniqueness. Indeed, we note that the VLR does not de-allocate the old value TMSI_o without receiving the acknowledge message. So if an adversary drops this latter, the VLR has the both allocated temporary values: the old one TMSI_o and the new one TMSI_n.

For the next identification, the VLR will accept the both values to identify the user. So an adversary dropping the acknowledge message can easily force the replay of the same TMSI during different sessions. Moreover, we note that the new temporary identifier TMSI_o is only sent encrypted with the session key CK and without any integrity service. The mobile subscriber can not verify whether the received value is sent by the VLR or any malicious entity. Thus, any adversary can impose a false random value as new TMSI sending a random value instead of the ciphered temporary identifier.

As specified previously, each TMSI is associated to a specific location area denoting the VLR having generated this temporary identifier. The related LAI is sent in cleartext with the TMSI. Due to the previous detailed failure, an adversary can associated the IMSI with an LAI: that permits to know where the subscriber with the permanent identifier IMSI was located during the previous session.

Therefore, the use of temporary identifiers as specified by the specifications can not assure any privacy as specified.

We proceed to describe two of the more promising improvements to the AKA protocol, and show that these variants are vulnerable to client-unlinkability attacks.

The Arapinis variant. Arapinis et al. [21] propose an AKA variant which is supposed to ensure client unlinkability by avoiding failure-message-based linking, as described in the previous section. To avoid this attack, they propose to replace the both failure messages by two indistinguishable messages. The failure message is now encrypted with a public key of the network and includes the IMSI, a constant Fail, a random value R and the current sequence number Sqn_C. This latter is specially encrypted with an unlinkability key $UK = f_{sk_C}(R)$

in order to authenticate the error message. After receiving this generic error message, the network can deduce the cause of the failure from the IMSI and the sequence number, and sends the appropriate answer.

Moreover, they propose a fixed version of the identification protocol avoiding to expose the IMSI (basically the IMSI is sent in cleartext upon request by the network). It breaches both user identity confidentiality and untraceability. Thus, the IMSI is encrypted with an asymmetric-randomized-encryption (the same one as this one of the error messages) and they do not use any temporary value as usual. That permits to cancel the reallocation step. We suppose that if the VLR cannot recover the IMSI, the protocol is aborted (and they do not use a Permanent Identity Request). We note that this variant owns different practical issues. Indeed, authors propose a simply representation of the network by concatenating the VLR and HLR as an unique entity. For the security analysis, it is not an issue but it is the case for practical considerations. Contrary to their formal security proof, we consider that the fixed version from Arapinis does not assure the untraceability of the mobile subscriber. Indeed, we propose a new attack permitting to trace a victim client based on the knowledge of a user permanent identity. Its knowledge is reasonable since we can not consider this data as private. In a specific area, an adversary can trace a victim client C_v , which has its permanent identifier $IMSI_v$ known by the adversary. This attack consists to replace the answer of the user identity request by a response with the permanent identifier of the victim client and analyze the behavior of the session. If the authentication data answer Res contains a failure message, the attacker considers the tested client as different from the victim client. Otherwise, the victim client has been found. This attack is detailed in the figure 4. We state the following result:

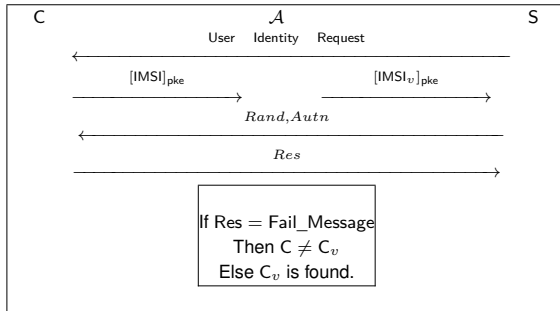


Fig. 4. Attack on the fixed version of Arapinis.

Lemma 1. *Let Π the protocol proposed by Arapinis and al. in [21]. Consider a $(t, 1, 1, 0)$ -adversary \mathcal{A} against the weak-privacy of the protocol Π running in time t , creating at most*

one party instance, running one user identification per each instance and making no extra query to the related internal cryptographic functions. The advantage of a such adversary is denoted $\text{Adv}_{\Pi}^{\text{ww-unlink}}(\mathcal{A})$. The protocol Π cannot guarantee the weak-privacy since there exists a such adversary \mathcal{A} with a no-negligeable advantage.

Proof. The attack we describe goes as follows:

1. At time clock = 0, the challenger sets up the server.
2. At time clock = 1, the adversary creates two clients C_0, C_1 from the oracles $\text{CreateClient}(IMSI^{\{0\}}, st_0^{\{0\}})$ and $\text{CreateClient}(IMSI^{\{1\}}, st_0^{\{1\}})$.
3. At time clock = 2 \rightarrow 3, the adversary uses the oracle $\text{DrawClient}(C_0, C_1)$ which returns at time clock = 3 a virtual client $vc = C_b$ following a chosen bit b .
4. At time clock = 4 \rightarrow 10, the adversary runs an instance of the protocol Π between the virtual client and the terminal as follows: At time clock = 4, the terminal sends a user identity request to the virtual client; at time clock = 5, the virtual client answers with the value UID such as: $UID = [IMSI^{\{b\}}]_{pke_b}$; at time clock = 6, the adversary intercepts this message (before arriving to the terminal) and forges a user identification $UID = [IMSI^{\{0\}}]_{pke_0}$ for the client C_0 which will be sent to the terminal at time clock = 7 (instead of the user identifier computed by the virtual client); at time clock = 8, the terminal receives the related user identifier and generates an authentication challenge which is sent to the virtual client vc at time clock = 9; this latter sends its authentication answer Res at time clock = 10 which is eavesdropped by the adversary.
5. At time clock = 11, since the eavesdropped authentication answer Res , the adversary tries to guess the bit b as follows: if the message Res contains a failure message, then the adversary considers the virtual client as the client C_0 (that implies $b' = 0$). Otherwise, it considers the virtual client as the client C_1 (that implies $b' = 1$).
6. At time clock = 12, the adversary returns its guess b' .

This attack is detailed in the figure 4. It is clear that the advantage of a such adversary is 1. Thus, the protocol cannot guarantee the weak-privacy. We note that in practice if the adversary \mathcal{A} cannot detect the de-synchronization, some "false positive" can be appear: indeed, during the tested session, if the victim client is de-synchronized, the adversary could not trace it. But the frequency of a such event (i.e this one of a synchronization procedure) is very weak. Moreover, we can use this attack procedure at least twice to reduce the probability to obtain a false positive. \square

The van den Broek variant. Van den Broek et al. [11] recently proposed an IMSI catcher countermeasure; in this improved variant, avoid sending the IMSI in clear by replacing (IMSI, TMSI) tuples by an upgradeable pseudonym denoted PMSI. Their modified identification phase is exclusively done by means of these pseudonyms. The PMSI is chosen by the operator and sent with the authentication challenge in the preparation phase, encrypted together with the sequence number with a new secret key that is assumed to be shared by clients and their operators. The ciphertext is used as the random value R in the authentication challenge. Indeed, a successful session of the AKA protocol, ending in the establishment of new session keys, can only be attained if the PMSI is correctly updated. This variant is described in detail in [11].

From a practical point of view, using the operator at each key-exchange session is costly, and something that the original AKA design tries to avoid. Furthermore, though this variant successfully prevents IMSI catchers, it does not address client unlinkability. The pseudonym PMSI can be intercepted in one session; if this session is then aborted, the PMSI can be replayed in a second session, thus leading to user linkability. Furthermore, the protocol is vulnerable to the attack based on linking failure messages, as presented by Arapinis et al. Thus, if Π denotes the protocol proposed by van den Broek et al., it holds that:

Lemma 2. *There exists a $(t, 2, 1, 0)$ -adversary \mathcal{A} against the narrow-weak-client-unlinkability of Π running in time t , initiating two protocol sessions, and making no query to the internal cryptographic function G , which has an advantage $\text{Adv}_{\Pi}^{\text{ww-unlink}}(\mathcal{A}) = \frac{1}{2}$ (and a probability of 1) to win the game.*

Proof. The attack we describe goes as the Arapinis attack based on the linkability of failure messages:

1. At time clock = 0, the challenger sets up the server.
2. At time clock = 1, the adversary creates two clients C_0, C_1 from the oracles $\text{CreateClient}(\text{IMSI}^{\{0\}}, \text{st}_0^{\{0\}})$ and $\text{CreateClient}(\text{IMSI}^{\{1\}}, \text{st}_0^{\{1\}})$.
3. At time clock = 2 \rightarrow 3, the adversary uses the oracle $\text{DrawClient}(C_0, C_1)$ which returns at time clock = 3 a virtual client $vc = C_b$ following a chosen bit b .
4. At time clock = 4 \rightarrow 5, the adversary runs an instance of the protocol Π between the virtual client and the terminal via the oracle $\text{Execute}(vc, s)$ which returns the transcript of an instance s at time clock = 5. This transcript includes the authentication challenge denoted (R, Autn) .
5. At time clock = 6 \rightarrow 12, the adversary runs an instance $s + 1$ of the protocol Π between the virtual client and the terminal as follows: at time clock = 6, the terminal sends a user identity request to the virtual client; at time clock = 7, the virtual client sends the user identifier UID

to the terminal; at time clock = 8, the terminal receives the related user identifier and generates an authentication challenge which is sent to the virtual client vc ; at time clock = 9, the adversary intercepts the fresh authentication challenge (before arriving to the virtual client) and sends the previous authentication challenge (R, Autn) to the virtual client at clock = 10 (obtains in a previous session on the client $C_0 := \text{DrawClient}(C_0, C_0)$); at time clock = 11, the virtual client sends its authentication answer Res at clock = 12 which is eavesdropped by the adversary.

6. At time clock = 13, since the eavesdropped authentication answer Res , the adversary tries to guess the bit b as follows: if the message Res contains a "desynchronization message", then the adversary considers the virtual client as the client C_0 (that implies $b' = 0$). Otherwise, it considers the virtual client as the client C_1 (that implies $b' = 1$).
7. At time clock = 14, the adversary returns its guess b' .

□

4 Our proposal: PrivAKA

In this section, we propose a new fixed variant of the AKA protocol which respects the weak-unlinkability and key-indistinguishability as defined in our privacy model 2. Obviously, it also guarantees the client-, server-impersonation. Moreover, our fixed variant takes into account practical requirements of the VLR in addition to provide the security properties. The VLR is an interoperable server in a cellular network that supports roaming functions for subscribers outside the coverage area of their own HLR. The VLR needs to manage the generation of the temporary identifier to avoid collision between subscriber from different operators. Indeed, if the HLR generates the new temporary identifiers, it cannot avoid collision between identifiers of user from different HLR but only avoids collision between its own subscribers. Then, the VLR cannot store any secret data shared with any user due to its interoperability feature, except the storage of the relation between the permanent and temporary identifiers for each subscriber. Indeed, the VLR communicates with user from any operator. So, this feature implies that the secret data also needs to be inter-operable, i.e a same secret data common for all the operators. We note that our fixed protocol does not increase the number of exchanged between the USIM and ME and respect the 3GPP's desire to limit communication with the HLR. Moreover, we are the first to seriously consider such practical requirements.

4.1 Description of our variant

Instead of five phases, our variant only consists of three. Our protocol is designed to not require re-synchronization (which was phase 4 in AKA), and we include the TMSI reallocation (which was phase 5 in AKA) as part of the key-exchange phase (phase 3). In our construction, we use a public-key encryption scheme $\text{PKE} = (\text{PKE.KGen}, \text{PKE.Enc}, \text{PKE.Dec})$, such that each operator has a certified public and secret key-pair denoted as $(\text{pke}_{\text{Op}}, \text{ske}_{\text{Op}})$. We assume that the client stores only its own operator's public key (and its certificate) internally. In particular, we do not give encryption keys to the servers in order to minimize key-management issues. We also use a secure authenticated encryption scheme $\text{AE} = (\text{AE.KGen}, \text{AE.Enc}, \text{AE.Dec})$. Though these can be instantiated generically, we use **AES-GCM** [9] for the AE scheme and (EC)IES [27] for the PKE scheme. We depict our variant in Figure 5, and indicate in the grey boxes the differences to the classical AKA procedure. Just as the original protocol, our variant starts with an *identification* phase, run by the client and the server over the insecure channel. The server sends an identification request which includes a random value that we denote R_{id} . The client forges an user identification answer following a flag $\text{flag}_{\text{TMSI}}$ managed by the client: it uses either a pre-exchanged fresh temporary identifier TMSI (if the previous protocol has been accepted and if the user stayed in the same area: $\text{flag}_{\text{TMSI}} = 1$) or it sends a public-key encryption of the concatenation of the received random value R_{id} and the evaluation of the function \mathcal{F}_5 on input the client secret key, the operator's secret key, the random value R_{id} and the client's IMSI (only if it does not own a fresh temporary identifier or if the user is located in a new area: $\text{flag}_{\text{TMSI}} = 0$). We demand that the output size of the PKE scheme for this input size is always equal to the length of the TMSI. A such flag $\text{flag}_{\text{TMSI}}$ is only used to restrict the computation of an encrypted permanent identity and not for security reasons. In both cases, the client also appends the identity of the operator Op it subscribes to, to the message. Moreover, we assume that the client can detect when it moves in a new area. In this case, it allocates the flag at one.

Intuitively, if the client stayed in the same area ($\text{flag}_{\text{TMSI}} = 0$), then the TMSI the client stored came from the server it is currently communicating with, hence the server can find the (TMSI, IMSI) association in its database. Otherwise, the client does not use its TMSI value, but rather encrypts a function of the IMSI with the operator's public key. The function is symmetric-key, requiring knowledge of the client- and operator keys, and it is not replayable due to the fresh identification randomness R_{id} . Upon receiving a string of the form (m, Op) , the server first checks whether the message m is a TMSI present in its database; if so, it retrieves the IMSI to

which this value corresponds; else, it assumes that m is a ciphertext, and it sends it to the operator Op for decryption.

Phase 2, *preparation*, is run over a secure channel, between the server and the operator. If the server received a valid TMSI in the previous phase, the preparation phase being with the server sending the corresponding IMSI to the operator; else, the server forwards the received ciphertext and the associated random value R_{id} . The operator proceeds similarly to the standard AKA preparation procedure, with the following differences:

- We add as input to each cryptographic function a server-specific value Res_S for a server with identity $S = \text{VLR}$. This is to prevent attacks in which the adversary replays authentication vectors from one network to another, as presented by Zhang [18]. We also use the constant AMF which is sent in clear, as an additional input.
- We add the sequence number $\text{Sqn}_{\text{Op,C}}$ to each of the cryptographic functions apart from \mathcal{F}_5 . Since the sequence number is an ephemeral value, which is updated, this guarantees freshness even if the randomness R is repeated.
- We introduce an index value $\text{id}_{\text{Op,C}}$ which essentially prevents the repetition of a challenge using the same sequence number. This value is essential in preventing a desynchronization of sequence number values. We note that the client also keeps track of a similar index id_C , which will play a role in the key-exchange phase, as detailed below.

In the final phase, *authenticated-key-exchange*, which is run between the client and the server, the server sends a random value R , the authentication string Autn , an authenticated encryption of the new TMSI, using the keys (CK, IK) derived for this session. The client proceeds similarly to the original AKA procedure, recovering AK by using the random value R , then checking Mac_S . If the authentication cannot be verified, the procedure is aborted, and the new TMSI is disregarded. Else, the user computes CK, IK and decrypts the received authenticated encryption string to find the TMSI value and the operator's index $\text{id}_{\text{Op,C}}$. Then, C checks the freshness of the sequence number, i.e. it verifies if one of the following two conditions is correct:

- $\text{Sqn}_C = \text{Sqn}^{\{i\}}$.
- $\text{Sqn}_C = \text{inc}(\text{Sqn}^{\{i\}})$ and $\text{id}_{\text{Op,C}} = \text{id}_C + 1$,

If the protocol is run normally, the first of these conditions is the one that will hold. However, if the previous session is aborted after receiving the server's authentication challenge, then the two sequence numbers may become desynchronized by one step (the second condition). Further desynchronization is prevented by the use of an index, which indicate whether




Fig. 5. Our fixed AKA Procedure.

the authentication string for a particular $Sqn_{Op,C}$ has already been used or not. If the first condition holds, then the client's internal index is reset; else, the index is incremented by 1. The client updates the sequence number only upon success-

ful authentication. If none of these conditions are verified, the procedure is aborted and does not use of a resynchronization procedure.

Finally, the user computes a response $\text{Res} := \mathcal{F}_2(\text{keys}, R^{\{i\}}, \text{Sqn}^{\{i\}}, \text{Res}_S, \text{AMF})$, sends Res , then stores the TMSI and the new index value. The server checks Res against the prepared value Mac_C (else, if no response is received, the procedure is aborted).

One notable exception to the original AKA protocol is that whenever an abort occurs on the server's side, the second phase – *preparation* – is used instead of simply querying the next vector in the prepared batch. Though this might seem more inefficient, we note that an abort only occurs in the presence of an adversary, which is considered to be a rare event. We detail the procedure upon aborts in Figure 6.

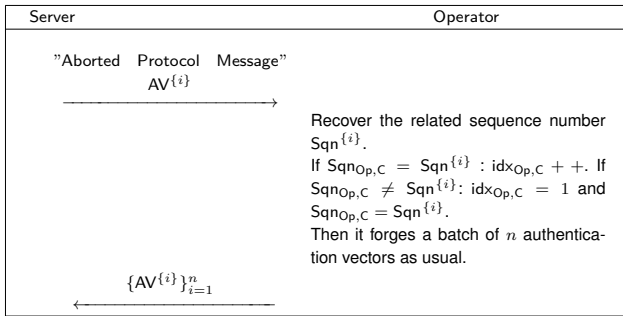


Fig. 6. Procedure after an abort.

Internal cryptographic algorithms:

In our variant, we have modified the inputs of the internal cryptographic algorithms, notably to include the sequence number and the new value Res_S . Thus, we need to provide an update of these algorithms to consider these modifications. As specified in specifications, the AKA protocol can be based on the different sets of algorithms: TUAK and MILENAGE.

To preserve backwards compatibility, we propose to keep and update these both sets. Moreover, our variant requires an algorithm of authenticated encryption. We require to use the well-know standard **AES-GCM** [9]. It is denoted **AE**.

Considering our modifications, we do not use the function \mathcal{F}_1^* and \mathcal{F}_5^* since we have dropped the resynchronization.

The seven internal cryptographic functions takes in inputs the following values:

- keys: the couple of the both 128-bit (or 256-bit) keys: the subscriber key sk and the operator key sk_{op} .
- Sqn (except for the function \mathcal{F}_5): a 48-bit sequence number.
- AMF (except for the function \mathcal{F}_5): a 16-bit authentication field management.
- R : a 128-bit random value.

- Res_S : a 128-bit (public?) value characterizing the visited network.

We note that the function \mathcal{F}_5 behaves differently because they do not consider the sequence number in inputs (contrary to the currently version, where \mathcal{F}_1 and \mathcal{F}_1^* behave differently).

Update of the MILENAGE algorithms: MILENAGE is the original set of algorithms which is currently implemented as detailed the specification 35.206 [1].

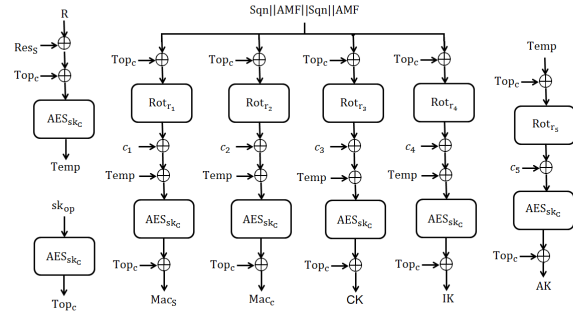


Fig. 7. Updated Milenage.

To assure a stronger security, we also modify the MILENAGE algorithms to output 128-bit MAC and session keys CK and IK .

Based on the Advanced Encryption Standard (AES), these functions compute firstly the both values Top_C and Temp as follows:

$$\text{Top}_C = sk_{op} \oplus \text{AES}_{sk_C}(sk_{op}),$$

$$\text{Temp} = \text{AES}_{sk_C}(R \oplus \text{Top}_C \oplus \text{Res}_S).$$

Then, we obtain the output of the seven functions as follows:

- **Output \mathcal{F}_1 :** $\text{Mac}_S = \text{AES}_{sk_C}(\text{Temp} \oplus \text{Rot}_{r_1}(\text{Sqn}||\text{AMF}||\text{Sqn}||\text{AMF}) \oplus c_1) \oplus \text{Top}_C$,
- **Output \mathcal{F}_2 :** $\text{Mac}_C = \text{AES}_{sk_C}(\text{Temp} \oplus \text{Rot}_{r_2}(\text{Sqn}||\text{AMF}||\text{Sqn}||\text{AMF}) \oplus c_2) \oplus \text{Top}_C$,
- **Output \mathcal{F}_3 :** $\text{CK} = \text{AES}_{sk_C}(\text{Temp} \oplus \text{Rot}_{r_3}(\text{Sqn}||\text{AMF}||\text{Sqn}||\text{AMF}) \oplus c_3) \oplus \text{Top}_C$,
- **Output \mathcal{F}_4 :** $\text{IK} = \text{AES}_{sk_C}(\text{Temp} \oplus \text{Rot}_{r_4}(\text{Sqn}||\text{AMF}||\text{Sqn}||\text{AMF}) \oplus c_4) \oplus \text{Top}_C$,
- **Output \mathcal{F}_5 :** $\text{AK} = \lfloor \text{AES}_{sk_C}(\text{Rot}_{r_5}(\text{Temp} \oplus \text{Top}_C) \oplus c_5) \oplus \text{Top}_C \rfloor_{0..47}$,

with the five integers $r_1 = 0, r_2 = 24, r_3 = 48, r_4 = 64$ and $r_5 = 96$ in the range $\{0, 127\}$, which define the number of

positions the intermediate variables are cyclically rotated by the right, and the five 128-bit constants c_i such as:

- $c_1[i] = 0, \forall i \in \{0, 127\}$.
- $c_2[i] = 0, \forall i \in \{0, 127\}$, except that $c_2[127] = 1$.
- $c_3[i] = 0, \forall i \in \{0, 127\}$, except that $c_3[126] = 1$.
- $c_4[i] = 0, \forall i \in \{0, 127\}$, except that $c_4[125] = 1$.
- $c_5[i] = 0, \forall i \in \{0, 127\}$, except that $c_5[124] = 1$.

This is also described in the Figure 7.

Update of the TUAK algorithms: TUAK is an alternative set of MILENAGE based on the internal permutation of the Keccak [12]. The specification TS 35.231 [2] details the internal algorithms of this set. We update these algorithms by only modifying the inputs of the second permutation.

As the no-updated TUAK algorithms, we compute firstly the value Top_C as follows:

$$\text{Top}_C = [f_{\text{Keccak}}(\text{sk}_{\text{op}} \parallel \text{Inst} \parallel \text{AN} \parallel 0^{192} \parallel \text{Key} \parallel \text{Pad} \parallel 1 \parallel 0^{512})]_{1..256},$$

We note that the values AN, Inst', Inst, Pad are the same as used in the no-updated TUAK algorithms and Key the (padded) subscriber key.

At this point, the behavior of the functions \mathcal{F}_5 diverges from that of the other functions. To generate the related output, we compute the value Val_1 and for the others ones, we compute the value Val_2 which differ as follows:

$$\begin{aligned} \text{Val}_1 &= f_{\text{Keccak}}(\text{Top}_C \parallel \text{Inst}' \parallel \text{AN} \parallel \text{R} \parallel 0^{64} \parallel \text{Key} \parallel \text{Res}_5 \parallel \text{Pad} \parallel 10^{512}), \\ \text{Val}_2 &= f_{\text{Keccak}}(\text{Top}_C \parallel \text{Inst}' \parallel \text{AN} \parallel \text{R} \parallel \text{AMF} \parallel \text{SqN} \parallel \text{Key} \parallel \text{Res}_5 \parallel \text{Pad} \parallel 10^{512}). \end{aligned}$$

Then, we obtain the output of the seven functions truncating the related value as follows:

- **Output \mathcal{F}_1 :** $\text{Mac}_5 = [\text{Val}_2]_{0..127}$,
- **Output \mathcal{F}_2 :** $\text{Mac}_C = [\text{Val}_2]_{256..383}$,
- **Output \mathcal{F}_3 :** $\text{CK} = [\text{Val}_2]_{512..639}$,
- **Output \mathcal{F}_4 :** $\text{IK} = [\text{Val}_2]_{768..895}$,
- **Output \mathcal{F}_5 :** $\text{AK} = [\text{Val}_1]_{0..47}$.

This is also depicted in Figure 8.




Fig. 8. Updated TUAK

We note that the multi-output property is, as the no-updated version, not an issue for the security of the master key, since during one session we can have as many as four calls to the same function with similar inputs (and a different truncation).

4.2 Privacy and Security Analysis

4.2.1 Weak Client Unlinkability of our fixed variant

In this section, we prove that our fixed variant of the AKA protocol achieves weak client unlinkability if we assume that the internal functions assures indistinguishability, pseudorandomness and unforgeability properties. The weak-client-unlinkability resistance of our fixed variant is proved as follows:

Theorem 1. [ww-unlink – Resistance.]

Let $G: \{0, 1\}^\kappa * \{0, 1\}^d * \{0, 1\}^t * \{0, 1\}^t \rightarrow \{0, 1\}^n$ be our unitary function described in section 4.2.1 and Π our fixed variant of the AKA protocol specified in section 4. Consider a $(t, q_{\text{exec}}, q_{\text{id}}, q_G, q_{\text{AE}}, q_{\text{PKE}})$ -adversary \mathcal{A} against the weak privacy ww-unlink-security of the protocol Π running in time t , creating at most q_{exec} party instance, with at most q_{id} user identification per instance and making at most $q_G, q_{\text{AE}}, q_{\text{PKE}}$ queries to respectively the functions G, AE and PKE . The advantage of this adversary is denoted $\text{Adv}_{\Pi}^{\text{ww-unlink}}(\mathcal{A})$. Then, there exist $(t' \sim O(t), q' = q_{\text{exec}} + q_G)$ -adversary \mathcal{A}_1 and $(t' \sim O(t), q' = 2 \cdot q_{\text{exec}} + q_G)$ - \mathcal{A}_2 against respectively the unforgeability and pseudorandomness of the function G , an $(t'' \sim O(t), q'' = q_{\text{exec}} + q_{\text{AE}})$ -adversary \mathcal{A}_3 against the ae-security of the function AE , $(t''' \sim O(t), q''' = q_{\text{exec}} \cdot q_{\text{id}} + q_{\text{PKE}})$ -adversaries \mathcal{A}_4 against the indistinguishability of the function PKE , $(t, q_{\text{exec}}, q_{\text{id}}, 0, 0, q_G, q_{\text{AE}}, q_{\text{PKE}})$ -adversary \mathcal{A}_5 against the key-secrecy of the protocol Π such that:

$$\begin{aligned} \text{Adv}_{\Pi}^{\text{ww-unlink}}(\mathcal{A}) &\leq \text{Adv}_{\Pi}^{\text{K.Ind}}(\mathcal{A}_5) + \frac{1 + (q_{\text{exec}} \cdot q_{\text{id}})^2}{2^{|\text{TMSI}|}} + \\ &\frac{q_{\text{exec}}^2}{2^{|\text{R}|}} + \frac{(q_{\text{exec}} \cdot q_{\text{id}})^2}{2^{|\text{R}_{\text{id}}|}} + n_C \cdot (3 \cdot \text{Adv}_G^{\text{mac}}(\mathcal{A}_1) + \\ &\text{Adv}_G^{\text{prf}}(\mathcal{A}_2) + 4 \cdot \text{Adv}_{\text{AE}}^{\text{ae}}(\mathcal{A}_3) + \text{Adv}_{\text{PKE}}^{\text{ind-cca2}}(\mathcal{A}_4)). \end{aligned}$$

Proof. Our proof has the following hops:

Game \mathbb{G}_0 : This game works as the ww-unlink-game stipulated in section 2.

Game \mathbb{G}_1 : We modify the original game \mathbb{G}_0 to ensure that the random values R_{id} and R used by honest server instances are always unique. The related security loss due to the collisions between each respective random in two different instances is given by the following expression:

$$\left| \Pr[\mathcal{A}_{\mathbb{G}_0} \text{ wins}] - \Pr[\mathcal{A}_{\mathbb{G}_1} \text{ wins}] \right| \leq \frac{q_{\text{exec}}^2}{2^{|\text{R}|}} + \frac{(q_{\text{exec}} \cdot q_{\text{id}})^2}{2^{|\text{R}_{\text{id}}|}}.$$

Game \mathbb{G}_2 : The game \mathbb{G}_2 behaves as the game \mathbb{G}_1 with the restriction consisting to abort the protocol if one of the three following events happen:

- **Event 1:** The adversary \mathcal{A} has forged the user identification answer.
- **Event 2:** The adversary \mathcal{A} has forged the authentication challenge.
- **Event 3:** The adversary \mathcal{A} has forged the authentication response.

Briefly, we modify the original game \mathbb{G}_1 to ensure that the adversary cannot forge these three messages. Indeed, a classic attack permits to recover the user identity of the client from the ability to forge any of these three messages. A such adversary behaves as follows: the adversary chooses one of the both clients (C_0 or C_1) related to the virtual client VC and forges a message x related to the current transcript. During the session, it drops the related message from the honest server and sends the forged message. If with a such modification the session is accepted, then the chosen client is the client to this current session. Otherwise, the other client is the good one.

Now, focus on each event. Firstly, if the adversary wants to forge a fresh user identification response (**event 1**), she has two options: either guess the fresh temporary identifier TMSI or an encrypted version of the permanent identifier related to the random value R_{id} included in the user identification request. To guess the fresh temporary value, either it guess the fresh one from the old previous ones or it decrypts the authenticated encryption of the fresh one from the previous session. The TMSIs are independently chosen, so they are no way to guess the fresh TMSI from the old one, except randomly, or replaying one of these previous TMSI. The probability of a such success is at least $(1 + (q_{exec} \cdot q_{id})^2) / 2^{|\text{TMSI}|}$. The ability to recover the fresh TMSI from the encrypted version is restricted by the security of the authenticated encryption and the key-indistinguishability of the session keys. Indeed, the fresh TMSI is sent authenticated and encrypted by the algorithm AE, which is based on session keys and not long-term keys. Moreover, we note that the fresh TMSI is encrypted with the index which is predictable. So, if an adversary can forge a correct output of the authenticated encryption AE (i.e the related input includes a correct index), it can impose its own temporary value and uses it for the next user identification. So the probability of a such recovery is at most $\text{Adv}_{\Pi}^{\text{K,Ind}}(\mathcal{A}) + n_C \cdot (2 \cdot \text{Adv}_{\text{AE}}^{\text{ae}}(\mathcal{A}))$. Then, the ability to forge a correct encrypted permanent identity is restricted by the ability to forge a correct output of the function G with the IMSI and without the private key. The success probability of a such ability is $n_C \cdot \text{Adv}_G^{\text{mac}}(\mathcal{A})$. So the security loss due to the ability to forge an fresh identification response is $\frac{1 + (q_{exec} \cdot q_{id})^2}{2^{|\text{TMSI}|}} + \text{Adv}_{\Pi}^{\text{K,Ind}}(\mathcal{A}) + n_C \cdot (2 \cdot \text{Adv}_{\text{AE}}^{\text{ae}}(\mathcal{A}) + \text{Adv}_G^{\text{mac}}(\mathcal{A}))$.

Now, focus on the second event involving to forge an authentication challenge. We recall that a such challenge is split in four parts: the random value R , a masked value of the cur-

rent sequence number $val = AK \oplus \text{Sqn}$, a message authentication code mac generated by the function G which takes in inputs the private keys $keys$, the random value R and the sequence number Sqn , and the couple of the next temporary identifier and the current index encrypted by an authenticated encryption AE. In a fresh instance, we have two options to forge an authentication challenge: either the adversary guesses a fresh authentication challenge based on the current sequence number or it replays an old challenge based on a previous used sequence number. For the first option, the complexity to forge a such challenge is restricted by the unforgeability of the function G . Indeed even if the adversary knows the fresh sequence number, which is in practice masked by one-time-pad, it cannot forge a fresh message authentication code without the related message authentication code. Moreover, we note that the index is only implied for the second condition of the freshness verification and the new TMSI will be only used for the next session. With the second condition, the best option for the adversary consists to replay the three first parts of the previous authentication challenge included in an aborted session and tries to forge an authenticated and encrypted version of fresh index and fresh TMSI. Indeed, when a protocol is aborted after that the server has sent the authentication challenge, the next authentication challenge is based on the same sequence number. To forge a fresh authenticated and encrypted value, the adversary chooses any non-nil value for the index value. We note that the adversary needs to know it because a such challenge including the index idx is considered as fresh only if there are exactly idx previous sessions aborted by a drop of the related authentication response. Moreover, we do not need to the chosen temporary value. The ability to forge a such challenge is restricted by the security of the chosen algorithm of authenticated encryption. So, considering both conditions, the success probability is at most $n_C \cdot (\text{Adv}_G^{\text{mac}}(\mathcal{A}) + \text{Adv}_{\text{AE}}^{\text{ae}}(\mathcal{A}))$.

Finally, we focus on the third event involving to forge an authentication response. This value is only composed by the value Res , which is an output of the function G . So the success probability is at most $n_C \cdot \text{Adv}_G^{\text{mac}}(\mathcal{A})$. Thus, we obtain

$$\begin{aligned} & \left| \Pr[\mathcal{A}_{\mathbb{G}_1} \text{ wins}] - \Pr[\mathcal{A}_{\mathbb{G}_2} \text{ wins}] \right| \leq \frac{1 + (q_{exec} \cdot q_{id})^2}{2^{|\text{TMSI}|}} \\ & + \text{Adv}_{\Pi}^{\text{K,Ind}}(\mathcal{A}) + 3 \cdot n_C \cdot (\text{Adv}_{\text{AE}}^{\text{ae}}(\mathcal{A}) + \text{Adv}_G^{\text{mac}}(\mathcal{A})). \end{aligned}$$

Game \mathbb{G}_3 : We modify the game \mathbb{G}_3 to replace outputs to call to the functions by truly randoms, i.e consistent values which are independent of the input, but the same input gives same output. We argue that the security loss is precisely the advantage of the adversary against the pseudorandomness of the internal cryptographic functions and the related security of

PKE and AE. It holds that:

$$\left| \Pr[\mathcal{A}_{\mathbb{G}_2} \text{ wins}] - \Pr[\mathcal{A}_{\mathbb{G}_3} \text{ wins}] \right| \leq n_C \cdot (\text{Adv}_G^{\text{prf}}(\mathcal{A}) + \text{Adv}_{\text{PKE}}^{\text{ind-cca2}}(\mathcal{A}) + \text{Adv}_{\text{AE}}^{\text{ae}}(\mathcal{A})).$$

Winning the game \mathbb{G}_3 : At this point, the adversary plays a game which consider to recover the bit b with a whole-randomized protocol. Assume that the adversary cannot learn any information about the related client in a such transcript. Thus, the adversary has only one option: guess the bit b without any specific information. So we obtain the following probability of winning the game \mathbb{G}_3 :

$$\Pr[\mathcal{A}_{\mathbb{G}_3} \text{ wins}] = \frac{1}{2}.$$

Security Statement This yields the following result:

$$\begin{aligned} \text{Adv}_{\Pi}^{\text{ww-unlink}}(\mathcal{A}) &\leq \text{Adv}_{\Pi}^{\text{K.Ind}}(\mathcal{A}_5) + \frac{1 + (q_{\text{exec}} \cdot q_{\text{id}})^2}{2^{|\text{TMSI}|}} + \\ &\frac{q_{\text{exec}}^2}{2^{|\text{R}|}} + \frac{(q_{\text{exec}} \cdot q_{\text{id}})^2}{2^{|\text{R}_{\text{id}}|}} + n_C \cdot (3 \cdot \text{Adv}_G^{\text{mac}}(\mathcal{A}_1) + \\ &\text{Adv}_G^{\text{prf}}(\mathcal{A}_2) + 4 \cdot \text{Adv}_{\text{AE}}^{\text{ae}}(\mathcal{A}_3) + \text{Adv}_{\text{PKE}}^{\text{ind-cca2}}(\mathcal{A}_4)). \end{aligned}$$

□

4.2.2 Key-Indistinguishability of our fixed variant

We consider the key-indistinguishability property, denoted K.Ind , as the guarantee that the session keys of honest sessions are indistinguishable from random. In the model previously detailed in section 2.2, we consider the session ID sid of each instance as follows: UID , ID_{S_i} , R , R_{id} , $\text{id}_{\times C}$ and the value Sq_{n_C} , that are agreed upon the session. As stipulated we can prove the key secrecy of our fixed variant of the AKA protocol, under indistinguishability, unforgeability and pseudorandomness properties of the different internal functions. This property is defined as follows:

Theorem 2. [K.Ind – resistance.] *Let $G : \{0, 1\}^{\kappa} \times \{0, 1\}^d \times \{0, 1\}^t \times \{0, 1\}^t \rightarrow \{0, 1\}^n$ be our specified function specified in section 4.2.1 and Π our fixed variant of the AKA protocol specified in section 4.2.1. Consider a $(t, q_{\text{exec}}, q_{\text{id}}, q_{\text{serv}}, q_{\text{Op}}, q_G, q_{\text{AE}}, q_{\text{PKE}})$ -adversary \mathcal{A} against the K.Ind -security of the protocol Π , running in time t and creating at most q_{exec} party instance, with at most q_{id} user identification per instance, corrupting at most q_{serv} servers, making at most q_{Op} OpAccess queries per operator per corrupted server, and making at most $q_G, q_{\text{AE}}, q_{\text{PKE}}$ queries to respectively the functions G, AE and PKE . Denote the advantage of this adversary as $\text{Adv}_{\Pi}^{\text{K.Ind}}(\mathcal{A})$. Then there exist a $(t' \approx O(t), q' = q_G + q_{\text{exec}})$ -mac-adversary \mathcal{A}_1 on G , a $(t' = O(t), q' = q_G + 2 \cdot q_{\text{exec}} + 5 \cdot q_{\text{serv}} \cdot q_{\text{Op}})$ -prf-adversary \mathcal{A}_2 on G , a $(t' = O(t), q' = q_G + q_{\text{exec}})$ -ind-cpa-adversary*

\mathcal{A}_3 on G , a $(t' = O(t), q' = q_{\text{exec}} + q_{\text{AE}})$ -ae-adversary \mathcal{A}_4 on AE , and a $(t' = O(t), q' = q_{\text{exec}} \cdot q_{\text{id}} + q_{\text{PKE}})$ -ind-cca2-adversary \mathcal{A}_5 on PKE such that:

$$\begin{aligned} \text{Adv}_{\Pi}^{\text{K.Ind}}(\mathcal{A}) &\leq n_C \cdot \left(\frac{(q_{\text{exec}} \cdot q_{\text{id}})^2}{2^{|\text{R}_{\text{id}}|}} + \frac{(q_{\text{exec}} + q_{\text{serv}} \cdot q_{\text{Op}})^2}{2^{|\text{R}|}} + \right. \\ &\text{Adv}_G^{\text{mac}}(\mathcal{A}_1) + \text{Adv}_G^{\text{prf}}(\mathcal{A}_2) + \text{Adv}_G^{\text{ind-cpa}}(\mathcal{A}_3) \\ &\left. + \text{Adv}_{\text{AE}}^{\text{ae}}(\mathcal{A}_4) + \text{Adv}_{\text{PKE}}^{\text{prf}}(\mathcal{A}_5) \right). \end{aligned}$$

Proof. Our proof has the following hops.

Game \mathbb{G}_0 : This game works as the K.Ind -game stipulated in our security model 2.2, but including the new oracles. The goal of the adversary $\mathcal{A}_{\mathbb{G}_0}$ is to distinguish, for a fresh instance that ends in an accepting state, the fresh session keys from random ones.

Game \mathbb{G}_1 : We modify \mathbb{G}_0 to only consider the new query $\text{Corrupt}(\text{P}, \text{type})$ but both games have the same goal. We note that this new query permits to consider the corruption of the key operator independently to the corruption of the subscriber keys. This new query behaves as follows:

$\text{Corrupt}(\text{P}, \text{type})$: yields to the adversary the long-term keys of party $\text{P} \neq \text{S}$ (else, if the oracle takes as input $\text{P} = \text{S}$, then it outputs \perp). The output of the oracle depends on the value $\text{type} \in \{\text{sub}, \text{op}, \text{all}\}$. If $\text{type} = \text{sub}$, then the returned value is sk_{P} . If $\text{type} = \text{op}$, then the oracle returns sk_{op} . Then, for $\text{type} = \text{all}$, we return the both values $\text{sk}_{\text{P}}, \text{sk}_{\text{op}}$. If $\text{type} \in \{\text{sub}, \text{all}\}$, then P (and all its instances, past, present, or future), are considered to be adversarially controlled.

We argue that given any adversary \mathcal{A} playing the game \mathbb{G}_1 and winning w.p $\epsilon_{\mathcal{A}}$ the same adversary wins the game \mathbb{G}_0 w.p at least $\epsilon_{\mathcal{A}}$ (this is trivial since in game \mathbb{G}_1 , \mathcal{A} has more information).

$$\Pr[\mathcal{A}_{\mathbb{G}_0} \text{ wins}] \leq \Pr[\mathcal{A}_{\mathbb{G}_1} \text{ wins}].$$

Game \mathbb{G}_2 : We modify \mathbb{G}_1 to only allow interactions with a single client (any future UReg calls for a client would be answered with an error symbol \perp). The challenger generates only a single operator key, which is associated with the operator chosen for the registered client and chooses a bit $b \in \{0, 1\}$. We proceed as follows: for any adversary $\mathcal{A}_{\mathbb{G}_1}$ winning the game \mathbb{G}_1 with a non-negligible success probability $\epsilon_{\mathbb{G}_1}$, we propose to construct a generic adversary $\mathcal{A}_{\mathbb{G}_2}$ winning the game \mathbb{G}_2 with a black-box access to the adversary $\mathcal{A}_{\mathbb{G}_1}$.

Adversary $\mathcal{A}_{\mathbb{G}_2}$ begins by choosing a single client C . For every user registration request that $\mathcal{A}_{\mathbb{G}_1}$ sends to its challenger, $\mathcal{A}_{\mathbb{G}_2}$ responds as follows: if the registered client is C , then it forwards the exact UReg query that $\mathcal{A}_{\mathbb{G}_1}$ makes to its own UReg oracle. Else, if $\mathcal{A}_{\mathbb{G}_1}$ registers any client $\text{C}^* \neq \text{C}$, $\mathcal{A}_{\mathbb{G}_2}$

simulates the registration, generating sk_{C^*} and Sqn_{C^*} , returning the latter value. Adversary \mathcal{A}_{G_2} also generates $n_{Op} - 1$ operator keys rsk_{Op^*} (for all operator Op^* such that $Op^* \neq Op$), and associates them with the clients as follows: the target client C is associated with the same operator given as input by \mathcal{A}_{G_1} to the UReg query (thus with the operator key sk_{Op} generated by the challenger of game \mathbb{G}_2). Let this target operator be denoted as Op . Adversary \mathcal{A}_{G_2} queries $\text{Corrupt}(C, Op)$ and stores sk_{Op} .

We distinguish two types of clients: the *brothers* of the target client (i.e the clients which have the same operator as the target client) and the others ones. For these latter, denoted C^* , which are registered by \mathcal{A}_{G_1} with an operator $Op^* \neq Op$, adversary \mathcal{A}_{G_2} associates Op^* with one of its generated keys rsk_{Op^*} . Recall that, since adversary \mathcal{A}_{G_1} plays the game in the presence of n_{Op} operators, there are $n_{Op} - 1$ keys which will be used this way. We call all clients $C^* \neq C$ registered by \mathcal{A}_{G_0} with the target operator Op the *brothers* of the target client C . Adversary \mathcal{A}_{G_2} associates each brother of C with the corrupted key sk_{Op} it learns from its challenger.

In the rest of the simulation, whenever \mathcal{A}_{G_1} makes a query to an instance of some party C^* , not a brother of C , the adversary \mathcal{A}_{G_2} simulates the response using the values sk_{C^*} , rsk_{Op^*} , and the current value of Sqn . For the brothers of C , the simulation is done with sk_{C^*} , sk_{Op} , and the current Sqn . For the target client C , any queries are forwarded by \mathcal{A}_{G_2} to its challenger.

Any corruption or reveal queries are dealt with in a similar way. Note that \mathcal{A}_{G_2} cannot query Corrupt to its adversary (this is a condition of freshness). The simulation is thus perfect up to the Test query.

In the Test query, \mathcal{A}_{G_1} chooses a fresh session and sends it to \mathcal{A}_{G_2} (acting as a challenger). Note that \mathcal{A}_{G_2} will be able to test whether this instance is fresh, as freshness is defined in terms of \mathcal{A}_{G_1} 's queries. If \mathcal{A}_{G_1} queries Test with a client other than the target client C , then \mathcal{A}_{G_2} aborts the simulation, tests a random, fresh instance of the client C (creating one if necessary), and guesses the bit d , winning with probability $\frac{1}{2}$. Else, if \mathcal{A}_{G_1} queried a fresh instance of C , \mathcal{A}_{G_2} forwards this choice to its challenger and receives the challenger's input. The adversary \mathcal{A}_{G_2} forwards the input of the challenger to \mathcal{A}_{G_1} and then receives \mathcal{A} 's output d , which will be \mathcal{A}_{G_2} 's own response to its own challenger.

Denote by E_1 the event that adversary tests C in game \mathbb{G}_1 , while \bar{E}_1 denotes the event that \mathcal{A}_{G_1} chooses to test $C^* \neq C$.

It holds that:

$$\begin{aligned} \Pr[\mathcal{A}_{G_2} \text{ wins}] &= \Pr[\mathcal{A}_{G_2} \text{ wins} \mid E_1] \cdot \Pr[E_1] + \\ &\quad \Pr[\mathcal{A}_{G_2} \text{ wins} \mid \bar{E}_1] \cdot \Pr[\bar{E}_1] \\ &\geq \frac{1}{n_C} \Pr[\mathcal{A}_{G_1} \text{ wins}] + \frac{1}{2} \cdot \left(1 - \frac{1}{n_C}\right) \\ &\geq \frac{1}{n_C} \Pr[\mathcal{A}_{G_0} \text{ wins}] + \frac{1}{2} \cdot \left(1 - \frac{1}{n_C}\right). \end{aligned}$$

Note that adversary \mathcal{A}_{G_2} makes one extra query with respect to \mathcal{A}_{G_1} , since we need to learn the key of the target operator.

Game \mathbb{G}_3 : We modify \mathbb{G}_2 to ensure that the random values sampled by honest server instances are always unique.

This gives us a security loss (related to the respective collisions between the R and R_{id} in two different instances) of

$$\begin{aligned} |\Pr[\mathcal{A}_{G_2} \text{ wins}] - \Pr[\mathcal{A}_{G_3} \text{ wins}]| &\leq \frac{(q_{exec} \cdot q_{id})^2}{2^{|R_{id}|}} + \\ &\quad \frac{(q_{exec} + q_{serv} \cdot q_{Op})^2}{2^{|R|}}. \end{aligned}$$

Game \mathbb{G}_4 : This game behaves as the game \mathbb{G}_3 with the restriction to only interact with only one server (any future UReg calls for a server would be answered with an error symbol \perp). The benefices lost is the ability to obtain some authentication vectors from corrupted servers. Such authentication vectors can either give information about the used sequence number and the long term keys or forge a fresh challenge replaying some parts of these vectors. We recall that the challenge is split in four parts: a random value, a masked version of the fresh sequence number (an one-time-pad based on an anonymity key generated by the function G), a mac computed with the function G and an authenticated and encrypted version of the next temporary identifier and the current index. Moreover, we note that all the call of the function G take in input a specific value of the related server, denoted ID_S . The corrupted servers permit to obtain vectors based on the fresh sequence number but different random and different server identifier. So the related security loss is given by the collision on two outputs of the same function G with two different inputs (the only differences between the both inputs are at least the value of the network identifier) and by the indistinguishability of the function G . We recall that the Test Phase of the game can be only focus on a network which is or was never corrupted. This give us a security loss

$$|\Pr[\mathcal{A}_{G_4} \text{ wins}] - \Pr[\mathcal{A}_{G_3} \text{ wins}]| \leq \text{Adv}_G^{\text{ind-cpa}}(\mathcal{A}) + \text{Adv}_G^{\text{mac}}(\mathcal{A}).$$

Game \mathbb{G}_5 : We modify \mathbb{G}_4 to replace outputs of the internal cryptographic functions by truly random, but consistent values (they are independent of the input, but the same input gives the same output). We argue that the security loss is precisely the advantage of the adversary \mathcal{A} against the pseudorandomness of function G , and the security of respectively PKE and AE. Note that the total number of queries to the related functions are at most $2 \cdot q_G$ and one q_{AE} and one q_{PKE} per honest instance (thus totaling at most $q_G + (2 \cdot q_{exec})$ queries to the function G , q_{exec} queries to the function PKE and q_{exec} queries to the function AE).

$$\left| \Pr[\mathcal{A}_{\mathbb{G}_4} \text{ wins}] - \Pr[\mathcal{A}_{\mathbb{G}_5} \text{ wins}] \right| \leq \text{Adv}_{\text{PKE}}^{\text{ind-cca2}}(\mathcal{A}) + \text{Adv}_G^{\text{prf}}(\mathcal{A}) + \text{Adv}_{\text{AE}}^{\text{ae}}(\mathcal{A}).$$

Winning \mathbb{G}_5 : At this point, the adversary plays a game in the presence of a single client C . The goal of this adversary is to distinguish a random session key to a fresh session key. But, in game \mathbb{G}_5 , queries to G return truly random, consistent values. In this case, the adversary can do no better than guessing. Thus, we have:

$$\Pr[\mathcal{A}_{\mathbb{G}_5} \text{ wins}] = \frac{1}{2}.$$

Security statement: This yields the following result:

$$\begin{aligned} \text{Adv}_{\Pi}^{\text{K.Ind}}(\mathcal{A}_{\mathbb{G}_0}) &\leq n_C \cdot \left(\frac{(q_{\text{exec}} \cdot q_{\text{id}})^2}{2^{|\text{R}_{\text{id}}|}} + \frac{(q_{\text{exec}} + q_{\text{serv}} \cdot q_{\text{Op}})^2}{2^{|\text{R}|}} \right) \\ &\quad + \text{Adv}_G^{\text{mac}}(\mathcal{A}_1) + \text{Adv}_G^{\text{prf}}(\mathcal{A}_2) + \text{Adv}_G^{\text{ind-cpa}}(\mathcal{A}_3) \\ &\quad + \text{Adv}_{\text{AE}}^{\text{ae}}(\mathcal{A}_4) + \text{Adv}_{\text{PKE}}^{\text{prf}}(\mathcal{A}_5). \end{aligned}$$

This concludes the proof. \square

4.2.3 Impersonation of our fixed variant

We present first the client-impersonation resistance proof, then the equivalent statement for server impersonations.

Theorem 3. [C.Imp – resistance.] *Let $G : \{0, 1\}^\kappa \times \{0, 1\}^d \times \{0, 1\}^t \times \{0, 1\}^t \rightarrow \{0, 1\}^n$ be our specified function specified in section 4.2.1 and Π our fixed variant of the AKA protocol specified in section 4.2.1. Consider a $(t, q_{\text{exec}}, q_{\text{id}}, q_{\text{serv}}, q_{\text{Op}}, q_G, q_{\text{AE}}, q_{\text{PKE}})$ -adversary \mathcal{A} against the C.Imp-security of the protocol Π , running in time t and creating at most q_{exec} party instance, with at most q_{id} user identification per instance, corrupting at most q_{serv} servers, making at most q_{Op} OpAccess queries per operator per corrupted server, and making at most $q_G, q_{\text{AE}}, q_{\text{PKE}}$ queries to respectively the functions G, AE and PKE . Denote the advantage of this adversary as $\text{Adv}_{\Pi}^{\text{C.Imp}}(\mathcal{A})$. Then, there exist a $(t' \approx O(t), q' = q_G + 2 \cdot q_{\text{exec}} + 5 \cdot q_{\text{serv}} \cdot q_{\text{Op}})$ -prf-adversary \mathcal{A}_1 , a $(t' = O(t), q' = q_{\text{exec}} + q_{\text{AE}})$ -ae-adversary \mathcal{A}_2 on AE and a $(t' = O(t), q' = q_{\text{exec}} \cdot q_{\text{id}} + q_{\text{PKE}})$ -ind-cca2-adversary \mathcal{A}_3 on PKE such that:*

$$\begin{aligned} \text{Adv}_{\Pi}^{\text{C.Imp}}(\mathcal{A}) &\leq n_C \cdot (\text{Adv}_G^{\text{prf}}(\mathcal{A}_1) + \text{Adv}_{\text{AE}}^{\text{ae}}(\mathcal{A}_2) \\ &\quad + \text{Adv}_{\text{PKE}}^{\text{ind-cca2}}(\mathcal{A}_3) + \frac{(q_{\text{exec}} \cdot q_{\text{id}})^2}{2^{|\text{R}_{\text{id}}|}} \\ &\quad + \frac{(q_{\text{exec}})^2}{2^{|\text{R}|}} + \frac{q_{\text{exec}}}{2^{|\text{Res}|}} + \frac{1}{2^\kappa} + \frac{q_{\text{exec}} \cdot q_{\text{id}}}{2^{|\text{ID}|}}). \end{aligned}$$

Proof. Our proof has the following hops:

Game \mathbb{G}_0 : This game works as the C.Imp-game: When \mathcal{A} stops, she wins if there exists an instance S_i that ends in an accepting state with session ID sid and partner ID pid such that: (a) pid is not adversarially controlled (sk_{pid} has not been corrupted), (b) no other instance C_i exists for $\text{pid} = S_i$ that ends in an accepting state, with session ID sid .

Game \mathbb{G}_1 : We modify the game to allow the new $\text{Corrupt}(P, \text{type})$ query from the previous proof. It holds that:

$$\Pr[\mathcal{A}_{\mathbb{G}_0} \text{ wins}] \leq \Pr[\mathcal{A}_{\mathbb{G}_1} \text{ wins}].$$

Game \mathbb{G}_2 : We modify \mathbb{G}_1 to only interact with a single client, as in the previous proof, giving a security loss of:

$$\Pr[\mathcal{A}_{\mathbb{G}_2} \text{ wins}] \geq \frac{1}{n_C} \Pr[\mathcal{A}_{\mathbb{G}_1} \text{ wins}] + \frac{1}{2} \cdot \left(1 - \frac{1}{n_C}\right)$$

Game \mathbb{G}_3 : We now restrict the adversary to using a single server As as detailed in the key-indistinguishability proof, the related security loss is given by: This give us a security loss

$$\left| \Pr[\mathcal{A}_{\mathbb{G}_3} \text{ wins}] - \Pr[\mathcal{A}_{\mathbb{G}_2} \text{ wins}] \right| \leq \text{Adv}_G^{\text{prf}}(\mathcal{A}).$$

Game \mathbb{G}_4 : We modify \mathbb{G}_4 to replace outputs to calls to all the internal cryptographic functions by truly random, but consistent values, and as before, we lose a term:

$$\left| \Pr[\mathcal{A}_{\mathbb{G}_3} \text{ wins}] - \Pr[\mathcal{A}_{\mathbb{G}_4} \text{ wins}] \right| \leq \text{Adv}_G^{\text{prf}}(\mathcal{A}) + \text{Adv}_{\text{AE}}^{\text{ae}}(\mathcal{A}) + \text{Adv}_{\text{PKE}}^{\text{ind-cca2}}(\mathcal{A}).$$

Game \mathbb{G}_5 : We modify \mathbb{G}_4 to ensure that the random values sampled by honest server instances are always unique. As in the unlinkability proof, this yields a loss of:

$$\left| \Pr[\mathcal{A}_{\mathbb{G}_4} \text{ wins}] - \Pr[\mathcal{A}_{\mathbb{G}_5} \text{ wins}] \right| \leq \frac{(q_{\text{exec}} \cdot q_{\text{id}})^2}{2^{|\text{R}_{\text{id}}|}} + \frac{q_{\text{exec}}^2}{2^{|\text{R}|}}.$$

Winning \mathbb{G}_5 : At this point, the adversary plays a game with a single client and server. A server instance S_i only accepts $\mathcal{A}_{\mathbb{G}_5}$, if the latter can generate a fresh identification response ID and an authentication response Res for some session sid . Assume that this happens against accepting instance S_i of the server, for some target session sid . Note that the values Res and ID computed by C_i are purely random, but consistent. Thus, the adversary has three options for each of these values: (a) forwarding a value already received from the honest client for the same input values, of which sk_C is unknown; (b) guessing the key sk_C ; or (c) guessing the value. The first option yields no result, since it implies there exists a previous client instance with the same session id sid as the client. The second option happens with a probability of $2^{-|\text{sk}_C|}$. The third option occurs

with a probability of $2^{-|\text{Res}|} + 2^{-|\text{ID}|}$ per session, thus a total of $q_{\text{exec}} \cdot (2^{-|\text{Res}|} + q_{\text{id}} \cdot 2^{-|\text{ID}|})$. Thus,

$$\Pr[\mathcal{A}_{\mathbb{G}_3} \text{ wins}] = 2^{-|\text{sk}_c|} + q_{\text{exec}} \cdot (2^{-|\text{Res}|} + q_{\text{id}} \cdot 2^{-|\text{ID}|}).$$

Security statement: This yields the following result:

$$\begin{aligned} \text{Adv}_{\Pi}^{\text{C.Imp}}(\mathcal{A}_{\mathbb{G}_0}) &\leq n_C \cdot (\text{Adv}_G^{\text{prf}}(\mathcal{A}_1) + \text{Adv}_{\text{AE}}^{\text{ae}}(\mathcal{A}_2) \\ &\quad + \text{Adv}_{\text{PKE}}^{\text{ind-cca2}}(\mathcal{A}_3) + \frac{(q_{\text{exec}} \cdot q_{\text{id}})^2}{2^{|\text{R}_{\text{id}}|}} \\ &\quad + \frac{(q_{\text{exec}})^2}{2^{|\text{R}|}} + \frac{q_{\text{exec}}}{2^{|\text{Res}|}} + \frac{1}{2^\kappa} + \frac{q_{\text{exec}} \cdot q_{\text{id}}}{2^{|\text{ID}|}}). \end{aligned}$$

Theorem 4. [S.Imp – resistance.] *Let $G : \{0, 1\}^\kappa \times \{0, 1\}^d \times \{0, 1\}^t \times \{0, 1\}^t \rightarrow \{0, 1\}^n$ be our specified function specified in section 4.2.1 and Π our fixed variant of the AKA protocol specified in section 4.2.1. Consider a $(t, q_{\text{exec}}, q_{\text{id}}, q_{\text{serv}}, q_{\text{Op}}, q_G, q_{\text{AE}}, q_{\text{PKE}})$ -adversary \mathcal{A} against the S.Imp-security of the protocol Π , running in time t and creating at most q_{exec} party instance, with at most q_{id} user identification per instance, corrupting at most q_{serv} servers, making at most q_{Op} OpAccess queries per operator per corrupted server, and making at most $q_G, q_{\text{AE}}, q_{\text{PKE}}$ queries to respectively the functions G, AE and PKE . Denote the advantage of this adversary as $\text{Adv}_{\Pi}^{\text{S.Imp}}(\mathcal{A})$. Then, there exist a $(t' \approx O(t), q' = q_G + 2 \cdot q_{\text{exec}} + 5 \cdot q_{\text{serv}} \cdot q_{\text{Op}})$ -prf-adversary \mathcal{A}_1 , a $(t' = O(t), q' = q_{\text{exec}} + q_{\text{AE}})$ -ae-adversary \mathcal{A}_2 on AE and a $(t' = O(t), q' = q_{\text{exec}} \cdot q_{\text{id}} + q_{\text{PKE}})$ -ind-cca2-adversary \mathcal{A}_3 on PKE such that:*

$$\begin{aligned} \text{Adv}_{\Pi}^{\text{S.Imp}}(\mathcal{A}) &\leq n_C \cdot \left(\frac{q_{\text{exec}}}{2^{|\text{Mac}_S|}} + \frac{1}{2^\kappa} + \text{Adv}_G^{\text{prf}}(\mathcal{A}_1) \right. \\ &\quad \left. + \text{Adv}_{\text{AE}}^{\text{ae}}(\mathcal{A}_2) + \text{Adv}_{\text{PKE}}^{\text{ind-cca2}}(\mathcal{A}_3) \right). \end{aligned}$$

Proof. Our proof has the following hops:

Game \mathbb{G}_0 : This game works as the S.Imp-game. The adversary \mathcal{A} wins if there exists an instance C_i that ends in an accepting state with session ID sid and partner ID and pid s.t.: (a) $\text{pid} = S$, (b) no instance S_j exists such as S_j and C_i has the same session ID sid , (c) C_i and these partners are not adversarially controlled.

Game \mathbb{G}_1 : We add the new query $\text{Corrupt}(P, \text{type})$ as in the previous proof, such that:

$$\Pr[\mathcal{A}_{\mathbb{G}_0} \text{ wins}] \leq \Pr[\mathcal{A}_{\mathbb{G}_1} \text{ wins}].$$

Game \mathbb{G}_2 : We modify \mathbb{G}_1 to only interact with a single client, as in the previous proofs, and lose:

$$\Pr[\mathcal{A}_{\mathbb{G}_2} \text{ wins}] \geq \frac{1}{n_C} \Pr[\mathcal{A}_{\mathbb{G}_1} \text{ wins}] + \frac{1}{2} \cdot \left(1 - \frac{1}{n_C}\right)$$

Game \mathbb{G}_3 : Again, we restrict the adversary to only one server. This give us a security loss

$$|\Pr[\mathcal{A}_{\mathbb{G}_3} \text{ wins}] - \Pr[\mathcal{A}_{\mathbb{G}_2} \text{ wins}]| \leq \text{Adv}_G^{\text{prf}}(\mathcal{A}).$$

Game \mathbb{G}_4 : We modify \mathbb{G}_3 to replace outputs to calls to all the internal cryptographic functions by truly random, but consistent values and as before, it holds that:

$$\begin{aligned} \square \quad |\Pr[\mathcal{A}_{\mathbb{G}_3} \text{ wins}] - \Pr[\mathcal{A}_{\mathbb{G}_4} \text{ wins}]| &\leq \text{Adv}_G^{\text{prf}}(\mathcal{A}) + \text{Adv}_{\text{AE}}^{\text{ae}}(\mathcal{A}) \\ &\quad + \text{Adv}_{\text{PKE}}^{\text{ind-cca2}}(\mathcal{A}). \end{aligned}$$

Winning \mathbb{G}_4 : At this point, the adversary plays a game with a single client C_i , which only accepts $\mathcal{A}_{\mathbb{G}_4}$, if the authentication challenge is verified for some session sid . Assume that this happens against accepting instance C_i of the target client, for some target session sid . Note that the Mac_S value computed (for verification) by C_i is purely random, but consistent. Thus, the adversary has three options: (a) forwarding a value already received from the honest server for the same input values $R; \text{Sqn}; \text{sk}_{\text{op}}; \text{sk}_c$, of which sk_c is unknown; (b) guessing the key sk_c ; or (c) guessing the response. The first option yields no result since there are no collision between the transcript of two different servers since all the servers have a different session ID. The second option happens with a probability of $2^{-|\text{sk}_c|}$. The third option occurs with a probability of $2^{-|\text{Mac}_S|}$ per session, thus a total of $q_{\text{exec}} \cdot 2^{-|\text{Mac}_S|}$. Thus,

$$\Pr[\mathcal{A}_{\mathbb{G}_4} \text{ wins}] = 2^{-|\text{sk}_c|} + q_{\text{exec}} \cdot 2^{-|\text{Mac}_S|}.$$

Security statement: This yields the following result:

$$\begin{aligned} \text{Adv}_{\Pi}^{\text{S.Imp}}(\mathcal{A}_{\mathbb{G}_0}) &\leq n_C \cdot \left(\frac{q_{\text{exec}}}{2^{|\text{Mac}_S|}} + \frac{1}{2^\kappa} + \text{Adv}_G^{\text{prf}}(\mathcal{A}_1) \right. \\ &\quad \left. + \text{Adv}_{\text{AE}}^{\text{ae}}(\mathcal{A}_2) + \text{Adv}_{\text{PKE}}^{\text{ind-cca2}}(\mathcal{A}_3) \right). \end{aligned}$$

□

4.2.4 Soundness and key-confidentiality of our fixed variant.

Theorem 5. [St.conf – resistance.] *Let G and G^* be our specified functions specified in section 4.2.1 and Π our fixed variant of the AKA protocol specified in section 4.2.1. Consider a $(t, q_{\text{exec}}, q_{\text{id}}, q_{\text{Op}}, q_G, q_{G^*}, q_{\text{AE}}, q_{\text{PKE}})$ -adversary \mathcal{A} against the St.conf-security of the protocol Π ,*

running in time t and executing q_{exec} sessions of the protocol Π , making at most q_{Op} queries to any operator, and making $q_G, q_{G^*}, q_{\text{AE}}, q_{\text{PKE}}$ queries to respectively the functions G, G^*, AE and PKE resp. q_{G^*} queries to the function G (resp. G^*). Denote the advantage of this adversary as $\text{Adv}_{\Pi}^{\text{St.conf}}(\mathcal{A})$. Then there exist a $(t' \approx O(t), q' = q_G + 5 \cdot q_{\text{Op}} + 2 \cdot q_{\text{exec}})$ -prf-adversary \mathcal{A}_1 on G , a $(t' = O(t), q' = q_{\text{exec}} + q_{\text{AE}})$ -ae-adversary \mathcal{A}_2 on AE , a $(t' = O(t), q' = q_{\text{exec}} \cdot q_{\text{id}} + q_{\text{PKE}})$ -ind-cca2-adversary \mathcal{A}_3 on PKE , and $(t' \approx O(t), q' = q_{G^*})$ -prf-adversary \mathcal{A}_4 on G^* a such that:

$$\begin{aligned} \text{Adv}_{\Pi}^{\text{St.conf}}(\mathcal{A}) \leq & n_C \cdot \left(\frac{1}{2^{|\text{sk}_C|}} + \frac{1}{2^{|\text{sk}_{\text{Op}}|}} + \frac{1}{2^{|\text{Sqn}_C|}} \right) \\ & + \text{Adv}_G^{\text{prf}}(\mathcal{A}_1) + \text{Adv}_{\text{AE}}^{\text{ae}}(\mathcal{A}_2) \\ & + \text{Adv}_{\text{PKE}}^{\text{ind-cca2}}(\mathcal{A}_3) + \text{Adv}_{G^*}^{\text{prf}}(\mathcal{A}_4). \end{aligned}$$

Proof. Our proof has the following hops.

Game \mathbb{G}_0 : This game works as the St.conf -game stipulated in our security model. The goal of the adversary $\mathcal{A}_{\mathbb{G}_0}$ is to recover at least one secret value, i.e the subscriber key sk_C , my operator key sk_{Op} or the subscriber sequence number Sqn_C for a fresh instance.

Game \mathbb{G}_1 : We modify \mathbb{G}_0 to only allow interactions with one operator. The challenger related to the game \mathbb{G}_1 only generates a single operator key, which is associated with the operator chosen for the registered client. We proceed as follows: for any adversary $\mathcal{A}_{\mathbb{G}_0}$ winning the game \mathbb{G}_0 with a non-negligible success probability $\epsilon_{\mathbb{G}_0}$, we propose to construct a generic adversary $\mathcal{A}_{\mathbb{G}_1}$ winning the game \mathbb{G}_1 with a black-box access to the adversary $\mathcal{A}_{\mathbb{G}_0}$.

Adversary $\mathcal{A}_{\mathbb{G}_1}$ begins by choosing a single operator Op . It generates $n_{\text{Op}} - 1$ operator keys, denoted rsk_{Op^*} . Then, for every user registration request that $\mathcal{A}_{\mathbb{G}_0}$ sends to its challenger, $\mathcal{A}_{\mathbb{G}_1}$ responds as follows: if the request $\text{CreateCl}(\cdot)$ takes in input the operator Op , then it forwards the same query to its own oracle. Else, if $\mathcal{A}_{\mathbb{G}_0}$ sends a registration request based on any operator $\text{Op}^* \neq \text{Op}$, $\mathcal{A}_{\mathbb{G}_1}$ simulates the registration, generating a subscriber key sk_{C^*} and a sequence number Sqn_{C^*} , returning the latter value. Moreover, each new client registered with the operator Op (resp. any Op^*) is associated with the related operator key sk_{Op} (resp. rsk_{Op^*}).

We distinguish two types of clients: the *brothers* of the target client (i.e the clients which have the same operator as the target client) and the others ones. For these latter, denoted C^* , which are registered by $\mathcal{A}_{\mathbb{G}_1}$ with an operator $\text{Op}^* \neq \text{Op}$, adversary $\mathcal{A}_{\mathbb{G}_2}$ associates Op^* with one of its generated keys rsk_{Op^*} .

In the rest of the simulation, whenever $\mathcal{A}_{\mathbb{G}_0}$ makes a query to an instance of some party C^* (from any operator except Op), the adversary $\mathcal{A}_{\mathbb{G}_1}$ simulates the response using the val-

ues sk_{C^*} , rsk_{Op^*} , and the current value of Sqn_{C^*} . For the other clients, the query is forwarded by $\mathcal{A}_{\mathbb{G}_1}$ to its own challenger.

Any corruption or reveal queries are dealt with in a similar way. Note that $\mathcal{A}_{\mathbb{G}_1}$ cannot query Corrupt to its adversary (this is a condition of freshness). The simulation is thus perfect up to the Test query.

In the Test query, $\mathcal{A}_{\mathbb{G}_0}$ chooses a fresh instance and sends it to $\mathcal{A}_{\mathbb{G}_1}$ (acting as a challenger). Note that $\mathcal{A}_{\mathbb{G}_1}$ will be able to test whether this instance is fresh, as freshness is defined in terms of $\mathcal{A}_{\mathbb{G}_0}$'s queries. If $\mathcal{A}_{\mathbb{G}_0}$ queries an instance C_i^* for the Test query, then $\mathcal{A}_{\mathbb{G}_1}$ aborts the simulation, tests a random tuple about any fresh instance of the client C (creating one if necessary), winning with probability $\frac{1}{2^{|\text{sk}_C|}} + \frac{1}{2^{|\text{sk}_{\text{Op}}|}} + \frac{1}{2^{|\text{Sqn}_C|}}$. Else, if $\mathcal{A}_{\mathbb{G}_0}$ sends a tuple of a fresh instance of C_i , $\mathcal{A}_{\mathbb{G}_1}$ forwards this choice to its challenger and receives the challenger's output which contains the result of this game.

Denote by E_1 the event that adversary $\mathcal{A}_{\mathbb{G}_0}$ tests an instance C_i (from the chosen operator Op), while \bar{E}_1 denotes the event that $\mathcal{A}_{\mathbb{G}_0}$ chooses to test C_i^* .

It holds that:

$$\begin{aligned} \Pr[\mathcal{A}_{\mathbb{G}_1} \text{ wins}] &= \Pr[\mathcal{A}_{\mathbb{G}_1} \text{ wins} \mid E_1] \cdot \Pr[E_1] + \\ &\Pr[\mathcal{A}_{\mathbb{G}_1} \text{ wins} \mid \bar{E}_1] \cdot \Pr[\bar{E}_1] \\ &\geq \frac{1}{n_{\text{Op}}} \Pr[\mathcal{A}_{\mathbb{G}_0} \text{ wins}] + \left(1 - \frac{1}{n_{\text{Op}}}\right) \cdot \\ &\left(\frac{1}{2^{|\text{sk}_C|}} + \frac{1}{2^{|\text{sk}_{\text{Op}}|}} + \frac{2}{2^{|\text{Sqn}_C|}}\right). \end{aligned}$$

Note that adversary $\mathcal{A}_{\mathbb{G}_1}$ makes no query with respect to $\mathcal{A}_{\mathbb{G}_0}$.

Game \mathbb{G}_2 : We modify \mathbb{G}_1 to only allow interactions with a single client (any future $\text{CreateCl}(\text{Op})$ calls for a client would be answered with an error symbol \perp). We recall that the two adversaries $\mathcal{A}_{\mathbb{G}_1}$ and $\mathcal{A}_{\mathbb{G}_2}$ interact with clients from a single operator key, denoted Op , which is associated with the operator key sk_{Op} . We proceed as follows: for any adversary $\mathcal{A}_{\mathbb{G}_1}$ winning the game \mathbb{G}_1 with a non-negligible success probability $\epsilon_{\mathbb{G}_1}$, we propose to construct a generic adversary $\mathcal{A}_{\mathbb{G}_2}$ winning the game \mathbb{G}_2 with a black-box access to the adversary $\mathcal{A}_{\mathbb{G}_1}$. Adversary $\mathcal{A}_{\mathbb{G}_2}$ begins by choosing a single client C . For every user registration request that $\mathcal{A}_{\mathbb{G}_1}$ sends to its challenger, $\mathcal{A}_{\mathbb{G}_2}$ responds as follows: for a new client $C^* \neq C$ it generates sk_{C^*} and Sqn_{C^*} , returning the latter value.

In the rest of the simulation, whenever $\mathcal{A}_{\mathbb{G}_1}$ makes a query to an instance of some party C^* , the adversary $\mathcal{A}_{\mathbb{G}_2}$ simulates the response using the oracle of the function G^* and the values sk_{C^*} and the current value of Sqn_{C^*} . Indeed, since the adversary can corrupt any operator key, she requires the oracle of G^* permitting to simulate all the queries of the brothers of the target client.

For the target client C , any queries are forwarded by $\mathcal{A}_{\mathbb{G}_2}$ to its challenger. Any corruption or reveal queries are dealt

with in a similar way. Note that \mathcal{A}_{G_2} cannot query `Corrupt` to its adversary (this is a condition of freshness). The simulation is thus perfect up to the `Test` query.

In the `Test` query, \mathcal{A}_{G_1} chooses a fresh instance and sends it to \mathcal{A}_{G_2} (acting as a challenger). Note that \mathcal{A}_{G_2} will be able to test whether this instance is fresh, as freshness is defined in terms of \mathcal{A}_{G_1} 's queries. If \mathcal{A}_{G_1} queries `Test` with a client other than the target client C , then \mathcal{A}_{G_2} aborts the simulation, tests a random tuple as the previous reduction. Else, if \mathcal{A}_{G_1} queried a fresh instance of C , \mathcal{A}_{G_2} forwards this choice to its challenger and receives the challenger's which contains the result of this game. It holds that:

$$\Pr[\mathcal{A}_{G_2} \text{ wins}] \geq \frac{1}{n_{C,Op}} \Pr[\mathcal{A}_{G_1} \text{ wins}] + \frac{1}{2} \cdot \left(1 - \frac{1}{n_{C,Op}}\right)$$

, with at most $n_{C,Op}$ clients by operator.

Note that adversary \mathcal{A}_{G_2} makes no extra query with respect to \mathcal{A}_{G_1} .

Game \mathbb{G}_3 : We modify \mathbb{G}_2 to replace outputs of the internal cryptographic functions by truly random, but consistent values (they are independent of the input, but the same input gives the same output). We argue that the security loss is precisely the advantage of the adversary \mathcal{A} against the pseudo-randomness of functions G and G^* , and related security of the functions PKE and AE.

$$\begin{aligned} |\Pr[\mathcal{A}_{G_3} \text{ wins}] - \Pr[\mathcal{A}_{G_2} \text{ wins}]| &\leq \text{Adv}_G^{\text{prf}}(\mathcal{A}) + \text{Adv}_{G^*}^{\text{prf}}(\mathcal{A}) \\ &\quad + \text{Adv}_{AE}^{\text{ae}}(\mathcal{A}) + \text{Adv}_{PKE}^{\text{ind-cca2}}(\mathcal{A}). \end{aligned}$$

Winning Game \mathbb{G}_3 : At this point, the adversary plays a game with an uncorruptible single client C_i in a protocol including truly but consistent values. She wins if she can output a tuple $(C_i, sk_C^*, sk_{Op}^*, Sqn_C^*, Sqn_{Op,C}^*)$ such as at least one of these values corresponds to the real related secret value of the instance C_i . Thus, the adversary has only one choice to win this game: guessing each value. So the probability that the adversary \mathcal{A}_{G_3} wins is as follows:

$$\Pr[\mathcal{A}_{G_3} \text{ wins}] = \frac{1}{2^{|sk_C|}} + \frac{1}{2^{|sk_{Op}|}} + \frac{2}{2^{|Sqn|}}.$$

Security statement: This yields the following result:

$$\begin{aligned} \text{Adv}_{\Pi}^{\text{St.conf}}(\mathcal{A}_{G_0}) &\leq n_C \cdot \left(\frac{1}{2^{|sk_C|}} + \frac{1}{2^{|sk_{Op}|}} + \frac{2}{2^{|Sqn|}} \right) \\ &\quad + \text{Adv}_G^{\text{prf}}(\mathcal{A}_1) + \text{Adv}_{AE}^{\text{ae}}(\mathcal{A}_2) \\ &\quad + \text{Adv}_{PKE}^{\text{ind-cca2}}(\mathcal{A}_3) + \text{Adv}_{G^*}^{\text{prf}}(\mathcal{A}_4). \end{aligned}$$

□

We provide the following theorem about the client-impersonation, denoted `S.sound-security`, of the fixed variant of the AKA protocol.

Theorem 6. [`S.sound – resistance`.] *Let $G : \{0, 1\}^{\kappa} \times \{0, 1\}^d \times \{0, 1\}^t \times \{0, 1\}^t \rightarrow \{0, 1\}^n$ be our specified function specified in section 4.2.1 and Π our fixed variant of the AKA protocol specified in section 4.2.1. Consider a $(t, q_{\text{exec}}, q_{Op}, q_G, q_{AE}, q_{PKE})$ -adversary \mathcal{A} against the `S.sound-security` of the protocol Π , running in time t and executing q_{exec} sessions of the protocol, making at most q_{Op} queries to any operator, and making q_G, q_{AE}, q_{PKE} queries to respectively the functions G, AE and PKE . Denote the advantage of this adversary as $\text{Adv}_{\Pi}^{\text{S.sound}}(\mathcal{A})$. Then there exist a $(t' \approx O(t), q' = q_G + q_{\text{exec}})$ -mac-adversary \mathcal{A}_1 on G , $(t' \approx O(t), q' = q_G + 2 \cdot q_{\text{exec}} + 5 \cdot q_{Op})$ -prf-adversary \mathcal{A}_2 on G , a $(t' = O(t), q' = q_G + q_{\text{exec}})$ -ind-cpa-adversary adv_3 on G , a $(t' = O(t), q' = q_{\text{exec}} + q_{AE})$ -ae-adversary adv_4 on AE and a $(t' = O(t), q' = q_{\text{exec}} + q_{PKE})$ -ind-cca2-adversary \mathcal{A}_5 on PKE such that:*

$$\begin{aligned} \text{Adv}_{\Pi}^{\text{S.sound}}(\mathcal{A}) &\leq n_C \cdot \left(\frac{q_{\text{exec}}}{2^{|\text{MacS}|}} + \frac{1}{2^{\kappa}} + \text{Adv}_G^{\text{mac}}(\mathcal{A}_1) \right) \\ &\quad + \text{Adv}_G^{\text{prf}}(\mathcal{A}_2) + \text{Adv}_G^{\text{ind-cpa}}(\mathcal{A}_3) \\ &\quad + \text{Adv}_{AE}^{\text{ae}}(\mathcal{A}_4) + \text{Adv}_{PKE}^{\text{ind-cca2}}(\mathcal{A}_5). \end{aligned}$$

Proof. **Game \mathbb{G}_0 :** This game works as the game `S.sound-game` stipulated in our security model. The goal of this adversary \mathcal{A}_{G_0} is similar as the `S.Imp-game` but with a different adversary; indeed in the `S.Imp-game` is a `MiM` adversary and in the `S.sound-game`, we have a *legitimate-but-malicious* adversary.

Game \mathbb{G}_1 : We consider the game \mathbb{G}_1 as the `S.Imp-game` (as previously detailed) but including the specific query `Corrupt(P, ttype)`, i.e with the presence of operator keys corruption. We have used a such query in some previous security proofs. We proceed as follows: for any adversary \mathcal{A}_{G_0} winning the game \mathbb{G}_0 with a non-negligible success probability ϵ_{G_0} , we propose to construct a generic adversary \mathcal{A}_{G_1} winning the game \mathbb{G}_1 with a black-box access to the adversary \mathcal{A}_{G_0} .

Both adversaries play her related game with oracles. The following oracles are similar in the two games: `Send`, `CreateCl`, `NewInstance`, `Execute`, `Reveal`, and `StReveal`. So for each query related to these oracles from the adversary \mathcal{A}_{G_0} , the adversary \mathcal{A}_{G_1} forwards these queries to its own challenger and sends to \mathcal{A}_{G_0} the related answers. Now focus on the two last oracles which can be used by the adversary \mathcal{A}_{G_0} : `OpAccess` and `Corrupt`.

At first, we recall that the OpAccess in the game \mathbb{G}_0 takes in input a client identifier and outputs, for our protocol, an authentication vector composed by the tuple $AV = (R, Autn, Mac_C, CK, IK)$. Simulating the answer of the oracle OpAccess(C_i), the \mathcal{A}_{G_1} uses the query Execute(S, C_i) (with the server related to the *legitimate-but-malicious* adversary) and Reveal(C, i).

Now, focus on the simulation of the Corrupt answer. We recall that we have two possible inputs: a client or an operator. In the Corrupt oracle takes in input a client, the adversary \mathcal{A}_{G_1} uses its own Corrupt oracle to obtain the related answer. If the input is an operator, \mathcal{A}_{G_1} needs to forge the following values: the operator key sk_{op} , and for each client of this operator the tuple $(UID, sk_{UID}, st_{Op,C})$. To simulate a such answer, \mathcal{A}_{G_1} uses its specific Corrupt(C) and StReveal(C, $i, 1$) for each client Cof this operator.

So at this point, the adversary \mathcal{A}_{G_1} can simulate any query from the adversary \mathcal{A}_{G_0} . At the end of the simulation, the adversary \mathcal{A}_{G_1} replays the impersonation's attempt from the adversary \mathcal{A}_{G_0} . Thus, we have:

$$\Pr[\mathcal{A}_{G_0} \text{ wins}] \leq \Pr[\mathcal{A}_{G_1} \text{ wins}].$$

Winning game \mathbb{G}_1 : This game follows the game \mathbb{G}_1 described in the reduction proof of the theorem S.Imp. Thus, we have :

$$\begin{aligned} \text{Adv}_{\Pi}^{\text{S,sound}}(\mathcal{A}_{G_0}) &\leq n_C \cdot \left(\frac{q_{\text{exec}}}{2^{|\text{Mac}_S|}} + \frac{1}{2^\kappa} + \text{Adv}_G^{\text{mac}}(\mathcal{A}_1) \right. \\ &\quad \left. + \text{Adv}_G^{\text{prf}}(\mathcal{A}_2) + \text{Adv}_G^{\text{ind-cpa}}(\mathcal{A}_3) \right. \\ &\quad \left. + \text{Adv}_{\text{AE}}^{\text{ae}}(\mathcal{A}_4) + \text{Adv}_{\text{PKE}}^{\text{ind-cca2}}(\mathcal{A}_5) \right). \end{aligned}$$

□

4.2.5 Updated TUAK algorithms security

In order to prove the pseudorandomness, unforgeability and indistinguishability of our updated TUAK algorithms, we assume that the truncated keyed internal Keccak permutation is a good pseudorandom function. We propose two generic constructions to model the updated TUAK algorithms: a first one, denoted G_{tuak} when the secret is based on the subscriber key sk and a second one, denoted G_{tuak}^* when is only based on the operator key.

It is worth noting that the construction of the TUAK functions is reminiscent of the Merkle-Damgård construction, where the output of the function f is an input of the next iteration of the function f . This is in contradiction with the Sponge construction used in the hash function Keccak given the internal permutation f_{Keccak} .

We model the truncated keyed internal permutation of Keccak by the function f and f^* :

$$f(K, x \| y, i, j) = [f_{\text{Keccak}}(x \| K \| y)]_{i..j},$$

$$f^*(K^*, x^* \| y^*, i, j) = [f_{\text{Keccak}}(K^* \| x^* \| y^*)]_{i..j},$$

with $x \in \{0, 1\}^{512}$, $K, K^* \in \{0, 1\}^\kappa$, $y \in \{0, 1\}^{1088-\kappa}$, $x^* \in \{0, 1\}^{512+\kappa}$, $y^* \in \{0, 1\}^{1088}$ and $i, j \in \{0, 1\}^t$ with $\log_2(t-1) < 1600 \leq \log_2(t)$. We note that $\forall K, x, x^*, y, y^*, i, j$ such as $x = K^* \| x^*$ and $y^* = K \| y$, we have $f(K, x \| y, i, j) = f^*(K^*, x^* \| y^*, i, j)$. The input x (resp. x^*) can be viewed as the chaining variable of the cascade construction of G_{tuak} given f (resp. f^*), y (resp. y^*) is an auxiliary input of the function, and i and j define the size of the truncation. The construction G_{tuak} acts as a generalization of the specific TUAK algorithms:

$$\begin{aligned} \mathcal{F}_1(sk_{op}, sk_C, R, Sqn, Res_S, AMF) &= G_{\text{tuak}}(sk_C, inp_1, 0, 127) \\ &= G_{\text{tuak}}^*(sk_{op}, inp_1^*, 0, 127), \\ \mathcal{F}_2(sk_{op}, sk_C, R, Sqn, Res_S, AMF) &= G_{\text{tuak}}(sk_C, inp_1, 256, 383) \\ &= G_{\text{tuak}}^*(sk_{op}, inp_1^*, 256, 383), \\ \mathcal{F}_3(sk_{op}, sk_C, R, Sqn, Res_S, AMF) &= G_{\text{tuak}}(sk_C, inp_1, 512, 639) \\ &= G_{\text{tuak}}^*(sk_{op}, inp_1^*, 512, 639), \\ \mathcal{F}_4(sk_{op}, sk_C, R, Sqn, Res_S, AMF) &= G_{\text{tuak}}(sk_C, inp_1, 768, 895) \\ &= G_{\text{tuak}}^*(sk_{op}, inp_1^*, 768, 895), \\ \mathcal{F}_5(sk_{op}, sk_C, R, Res_S) &= G_{\text{tuak}}(sk_C, inp_2, 0, 47) \\ &= G_{\text{tuak}}^*(sk_C, inp_2^*, 0, 47), \end{aligned}$$

with:

$$\begin{aligned} inp_1 &= sk_{op} \| cst_1 \| cst_3, \quad inp_2 = sk_{op} \| cst_1 \| cst_3, \\ inp_1^* &= cst_1 \| keys \| cst_3, \quad inp_2^* = cst_1 \| keys \| cst_3, \\ cst_1 &= Inst \| AN \| 0^{192} \| (Inst' \| AN \| R \| AMF \| Sqn), \\ cst_3 &= Res_S \| Pad \| 1 \| 0^{192}, \end{aligned}$$

We define the cascade constructions G_{tuak} and G_{tuak}^* based on the function f and f^* as follows:

$$\begin{aligned} G_{\text{tuak}}(K, val, i, j) &= f(K, f(K, val_1 \| val_3, 0, 256) \| val_2 \| val_3, i, j), \\ G_{\text{tuak}}^*(K^*, val^*, i, j) &= f^*(f^*(K^*, val_1^* \| val_3^*, 0, 256) \| val_2^* \| val_3^*, i, j), \end{aligned}$$

with G_{tuak} and G_{tuak}^* from $\{0, 1\}^\kappa \times \{0, 1\}^d \times \{0, 1\}^t \times \{0, 1\}^t$ to $\{0, 1\}^n$, $val = (val_1 \| val_2) \| val_3 \in \{0, 1\}^{512} \times \{0, 1\}^{256} \times \{0, 1\}^{(832-\kappa)}$, $val^* = (val_1^* \| val_2^*) \| val_3^* \in \{0, 1\}^{256} \times \{0, 1\}^{256} \times \{0, 1\}^{(1088-\kappa)}$ two known values with $n = j - i$, $d = 1600 - \kappa$, $\kappa = |K|$ and $\log_2(t-1) < 1600 \leq \log_2(t)$, K a secret value and $0 \leq i \leq j \leq 1600$.

We express the required security properties of the generalization G_{tuak} (resp. G_{tuak}^*) under the prf-security of the function f (resp. f^*). Since the construction of the two functions, while we cannot prove the latter property, we can conjecture that the advantage of a prf-adversary would be of the form:

$$\text{Adv}_{f^*}^{\text{prf}}(\mathcal{A}) = \text{Adv}_f^{\text{prf}}(\mathcal{A}) \leq c_1 \cdot \frac{t/T_f}{2^{|K|}} + c_2 \cdot \frac{q \cdot t/T_f}{2^{1600-m}},$$

for any adversary \mathcal{A} running in time t and making at most q queries at its challenger. Here, m is the output's size of our function f and T_f is the time to do one f computation on the fixed RAM model of computation and c_1 and c_2 are two constants depending only on this model. In other words, we assume that the best attacks are either an exhaustive key search or a specific attack on this construction. This attack uses the fact that the permutation is public and can be easily inverted. Even if the protocol truncates the permutation, if the output values are large, and an exhaustive search on the missing bits is performed, it is possible to invert the permutation and recover the inputs. Since the secret keys is one of the inputs as well as some known values are also inputs, it is then possible to determine which guesses of the exhaustive search are correct guess or incorrect ones. Finally, if the known inputs are shorter than the truncation, false positives can happen due to collisions and we have to filter the bad guesses. However, if the number of queries is large enough, it is possible to filter these bad guesses and uniquely recover the keys.

Pseudorandomness and Unforgeability of TUAK algorithms. We begin by reducing the prf-security of G_{tuak} to the prf-security of the function f . This implies the mac-security of each TUAK algorithm. Recall that our main assumption is that the function f is prf-secure if the Keccak permutation is a good random permutation.

Theorem 7. [prf – security for G_{tuak}^* .] *Let $G_{\text{tuak}}^* : \{0, 1\}^\kappa \times \{0, 1\}^e \times \{0, 1\}^{d-e} \times \{0, 1\}^t \times \{0, 1\}^t \rightarrow \{0, 1\}^n$ and $f^* : \{0, 1\}^\kappa \times \{0, 1\}^e \times \{0, 1\}^d \times \{0, 1\}^t \times \{0, 1\}^t \rightarrow \{0, 1\}^m$ be the two functions specified above. Consider a (t, q) -adversary \mathcal{A} against the prf-security of the function G_{tuak}^* , running in time t and making at most q queries to its challenger. Denote the advantage of this adversary as $\text{Adv}_{G_{\text{tuak}}^*}^{\text{prf}}(\mathcal{A})$. Then there exists a $(t' \approx O(t), q' = q)$ -adversary \mathcal{A}' with an advantage $\text{Adv}_{f^*}^{\text{prf}}(\mathcal{A}')$ of winning against the pseudorandomness of f^* such that:*

$$\text{Adv}_{G_{\text{tuak}}^*}^{\text{prf}}(\mathcal{A}) = \text{Adv}_{f^*}^{\text{prf}}(\mathcal{A}'),$$

Proof. We construct the adversary \mathcal{A}_{f^*} using a prf-adversary \mathcal{A}_{G^*} . The latter uses \mathcal{A}_{f^*} as a challenger for a prf-game $\mathbb{G}_{\text{prf}(f^*)}$ and can only communicate with \mathcal{A}_{f^*} whereas \mathcal{A}_{f^*} has access to a challenger for f^* . To begin with, the challenger $\mathcal{C}_{f^*}^{\text{prf}}$ chooses a bit b and a private $\text{sk}_{\text{op}} \in \{0, 1\}^\kappa$. If $b = 0$, it assigns f^* to a random function and if $b = 1$, it assigns f^* to the specific internal function.

The adversary \mathcal{A}_{f^*} waits for queries from \mathcal{A}_{G^*} of the form (m, a, b) , with $m = m^{(1)} \| m^{(2)} \| m^{(3)} \in \{0, 1\}^d$, and $a, b \in \{0, 1\}^t$ and responds as follows:

- It queries its challenger $\mathcal{C}_{f^*}^{\text{prf}}$ for inputs $(m^{(1)} \| m^{(3)}, 0, 256)$ and receives the value Out_1 .

- Then, it computes $\text{Out}_2 = f^*(\text{Out}_1, m^{(2)} \| m^{(3)}, a, b)$.
- It returns the value Out_2 .

We note that the two first bullets permits to generate $G^*(\text{sk}_{\text{op}}, m, a, b) = \text{Out}_2$. This step is repeated up to a total of q queries from \mathcal{A}_{G^*} , with a and b fixed.

At some point, \mathcal{A}_{G^*} halts and outputs a guess d of the bit b . The prf-adversary \mathcal{A}_{f^*} chooses its guess b' as $b' = d$ and forwards it to $\mathcal{C}_{f^*}^{\text{prf}}$, which verifies if $b = b'$.

We analyze this simulation. Recall that the challenger responded either with a random value (if its internal bit b was set to 0) or with the output of the function $f^*(\text{sk}_{\text{op}}, m^{(1)} \| m^{(3)}, 0, 256)$ (if its internal bit was set as 1).

Thus, the output Out_2 matches either the output of a random function or the output of the function $G^*(\text{sk}, m, a, b)$. So the prf-adversary \mathcal{A}_{f^*} simulates perfectly a prf-challenger of G . Thus, we have:

$$\begin{aligned} \text{Adv}_{f^*}^{\text{prf}}(\mathcal{A}_{f^*}) &= \left| \Pr[\mathcal{A}_{f^*} \rightarrow 1 \mid b = 1] - \Pr[\mathcal{A}_{f^*} \rightarrow 1 \mid b = 0] \right| \\ &= \left| \Pr[b = b' \mid b = 1] - \Pr[b = b' \mid b = 0] \right| \\ &= \left| \Pr[b = d \mid b = 1] - \Pr[b = d \mid b = 0] \right| \\ &= \left| \Pr[d' = d \mid b = 1] - \Pr[d' = d \mid b = 0] \right| \\ &= \left| \Pr[\mathcal{A}_{G^*} \rightarrow 1 \mid b = 1] - \Pr[\mathcal{A}_{G^*} \rightarrow 1 \mid b = 0] \right| \\ &= \left| \text{Adv}_{G^*}^{\text{prf}}(\mathcal{A}_{G^*}) \right|. \end{aligned}$$

□

Theorem 8. [prf – security for G_{tuak} .] *Let $G_{\text{tuak}} : \{0, 1\}^\kappa \times \{0, 1\}^e \times \{0, 1\}^{d-e} \times \{0, 1\}^t \times \{0, 1\}^t \rightarrow \{0, 1\}^n$ and $f : \{0, 1\}^\kappa \times \{0, 1\}^e \times \{0, 1\}^d \times \{0, 1\}^t \times \{0, 1\}^t \rightarrow \{0, 1\}^m$ be the two functions specified above. Consider a (t, q) -adversary \mathcal{A} against the prf-security of the function G , running in time t and making at most q queries to its challenger. Denote the advantage of this adversary as $\text{Adv}_{G_{\text{tuak}}}^{\text{prf}}(\mathcal{A})$. Then there exists a $(t' \approx 2 \cdot t, q' = 2 \cdot q)$ -adversary \mathcal{A}' with an advantage $\text{Adv}_f^{\text{prf}}(\mathcal{A}')$ of winning against the pseudorandomness of f such that:*

$$\text{Adv}_{G_{\text{tuak}}}^{\text{prf}}(\mathcal{A}) = \text{Adv}_f^{\text{prf}}(\mathcal{A}'),$$

Proof. We construct the adversary \mathcal{A}_f using a prf-adversary \mathcal{A}_G . The latter uses \mathcal{A}_f as a challenger for a prf-game $\mathbb{G}_{\text{prf}(f)}$ and can only communicate with \mathcal{A}_f whereas \mathcal{A}_f has access to a challenger for f . To begin with, the challenger $\mathcal{C}_f^{\text{prf}}$ chooses a bit b and a private $\text{sk} \in \{0, 1\}^\kappa$. If $b = 0$, it assigns f to a random function and if $b = 1$, it assigns f to the specific internal function.

The adversary \mathcal{A}_f waits for queries from \mathcal{A}_G of the form (m, a, b) , with $m = m^{(1)} \| m^{(2)} \| m^{(3)} \in \{0, 1\}^d$, and $a, b \in \{0, 1\}^t$ and responds as follows:

- It queries its challenger $\mathcal{C}_f^{\text{prf}}$ for inputs $(m^{(1)} \| m^{(3)}, 0, 256)$ and receives the value Out_1 .

- Then, it queries C_f^{prf} for inputs $(\text{Out}_1 \| m^{(2)} \| m^{(3)}, a, b)$ and receives the value Out_2 .
- It returns the value Out_2 .

We note that the two first bullets permits to generate $G(\text{sk}, m, a, b)$ computing

$f(\text{sk}, f(\text{sk}, m^{(1)} \| m^{(3)}, 0, 256) \| m^{(2)} \| m^{(3)}, a, b)$. This step is repeated up to a total of q queries from \mathcal{A}_g , with a and b fixed.

At some point, \mathcal{A}_g halts and outputs a guess d of the bit b . The prf-adversary \mathcal{A}_f chooses its guess b' as $b' = d$ and forwards it to C_f^{prf} , which verifies if $b = b'$.

We analyze this simulation. Recall that the challenger responded either with a random value (if its internal bit b was set to 0) or with the output of the function $f(\text{sk}, m^{(1)} \| m^{(3)}, 0, 256)$ (if its internal bit was set as 1).

Thus, the output Out_2 matches either the output of a random function or the output of the function $G(\text{sk}, m, a, b)$. So the prf-adversary \mathcal{A}_f simulates perfectly a prf-challenger of G . Thus, we have:

$$\begin{aligned} \text{Adv}_f^{\text{prf}}(\mathcal{A}_f) &= \left| \Pr[\mathcal{A}_f \rightarrow 1 \mid b = 1] - \Pr[\mathcal{A}_f \rightarrow 1 \mid b = 0] \right| \\ &= \left| \Pr[b = b' \mid b = 1] - \Pr[b = b' \mid b = 0] \right| \\ &= \left| \Pr[b = d \mid b = 1] - \Pr[b = d \mid b = 0] \right| \\ &= \left| \Pr[d' = d \mid b = 1] - \Pr[d' = d \mid b = 0] \right| \\ &= \left| \Pr[\mathcal{A}_G \rightarrow 1 \mid b = 1] - \Pr[\mathcal{A}_G \rightarrow 1 \mid b = 0] \right| \\ &= \left| \text{Adv}_G^{\text{prf}}(\mathcal{A}') \right|. \end{aligned}$$

□

We use the generic result specified in A.1 to reduce the mac-secure to the prf-secure of the function G .

Theorem 9. [Mac – security for G_{tuak} .] *Let $G_{\text{tuak}} : \{0, 1\}^\kappa \times \{0, 1\}^e \times \{0, 1\}^{d-e} \times \{0, 1\}^t \times \{0, 1\}^t \rightarrow \{0, 1\}^n$ and $f : \{0, 1\}^\kappa \times \{0, 1\}^e \times \{0, 1\}^d \times \{0, 1\}^t \times \{0, 1\}^t \rightarrow \{0, 1\}^m$ be the two functions specified above. Consider a (t, q) -adversary \mathcal{A} against the Mac-security of the function G_{tuak} , running in time t and making at most q queries to its challenger. Denote the advantage of this adversary as $\text{Adv}_{G_{\text{tuak}}}^{\text{mac}}(\mathcal{A})$. Then there exists a $(t' \approx 2 \cdot (t + O(n + d)), q' = 2 \cdot q)$ -adversary \mathcal{A}' with an advantage $\text{Adv}_f^{\text{prf}}(\mathcal{A}')$ of winning against the pseudorandomness of f such that:*

$$\text{Adv}_{G_{\text{tuak}}}^{\text{mac}}(\mathcal{A}) \leq \text{Adv}_f^{\text{prf}}(\mathcal{A}') + \frac{1}{2^n}.$$

Indistinguishability of TUAK algorithms. We begin by reducing the ind – cpa-security of G_{tuak} to the prf-security of the function f . This implies the ind – cpa-security of each TUAK algorithm. Recall that our main assumption is that the function f is prf-secure if the Keccak permutation is a good random permutation.

Theorem 10. [ind – cpa – security of G_{tuak} .] *Let $G_{\text{tuak}} : \{0, 1\}^\kappa \times \{0, 1\}^e \times \{0, 1\}^{d-e} \times \{0, 1\}^t \times \{0, 1\}^t \rightarrow \{0, 1\}^n$ and $f : \{0, 1\}^\kappa \times \{0, 1\}^d \times \{0, 1\}^t \times \{0, 1\}^t \rightarrow \{0, 1\}^m$ be the two functions specified in Section B. Consider a (t, q) -adversary \mathcal{A} against the ind – cpa-security of the function G , running in time t and making at most q queries to its challenger. Denote the advantage of this adversary as $\text{Adv}_{G_{\text{tuak}}}^{\text{ind-cpa}}(\mathcal{A})$. Then there exists a $(t' \approx 2 \cdot t, q' = 2 \cdot q)$ -adversary \mathcal{A}' with an advantage $\text{Adv}_f^{\text{prf}}(\mathcal{A}')$ of winning against the pseudorandomness of f such that:*

$$\text{Adv}_{G_{\text{tuak}}}^{\text{ind-cpa}}(\mathcal{A}) = \text{Adv}_f^{\text{prf}}(\mathcal{A}'),$$

Proof. To prove the ind – cpa-security of G , we reduce this security to the prf-security of G which is defined in the proof of the theorem 9. We show that a prf-adversary \mathcal{A} of G can simulate the ind – cpa-challenger $C_G^{\text{ind-cpa}}$. A such prf-adversary behaves as follows. At first, the challenger C_G^{prf} chooses a private key K and one random bit $b \in \{0, 1\}$. If $b = 0$, he assigns f to a random function and if $b = 1$, he assigns f to the specific function G . For each query $(M^{\{i\}}, a, b)$ with a fixed (a, b) from the ind – cpa-adversary \mathcal{A}' to the prf-adversary \mathcal{A} the latter forwards each one to C_G^{prf} . The answer $f(K, M^{\{i\}}, a, b)$ is sent to \mathcal{A} which forwards it to \mathcal{A}' . Then, \mathcal{A}' sends a specific query containing two values (M_0, M_1) to \mathcal{A} . The latter chooses randomly a bit d and forwards (M_d, a, b) to the prf-challenger. As usual, this challenger sends $f(K, M_d, a, b)$ to \mathcal{A} which forwards it to \mathcal{A}' .

The goal of the ind – cpa-adversary is to find the bit d chosen by the \mathcal{A} . To do so, it can ask again some queries $(M^{\{i\}}, a, b)$ as previously. Finally, it sends its guessing d' to \mathcal{A} . Upon receiving this guessing, it chooses its guessing b' of b as follows: if $d = d'$, it chooses $b' = 1$, else $b' = 0$.

$$\begin{aligned} \text{Adv}_G^{\text{prf}}(\mathcal{A}) &= \left| \Pr[b = b' \mid f \stackrel{\$}{\leftarrow} G(K, \cdot, \cdot, \cdot), K \stackrel{\$}{\leftarrow} \{0, 1\}^\kappa] \right. \\ &\quad \left. - \Pr[b = b' \mid f \stackrel{\$}{\leftarrow} \mathcal{R}] \right| \\ &= \left| \Pr[b = b' \mid b = 1] - \Pr[b = b' \mid b = 0] \right| \\ &= \left| \Pr[b' = 1 \mid b = 1] - \Pr[b' = 0 \mid b = 0] \right| \\ &= \left| \Pr[d' = d \mid b = 1] - \Pr[d' \neq d \mid b = 0] \right| \\ &= \left| \text{Adv}_G^{\text{ind-cpa}}(\mathcal{A}') + \frac{1}{2} - \frac{1}{2} \right| \\ &= \left| \text{Adv}_G^{\text{ind-cpa}}(\mathcal{A}') \right| \end{aligned}$$

In the last equality, $\Pr[d' = d \mid b = 1]$ is the probability that the ind – cpa-adversary correctly guess the bit d which is $\text{Adv}_G^{\text{ind-cpa}}(\mathcal{A}') + 1/2$ and $\Pr[d' \neq d \mid b = 0]$ which is equal to $1/2$ since when $b = 0$, G is a random function, that is its output is chosen independently from its inputs. Consequently, it is not related to its inputs and the adversary cannot guess correctly the bit d . □

4.2.6 Updated MILENAGE algorithms security

In order to prove the unforgeability and indistinguishability properties of the MILENAGE algorithms, we assume that the AES permutation is a good pseudo-random function.

We model the AES algorithm by the function f :

$$f(K, x, y) = \text{AES}_K(x \oplus y),$$

with $x, y \in \{0, 1\}^{128}$, $K \in \{0, 1\}^\kappa$. Contrary to the TUAK algorithms, the MILENAGE algorithms do not behave as the same way but as two different ways. Let the construction G_{mil1} , the generalization of the functions $\mathcal{F}_1, \mathcal{F}_1^*, \mathcal{F}_2, \mathcal{F}_3, \mathcal{F}_4$ and G_{mil2} the generalization of the functions \mathcal{F}_5 and \mathcal{F}_5^* as follows:

$$\begin{aligned} G_{\text{mil1}}(\text{sk}_C, \text{inp}_1, 0, 127) &= \mathcal{F}_1(\text{sk}_{\text{op}}, \text{sk}_C, R, \text{Sqn}, \text{Res}_S, \text{AMF}), \\ G_{\text{mil1}}(\text{sk}_C, \text{inp}_2, 0, 127) &= \mathcal{F}_1^*(\text{sk}_{\text{op}}, \text{sk}_C, R, \text{Sqn}, \text{Res}_S, \text{AMF}), \\ G_{\text{mil1}}(\text{sk}_C, \text{inp}_3, 0, 47) &= \mathcal{F}_2(\text{sk}_{\text{op}}, \text{sk}_C, R, \text{Sqn}, \text{Res}_S, \text{AMF}), \\ G_{\text{mil1}}(\text{sk}_C, \text{inp}_4, 0, 127) &= \mathcal{F}_3(\text{sk}_{\text{op}}, \text{sk}_C, R, \text{Sqn}, \text{Res}_S, \text{AMF}), \\ G_{\text{mil1}}(\text{sk}_C, \text{inp}_5, 0, 127) &= \mathcal{F}_4(\text{sk}_{\text{op}}, \text{sk}_C, R, \text{Sqn}, \text{Res}_S, \text{AMF}), \\ G_{\text{mil2}}(\text{sk}_C, \text{inp}_6, 0, 47) &= \mathcal{F}_5(\text{sk}_{\text{op}}, \text{sk}_C, R, \text{Res}_S), \\ G_{\text{mil2}}(\text{sk}_C, \text{inp}_6, 64, 111) &= \mathcal{F}_5^*(\text{sk}_{\text{op}}, \text{sk}_C, R, \text{Res}_S), \end{aligned}$$

with:

$$\begin{aligned} \forall i \in \{1, \dots, 5\}, \text{inp}_i &= \text{sk}_{\text{op}} \parallel R \parallel (\text{Sqn} \parallel \text{AMF}) \parallel \text{Res}_S \parallel c_i \parallel r_i, \\ \text{inp}_6 &= \text{sk}_{\text{op}} \parallel R \parallel \text{Res}_S \parallel c_6 \parallel r_6, \end{aligned}$$

These both constructions are constructed as follows:

$$\begin{aligned} G_{\text{mil1}}(K, \text{val}^{(1)}, a, b) &= \lfloor \text{Top}_C \oplus f(K, \text{val}_4, f(K, \text{Top}_C, \text{val}_2)) \\ &\quad \oplus \text{Rot}_{\text{val}_5}(\text{Top}_C \oplus (\text{val}_3 \parallel \text{val}_3)) \rfloor_{a..b}, \\ G_{\text{mil2}}(K, \text{val}^{(2)}, a, b) &= \lfloor \text{Top}_C \oplus f(K, \text{val}_4, \text{Rot}_{\text{val}_5}(\text{Top}_C \\ &\quad \oplus f(K, \text{Top}_C, \text{val}_2))) \rfloor_{a..b}, \end{aligned}$$

with $G_{\text{mil1}} : \{0, 1\}^\kappa \times \{0, 1\}^{d_1} \times \{0, 1\}^t \times \{0, 1\}^t \rightarrow \{0, 1\}^n$,
 $G_{\text{mil2}} : \{0, 1\}^\kappa \times \{0, 1\}^{d_2} \times \{0, 1\}^t \times \{0, 1\}^t \rightarrow \{0, 1\}^n$,
and $\text{val}^{(1)} = \text{val}_1 \parallel \text{val}_2 \parallel \text{val}_3 \parallel \text{val}_4 \parallel \text{val}_5$, $\text{val}^{(2)} = \text{val}_1 \parallel \text{val}_2 \parallel \text{val}_4 \parallel \text{val}_5$, $\text{val}_1, \text{val}_2, \text{val}_4 \in \{0, 1\}^{128}$, $\text{val}_3 \in \{0, 1\}^{64}$, $\text{val}_5 \in \{0, 1\}^7$ and $\text{Top}_C = \text{val}_1 \oplus f(K, \text{val}_1, 0)$.

We express the security properties of the generalizations G_{mil1} and G_{mil2} under the prf-security of the function f . While we cannot prove the latter property, we can conjecture that the advantage of a prf-adversary would be of the form:

$$\text{Adv}_f^{\text{prf}}(\mathcal{A}) \leq c_1 \cdot \frac{t/T_f}{2^{128}} + c_2 \cdot \frac{q^2}{2^{128}},$$

for any adversary \mathcal{A} running in time t and making at most q queries at its challenger. Here, m is the output's size of our function f and T_f is the time to do one f computation on the fixed RAM model of computation and c_1 and c_2 are two constants depending only on this model. In other words, we assume that the best attacks are either a exhaustive key search or a linear cryptanalysis.

Pseudorandomness and Unforgeability of MILENAGE algorithms. We begin by reducing the prf-security of G_{mil1} and G_{mil2} to the prf-security of the function f . This implies the Mac-security of each MILENAGE algorithm.

Theorem 11. [prf – security for G_{mil1} and G_{mil2} .] *Let $G_{\text{mil1}}, G_{\text{mil2}} : \{0, 1\}^\kappa \times \{0, 1\}^d \times \{0, 1\}^t \times \{0, 1\}^t \rightarrow \{0, 1\}^n$ and $f : \{0, 1\}^\kappa \times \{0, 1\}^{128} \times \{0, 1\}^{128} \rightarrow \{0, 1\}^{128}$ be the two functions specified above. Consider a (t, q) -adversary \mathcal{A} against the prf-security of the function G_{mil1} (respectively G_{mil2}), running in time t and making at most q queries to its challenger. Denote the advantage of this adversary as $\text{Adv}_{G_{\text{mil1}}}^{\text{prf}}(\mathcal{A})$ (respectively $\text{Adv}_{G_{\text{mil2}}}^{\text{prf}}(\mathcal{A})$). Then there exists a $(t' \approx 3 \cdot t, q' = 3 \cdot q)$ -adversary \mathcal{A}' with an advantage $\text{Adv}_f^{\text{prf}}(\mathcal{A}')$ of winning against the pseudorandomness of f such that:*

$$\text{Adv}_{G_{\text{mil1}}}^{\text{prf}}(\mathcal{A}) = \text{Adv}_{G_{\text{mil2}}}^{\text{prf}}(\mathcal{A}) = \text{Adv}_f^{\text{prf}}(\mathcal{A}').$$

Proof. We construct the adversary \mathcal{A}_f using a prf-adversary \mathcal{A}_{G_1} (respectively \mathcal{A}_{G_2}). The latter uses \mathcal{A}_f as a challenger for a prf-game $\mathbb{G}_{\text{prf}(f)}$ and can only communicate with \mathcal{A}_f whereas \mathcal{A}_f has access to a challenger for f . To begin with, the challenger $\mathcal{C}_f^{\text{prf}}$ chooses a bit b and a private $\text{sk} \in \{0, 1\}^\kappa$. If $b = 0$, it assigns f to a random function and if $b = 1$, it assigns f to the specific internal function.

The adversary \mathcal{A}_f waits for queries from \mathcal{A}_G of the form (m, a, b) , with $m = m^{(1)} \parallel m^{(2)} \parallel m^{(3)} \parallel m^{(4)} \parallel m^{(5)} \in \{0, 1\}^d$, and $a, b \in \{0, 1\}^t$ and responds as follows:

- It queries its challenger $\mathcal{C}_f^{\text{prf}}$ for inputs $(m^{(1)}, 0^{128})$ and receives the value Out_1 .
- Then, it computes $\text{Top}_C = m^{(1)} \oplus \text{Out}_1$ and it queries $\mathcal{C}_f^{\text{prf}}$ for inputs $(\text{Out}_1, m^{(2)})$ and receives the value Out_2 .
- It queries $(m^{(4)}, \text{Out}_2 \oplus \text{rot}(\text{Out}_1 \oplus (m^{(3)} \parallel m^{(3)}), m^{(5)}))$ (respectively $(m^{(4)}, \text{rot}(\text{Out}_1 \oplus \text{Out}_2, m^{(5)}))$) and receives the value Out_3 .
- It returns the value $\lfloor \text{Out}_1 \oplus \text{Out}_3 \rfloor_{a..b}$.

This step is repeated up to a total of q queries from \mathcal{A}_{G_1} (respectively \mathcal{A}_{G_2}), with a and b fixed.

At some point, \mathcal{A}_{G_1} (respectively \mathcal{A}_{G_2}) halts and outputs a guess d of the bit b . The prf-adversary \mathcal{A}_f chooses its guess b' as $b' = d$ and forwards it to $\mathcal{C}_f^{\text{prf}}$, which verifies if $b = b'$.

We analyze this simulation. Recall that the challenger responded either with a random value (if its internal bit b was set to 0) or with the output of the function $f(\text{sk}, \dots)$ (if its internal bit was set as 1).

Thus, the output Out_3 matches either the output of the output of the function $G_1(\text{sk}, m, a, b)$ (respectively $G_2(\text{sk}, m, a, b)$) or a random function (indeed, the combination of two random functions by a boolean addition gives a random function). So the prf-adversary \mathcal{A}_f simulates perfectly

a prf-challenger of G_1 (respectively G_2). Thus, we have:

$$\begin{aligned}
\text{Adv}_f^{\text{prf}}(\mathcal{A}_f) &= \left| \Pr[\mathcal{A}_f \rightarrow 1 \mid b = 1] - \Pr[\mathcal{A}_f \rightarrow 1 \mid b = 0] \right| \\
&= \left| \Pr[b = b' \mid b = 1] - \Pr[b = b' \mid b = 0] \right| \\
&= \left| \Pr[b = d \mid b = 1] - \Pr[b = d \mid b = 0] \right| \\
&= \left| \Pr[d' = d \mid b = 1] - \Pr[d' = d \mid b = 0] \right| \\
&= \left| \Pr[\mathcal{A}_{G_1} \rightarrow 1 \mid b = 1] - \Pr[\mathcal{A}_{G_1} \rightarrow 1 \mid b = 0] \right| \\
&= \text{Adv}_{G_1}^{\text{prf}}(\mathcal{A}').
\end{aligned}$$

(similar computation for G_2). \square

We use the generic result specified in A.1 to reduce the mac-secure to the prf-secure of the function G_1 (resp. G_2).

Theorem 12. [Mac – security for $G_{\text{mil}1}$ and $G_{\text{mil}2}$.]

Let $G_{\text{mil}1}, G_{\text{mil}2} : \{0, 1\}^\kappa \times \{0, 1\}^d \times \{0, 1\}^t \times \{0, 1\}^t \rightarrow \{0, 1\}^n$ and $f : \{0, 1\}^\kappa \times \{0, 1\}^{128} \times \{0, 1\}^{128} \rightarrow \{0, 1\}^{128}$ be the two functions specified above. Consider a (t, q) -adversary \mathcal{A} against the Mac-security of the function $G_{\text{mil}1}$ (respectively $G_{\text{mil}2}$), running in time t and making at most q queries to its challenger. Denote the advantage of this adversary as $\text{Adv}_{G_{\text{mil}1}}^{\text{mac}}(\mathcal{A})$ (respectively $\text{Adv}_{G_{\text{mil}2}}^{\text{mac}}(\mathcal{A})$). Then there exists a $(t' \approx 3 \cdot (t + O(n + d)), q' = 3 \cdot q)$ -adversary \mathcal{A}' with an advantage $\text{Adv}_f^{\text{prf}}(\mathcal{A}')$ of winning against the pseudorandomness of f such that: $\text{Adv}_{G_{\text{mil}1}}^{\text{mac}}(\mathcal{A}) = \text{Adv}_{G_{\text{mil}2}}^{\text{mac}}(\mathcal{A}) \leq \text{Adv}_f^{\text{prf}}(\mathcal{A}') + \frac{1}{2^n}$.

Indistinguishability of MILENAGE algorithms. We begin by reducing the ind – cpa-security of $G_{\text{mil}1}$ and $G_{\text{mil}2}$ to the prf-security of the function f . This implies the prf-security of each MILENAGE algorithm.

Theorem 13. [ind – cpa – security of $G_{\text{mil}1}$ and $G_{\text{mil}2}$.]

Let $G_{\text{mil}1}, G_{\text{mil}2} : \{0, 1\}^\kappa \times \{0, 1\}^d \times \{0, 1\}^t \times \{0, 1\}^t \rightarrow \{0, 1\}^n$ and $f : \{0, 1\}^\kappa \times \{0, 1\}^{128} \times \{0, 1\}^{128} \rightarrow \{0, 1\}^{128}$ be the two functions specified in Section B. Consider a (t, q) -adversary \mathcal{A} against the ind – cpa-security of the function $G_{\text{mil}1}$ (resp. $G_{\text{mil}2}$), running in time t and making at most q queries to its challenger. Denote the advantage of this adversary as $\text{Adv}_{G_{\text{mil}1}}^{\text{ind-cpa}}(\mathcal{A})$ (resp. $\text{Adv}_{G_{\text{mil}2}}^{\text{ind-cpa}}(\mathcal{A})$). Then there exists a $(t' \approx 3 \cdot t, q' = 3 \cdot q)$ -adversary \mathcal{A}' with an advantage $\text{Adv}_f^{\text{prf}}(\mathcal{A}')$ of winning against the pseudorandomness of f such that:

$$\text{Adv}_{G_{\text{mil}1}}^{\text{ind-cpa}}(\mathcal{A}) = \text{Adv}_{G_{\text{mil}2}}^{\text{ind-cpa}}(\mathcal{A}) = \text{Adv}_f^{\text{prf}}(\mathcal{A}'),$$

Proof. To prove the ind – cpa-security of G_1 (resp. G_2), we reduce this security to the prf-security of G_1 (resp. G_2) and then to the prf-security of the function f using the theorem 11

. The first reduction follows the exact lines of the proof of the prf-security of the TUAK algorithms and we obtain:

$$\text{Adv}_{G_1}^{\text{ind-cpa}}(\mathcal{A}) = \text{Adv}_{G_2}^{\text{ind-cpa}}(\mathcal{A}) = \text{Adv}_f^{\text{prf}}(\mathcal{A}'),$$

\square

4.3 Narrow-Forward Privacy is Impossible

Our variant of AKA preserves the structure of the original protocol, but also provably attains wide-weak client unlinkability. In this section we show that this degree of client-unlinkability is optimal with respect to the structure of AKA. In particular, narrow-forward privacy is impossible.

Our result covers similar ground as that by Païse and Vaudenay [23], as we address protocols with mutual authentication. We extend the impossibility result to symmetric-key AKE protocols which also use public-key primitives. We also explain why the original impossibility result in [23] is imprecise, and presents some problems.

The result of [23]. Païse and Vaudenay showed an impossibility result for *authentication* protocols – rather than AKE; the extension from one environment to the other is, however, easy. In the terminology of our paper, [23] proved that server-authentication essentially precludes narrow-forward client-unlinkability. Their attack follows these steps: (1) the adversary \mathcal{A} creates two clients; (2) \mathcal{A} runs an honest protocol session between one of them (chosen uniformly at random depending on a secret bit b) and the server, but stops the last message from the server to the client; (3) \mathcal{A} corrupts both clients, learning their long-term state; (4) \mathcal{A} distinguishes between the clients by simulating the protocol with the intercepted message.

However, this attack makes a tacit assumption on the client’s behaviour, namely that if a session is aborted, the state is not updated or that it is updated in a consistent way, depending on the client’s internal state. Say that upon an abort, the client reverts to a random state; assuming that the adversary cannot access the very short time-frame in which the reversion to random is done, \mathcal{A} only gets the random state in response. Simulating the protocol with the received message will not match that state, thus reducing the adversary’s success probability to $\frac{1}{2}$.

Another way to bypass this result is to update the client state before the “last message” is sent to the client; if such an update is done at every execution, the attack presented in [23] fails. This is, however, a rather artificial twist: indeed, mutual authentication implies that the prover *must* somehow identify the server’s state as “valid” *before* it reaches a state which precludes it from verifying the server’s authentication.

Two new attacks. In the AKA protocol, it is the server which first authenticates to the client. The values used in authentication are the sequence number $Sqn_{Op,C}$ and the nonce R . The value $Sqn_{Op,C}$ is ephemeral, being updated at every session; however, it is a long-term state, and compatible with the client's own state Sqn_C . In particular, corrupting a client yields Sqn_C allowing the verifier to link the client with the corresponding $Sqn_{Op,C}$ value.

For a better comprehension of our attacks and their impact, we define the following notations: we divide a party's state (for both clients and operators) into a *static* state $stat.st_p$ and an ephemeral state $eph.st_p$. Thus, an operator's static state may contain operator-specific information, such as the secret key for a PKE scheme, but it will also include state shared with clients, i.e. $stat.st_{Op,C}$ for a client C . The same for the ephemeral state $eph.st_{Op,C}$ which for the AKA protocol consists of the sequence number $Sqn_{Op,C}$.

We propose the following attack:

- The adversary \mathcal{A} creates two clients C and C' with the same operator Op and the same location LAI .
- \mathcal{A} uses $DrawCl$ on C, C' and receives the handle VC , corresponding to either C (if the hidden bit $b = 0$), or C' (otherwise).
- \mathcal{A} runs an honest execution between the server S at LAI and the client VC until \mathcal{A} receives the message $R, Autn = (Sqn_{Op,C} \oplus AK) \parallel AMF \parallel Mac_S$ from the server. Denote $Autn[1] := Sqn_{Op,C} \oplus AK$, $Autn[2] = AMF$, and $Autn[3] := Mac_S$.
- \mathcal{A} corrupts C and learns Sqn_C, sk_C , and sk_{op} .
- By using the values R, sk_C , and sk_{op} , the adversary computes a value AK_C and retrieves $Sqn^* := Autn[1] \oplus AK_C$. Note that if $b = 0$, then $AK_C = AK$ and $Sqn^* = Sqn_{Op,C}$, while if $b = 1$, then $Sqn^* \neq Sqn_{Op,C'}$ with overwhelming probability.
- The adversary verifies Mac_S , on input (sk_C, sk_{op}, Sqn^*) . If this verification succeeds, the adversary outputs a guess $d = 0$ for the bit b ; else, it outputs 1.

For the analysis note that with overwhelming probability Mac_S will not verify if it was computed for C' , i.e. $f_1(sk_{C'}, sk_{op}, R, AMF, Sqn_{Op,C'}) \neq f_1(sk_C, sk_{op}, R, AMF, Sqn^*)$. The key vulnerability here is that, while $Sqn_{Op,C}$ is never sent in clear, the masking authentication key AK only depends on the client's static state. We also use the fact that the validity of the sequence number is confirmed by the value Mac_S .

While the latter factor, namely the validity of Mac_S certainly helps an attacker, our second attack (a variation of the first one) does not use the MAC value at all. The attack is run exactly in the same way, until we reach the final step. At that point:

- \mathcal{A} compares the obtained value Sqn^* with the recovered sequence number of C , namely Sqn_C , verifying if $|Sqn^* - Sqn_{Op,C}| \leq \Delta$. Note that in the actual attack presented above, the client's state Sqn_C should be exactly equal to the operator's state with respect to that client; however, our attack is even stronger in the sense that we do not need to control the executions of the protocol in order to obtain exact equality.

Analysis and Impact. Since the original AKA protocol is not even weak-client-unlinkable, it is not surprising that this protocol is not narrow-forward unlinkable either. However, the same attack works on our variant of the protocol and indeed, on any other extension or improvement of the original procedure which retains the characteristic of exchanging a message of the type $f(eph.st_{Op,C}, stat.st_{Op,C}, X)$ in the presence of a function $Match$; or exchanging that same message together with a message $g(eph.st_{Op,C}, stat.st_{Op,C}, Y)$, such that:

- f is reversible and takes as input $eph.st_{Op,C}, stat.st_{Op,C} = stat.st_C$, and a set X of publicly-known variables, giving arbitrary values in the set $\{0, 1\}^{*8}$;
- $Match$ takes as input two ephemeral state values $eph.st_{C'}$ and $eph.st_{Op,C}$ and it outputs a boolean value: 1 if $C = C'$ and 0 otherwise⁹;
- g takes as input the state values $eph.st_{Op,C}, stat.st_C$ and a set Y of public values, and which has the property that, for randomly chosen x and $stat.st_{C'}$ it holds that $g(x, stat.st_{Op,C'}, Y) \neq g(eph.st_{Op,C}, stat.st_{Op,C}, Y)$ ¹⁰.

5 Practical considerations

In this section, we discuss some of our design choices for the improvement we propose of the AKA protocol. We also provide a detailed analysis of our countermeasures and their intuitive effects in Appendix C.

As opposed to the proposal of van den Broek et al. [11], we opted to continue using (TMSI, LAI) tuples for the identification phase of the protocol. This infrastructure is maintained strictly by servers, with no operator contribution; thus it is efficient and inexpensive. Moreover, TMSI values and their correspondence to the client's IMSI is easy to find. In our

⁸ In our previous example, this is the string $Autn1$, which depends on $eph.st_{Op,C} = Sqn_{Op,C}$, on $stat.st_{Op,C} = (sk_C, sk_{op})$, and on the random value R which is public.

⁹ In our case, the $Match$ function returns 1 if and only if $|Sqn_{Op,C} - Sqn_{C'}| \leq \Delta$.

¹⁰ In our example, this function is f_1 , and the output value is Mac_S .

proposal, we bypass IMSI catchers attacks by never sending IMSIs in clear, and we add a symmetric authentication step in the encryption, thus precluding the client-unlinkability attack we found against the AKA variant of Arapinis et al. [21]. For the encryption, we use an IND-CCA public-key encryption scheme, we require a minimum PKI, only for operators. A client only stores the public key (and certificate) of the operator it subscribes to, thus minimizing key-management problems. In the TMSI reallocation step, we add an implicit authentication step, preventing Denial-of-Service-based linkability. We also add a freshness index, which prevents replays of challenges based on old sequence numbers.

We do specify, however, that our variant can only guarantee client-unlinkability if the *size* of the TMSI is equal to the length of the output of the PKE scheme. This is a non-trivial requirement, since servers are expected to keep track of all the TMSIs they issue; while using a shorter TMSI does not leak anything about the IMSI value, it does allow mass-surveillance organisms to track users down by distinguishing between the length of the encrypted IMSI as opposed to the TMSI length. On the positive side, servers may store TMSI values for a shorter while, since as soon as the user leaves the area, the TMSI is no longer useful.

Moreover, we recommend using a field of 32 bits for the index values id_{x_C} , $id_{x_{Op,C}}$. In fact, every time a session is aborted, the index(s) is (are) increased. The only way to replay a challenge is to previously drop 2^{32} successive authentication challenges, which is in our opinion hard to do. We require that the size of all the variables (except the network variables Op_{id} and Res_s) is: 96 bits for a 64-bit security bound, 128 bits for 96-bit security bound and 154 bits for 128-bit security bound.

We make the assumption that clients are aware of their current LAI, and thus avoid client-tracking by means of an itinerary. This is not a very strong assumption, since mobile devices are often equipped to detect their LAI. Finally, we bypass distinguishing attacks that exploit the re-synchronization phase by ensuring that sequence numbers cannot be desynchronized (and replays of challenges using old sequence numbers are prevented). Keeping in mind the practical requirement of minimizing the communication between servers and operators, our variant ensures that operators are contacted only in case the protocol is abnormally run or an adversary is detected. We also simplify the rather complex AKA structure, including only three communication phases rather than five. We depict our countermeasures and discuss them more in detail in Appendix C.

Finally, we require to restrict the batch of authentication vectors at only one vector if the last message (sent from the server to update the operator sequence number) can be dropped.

References

- [1] 3GPP. 3G Security; Specification of the MILENAGE algorithm set: An example algorithm set for the 3GPP authentication and key generation functions f_1 , f_1^* , f_2 , f_3 , f_4 , f_5 and f_5^* ; Document 2: Algorithm specification. TS 35.206, 3rd Generation Partnership Project (3GPP), June 2007.
- [2] 3GPP. 3G Security; Specification of the TUAK algorithm set: A 2nd example for the 3GPP authentication and key generation functions f_1 , f_1^* , f_2 , f_3 , f_4 , f_5 and f_5^* — Document 1: Algorithm specification. TS 35.231, 3rd Generation Partnership Project (3GPP), June 2013.
- [3] 3GPP. 3G Security; Technical Specification Group (TSG) SA; 3G Security; Security Architecture. TS 33.102, 3rd Generation Partnership Project (3GPP), June 2013.
- [4] 3GPP. 3rd Generation Partnership Project; Technical Specification Group Services and System Aspects; Security related network functions (Release 12). TS 43.020, 3rd Generation Partnership Project (3GPP), June 2014.
- [5] J. Alwen, M. Hirt, U. Maurer, A. Patra, and P. Raykov. Anonymous authentication with shared secrets. In *Proceedings of LatinCrypt*, volume 8895 of *LNCS*, pages 219–236. Springer-Verlag, 1999.
- [6] G. Ateniese, A. Herzberg, H. Krawczyk, and G. Tsudik. Untraceable mobility or how to travel *incognito*. In *Elsevier Computer Networks*, volume 31, pages 871–884. Elsevier, 1999.
- [7] C. Badertscher, C. Matt, U. Maurer, P. Rogaway, and B. Tackmann. Augmented secure channels and the goal of the TLS 1.3 record layer. In *Provable Security - 9th International Conference, ProvSec 2015, Kanazawa, Japan, November 24-26, 2015, Proceedings*, pages 85–104, 2015.
- [8] BSI. A Proposal for: Functionality classes for random number generators. AIS 20 / AIS 31. Version 2.0 , Bundesamt für Sicherheit in der Informationstechnik (BSI), 2011.
- [9] David A. McGrew and John Viega. The Security and Performance of the Galois/Counter Mode of Operation (Full Version). *IACR Cryptology ePrint Archive*, 2004:193, 2004.
- [10] D.Strobel. IMSI Catcher. In *2007, Seminar Work, Ruhr-Universität Bochum*, 2007.
- [11] Fabian van den Broek and Roel Verdult and Joeri de Ruiter. Defeating IMSI Catchers. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security, USA, October 12-6, 2015*, pages 340–351, 2015.
- [12] Guido Bertoni and Joan Daemen and Michaël Peeters and Gilles Van Assche. Keccak. In *Advances in Cryptology - EUROCRYPT 2013, 32nd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Athens, Greece, May 26-30, 2013. Proceedings*, pages 313–314, 2013.
- [13] Jens Hermans and Andreas Pashalidis and Frederik Vercauteren and Bart Preneel. A New RFID Privacy Model. In *Computer Security - ESORICS 2011 - 16th European Symposium on Research in Computer Security, Leuven, Belgium, September 12-14, 2011. Proceedings*, pages 568–587, 2011.
- [14] M. S. A. Khan and C. J. Mitchell. Another look at privacy threats in 3G mobile telephony. In *Proceedings of ACISP*, volume 8544 of *Lecture Notes in Computer Science*, pages

- 386–396. Springer, 2014.
- [15] Mihir Bellare and David Pointcheval and Phillip Rogaway. Authenticated Key Exchange Secure against Dictionary Attacks. In *Advances in Cryptology - EUROCRYPT 2000, International Conference on the Theory and Application of Cryptographic Techniques*, pages 139–155, 2000.
- [16] Mihir Bellare and Joe Kilian and Phillip Rogaway. The Security of the Cipher Block Chaining Message Authentication Code. *J. Comput. Syst. Sci.*, 61(3):362–399, 2000.
- [17] Ming-Feng Lee and Nigel P. Smart and Bogdan Warinschi and Gaven J. Watson. Anonymity guarantees of the UMTS/LTE authentication and connection protocol. *Int. J. Inf. Sec.*, 13(6):513–527, 2014.
- [18] Muxiang Zhang. Provably-Secure Enhancement on 3GPP Authentication and Key Agreement Protocol. *IACR Cryptology ePrint Archive*, 2003:92, 2003.
- [19] Muxiang Zhang and Yuguang Fang. Security analysis and enhancements of 3gpp authentication and key agreement protocol. *IEEE Transactions on Wireless Communications*, 4(2):734–742, 2005.
- [20] Myrto Arapinis and Loretta Ilaria Mancini and Eike Ritter and Mark Ryan. Privacy through Pseudonymity in Mobile Telephony Systems. In *21st Annual Network and Distributed System Security Symposium, NDSS*, 2014.
- [21] Myrto Arapinis and Loretta Ilaria Mancini and Eike Ritter and Mark Ryan and Nico Golde and Kevin Redon and Ravishankar Borgaonkar. New privacy issues in mobile telephony: fix and verification. In *the ACM Conference on Computer and Communications Security, CCS'12, Raleigh, NC, USA, October 16-18, 2012*, pages 205–216, 2012.
- [22] S. provider. Personal communication with one of europe's largest service providers, 2015.
- [23] Radu-Ioan Paice and Serge Vaudenay. Mutual Authentication in RFID: Security and Privacy. In *Proc. on the 3rd ACM Symposium on Information, Computer and Communications Security (ASIACCS)*, pages 292–299. ACM, 2008.
- [24] Serge Vaudenay. On Privacy Models for RFID. In *ASIACRYPT '07*, volume 4883, pages 68–87, 2007.
- [25] A. Shaik, R. Borgaonkar, N. Asokan, V. Niemi, and J.-P. Seifert. Practical attacks against privacy and availability in 4g/lte mobile communication systems. In *Proceedings of NDSS*. Internet Society, 2016.
- [26] Ulrike Meyer and Susanne Wetzels. A man-in-the-middle attack on UMTS. In *Proceedings of the 2004 ACM Workshop on Wireless Security, Philadelphia, PA, USA, October 1, 2004*, pages 90–97, 2004.
- [27] Victor Shoup. A Proposal for an ISO Standard for Public Key Encryption. *Cryptology ePrint Archive*, Report 2001/112, 2001. <http://eprint.iacr.org/>.
- [28] Zahra Ahmadian and Somayeh Salimi and Ahmad Salahi. New attacks on UMTS network access. In *2009 Wireless Telecommunications Symposium, WTS 2009, Prague, Czech Republic, April 22-24, 2009*, pages 1–6, 2009.

A Security notions

A.1 Security notions

The security notions can be proved under known or chosen message attacks, denoted respectively kma and cma . In this paper, we define the security notions under the chosen messages attacks.

Pseudo-random function. A pseudo-random function (prf) is a family of functions with the property that the input-output behavior of a random instance of the family is computationally indistinguishable from that of a random function. This property is defined in terms of the following security game \mathbb{G}^{prf} :

1. The challenger $\mathcal{C}_f^{\text{prf}}$ chooses a bit $b \in \{0, 1\}$. If $b = 0$, it assigns f to a random function $\text{Rand} : \{0, 1\}^d \rightarrow \{0, 1\}^n$. Else if $b = 1$, it chooses a key $K \in \{0, 1\}^\kappa$ and assigns f to the function $f(K, \cdot)$.
2. The adversary \mathcal{A} sends one by one q messages $x_i \in \{0, 1\}^d$ to the challenger and receives $f(x_i)$.
3. Finally, \mathcal{A} outputs a guess d of the bit b to the $\mathcal{C}_f^{\text{prf}}$.

We can evaluate the prf-advantage of an adversary against f , denoted $\text{Adv}_f^{\text{prf}}(\mathcal{A})$ as follows, for a random function denoted $\text{Rand} : \{0, 1\}^d \rightarrow \{0, 1\}^n$:

$$\text{Adv}_f^{\text{prf}}(\mathcal{A}) = \left| \Pr[\mathcal{A} \rightarrow 1 | f \xleftarrow{\$} F(K, \cdot), K \xleftarrow{\$} \{0, 1\}^\kappa] - \Pr[\mathcal{A} \rightarrow 1 | f \xleftarrow{\$} \text{Rand}] \right|,$$

Definition 11. (*[Pseudo-Random Function.]*) A family f of functions from $\{0, 1\}^\kappa \times \{0, 1\}^d$ to $\{0, 1\}^n$ is said to be (t, q) -prf-secure if any adversary \mathcal{A} running in time t and making at most q queries to its challenger $\mathcal{C}_f^{\text{prf}}$, cannot distinguish f from a random function Rand with a non-negligible advantage.

IND-CPA security. A scheme is considered secure in terms of indistinguishability (against chosen message attacks) if no adversary can learn any information on the input of the scheme given the output. This is formalized in terms of the following security game (denoted $\mathbb{G}^{\text{ind-cpa}}$):

1. The challenger $\mathcal{C}_f^{\text{ind-cpa}}$ chooses a key $K \in \{0, 1\}^\kappa$.
2. The adversary \mathcal{A} chooses and sends q_1 messages x to the challenger, which returns $f(K, x)$.
3. \mathcal{A} chooses two messages x_0 and x_1 and sends them to $\mathcal{C}_f^{\text{ind-cpa}}$.
4. The challenger chooses a bit b and returns (K, x_b) to the adversary.

5. Upon receiving $f(K, x_b)$, the adversary can adaptively query the challenger with at most q_2 messages $x \notin \{x_0, x_1\}$, to which the challenger responds with $f(K, x)$.
6. Finally, the adversary halts and outputs a guess d of the bit b used by the challenger.

Definition 12. (*[General security of ind-cpa-scheme.]*) A scheme f is considered as ind-cpa-secure by indistinguishability if any adversary \mathcal{A} running in time t and making at most q queries to its challenger $\mathcal{C}_f^{\text{ind-cpa}}$, given an encryption of a message randomly chosen from a two-element message space, cannot distinguish efficiently the encryption of one of both messages. We can evaluate the ind-cpa-advantage of a such adversary, denoted $\text{Adv}_f^{\text{ind-cpa}}(\mathcal{A})$:

$$\begin{aligned} \text{Adv}_f^{\text{ind-cpa}}(\mathcal{A}) &= \Pr[\mathcal{A} \rightarrow 1 | b \leftarrow \{0, 1\}, \\ K &\leftarrow \{0, 1\}^\kappa, (x_0, x_1) \in \{0, 1\}^d \times \{0, 1\}^d, \\ \mathcal{C} &\rightarrow f(K, x_b)]. \end{aligned}$$

MAC security. A scheme is considered secure in terms of unforgeability (against chosen message attacks) if no adversary can forge an acceptable message authentication code for any chosen message. This is formalized in terms of the following security game (denoted \mathbb{G}^{Mac}):

1. The challenger $\mathcal{C}_f^{\text{Mac}}$ chooses a key $K \in \{0, 1\}^\kappa$.
2. The adversary \mathcal{A} chooses and sends q messages x to the challenger, which returns $f(K, x)$.
3. The adversary tries to forge an acceptable output σ for a chosen input x which was no a previous query and sends the couple (x, σ) to the challenger.
4. The challenger verifies if the message x was not previously requested and if $f(K, x) = \sigma$. The adversary wins this game if the both conditions are accepted. Otherwise, the adversary loses the game.

We define the adversary's advantage with respect to this game as:

$$\text{Adv}_f^{\text{mac}}(\mathcal{A}) = \Pr[\mathcal{A} \text{ wins}].$$

Definition 13. (*[General security of mac-scheme.]*) A MAC construction (f, ϑ) is considered as mac-secure by unforgeability if any adversary \mathcal{A} running in time t and making at most q queries to its challenger $\mathcal{C}_f^{\text{mac}}$, cannot forge efficiently any couple (message, mac) no-forged by $\mathcal{C}_f^{\text{mac}}$. We can evaluate the mac-advantage We can evaluate the ind-cpa-advantage of a such adversary, denoted $\text{Adv}_f^{\text{mac}}(\mathcal{A})$:

$$\begin{aligned} \text{Adv}_f^{\text{mac}}(\mathcal{A}) &= \Pr[\vartheta(x_{q+1}, \tau_{q+1}) = 1 | K \xleftarrow{\$} \{0, 1\}^\kappa, \\ \forall i \in \{1, \dots, q\}, \mathcal{C} &\rightarrow (x_i, \tau_i), \mathcal{A} \rightarrow (x_{q+1}, \tau_{q+1})]. \end{aligned}$$

Generic Results. As mentioned in [16], pseudorandom functions make good message authentication codes. The authors have determined the exact security of a such reduction by the following proposal:

Proposition 1. *Let $f : \{0, 1\}^\kappa * \{0, 1\}^d \rightarrow \{0, 1\}^s$ be a family of functions. Consider a (t, q) -adversary \mathcal{A} against the prf-security of the function f , running in time t and making at most q queries to its challenger. Denote the advantage of a such adversary $\text{Adv}_f^{\text{prf}}(\mathcal{A})$. then, there are a $(t' \sim t + O(s + d), q' = q)$ -adversary \mathcal{A}' against the mac-security of the function f with an advantage $\text{Adv}_f^{\text{mac}}(\mathcal{A}')$ such as:*

$$\text{Adv}_f^{\text{mac}}(\mathcal{A}) \leq \text{Adv}_f^{\text{prf}}(\mathcal{A}') + \frac{1}{2^s}.$$

IND-CCA2 security. A cryptosystem is indistinguishable under an adaptive chosen cyphertext attack if every probabilistic polynomial time adversary has only a negligible advantage over random guessing, i.e it wins the above game with a probability $1/2 + \epsilon(k)$, where $\epsilon(k)$ is a negligible function in the security parameter k . Although the adversary knows x_0, x_1 and pke the probabilistic nature of the encryption function means that the encryption of x_b will be only one of many valid ciphertexts, and therefore encrypting x_0, x_1 , and comparing the resulting ciphertexts with the challenge ciphertext does not afford any advantage to the adversary. The game $\mathbb{G}^{\text{ind-cca2}}$ behaves as follows:

1. The challenger $\mathcal{C}_f^{\text{ind-cca2}}$ chooses a key pair (pke, ske) and returns pke.
2. The adversary \mathcal{A} may perform q_1 number of encryption oracle or decryption oracle to the challenger.
3. \mathcal{A} submits two distinct chosen messages x_0 and x_1 and sends them to $\mathcal{C}_f^{\text{ind-cca2}}$.
4. The challenger uniformly chooses a bit $b \in \{0, 1\}$, and returns $f(\text{pke}, x_b)$ to the adversary.
5. Upon receiving $C = f(\text{pke}, x_b)$, the adversary may make further calls to the decryption oracle, but may not submit the challenge ciphertext C .
6. Finally, the adversary halts and outputs a guess d of the bit b used by the challenger.

Definition 14. (*[General security of ind-cca2-scheme.]*) A scheme f is considered as ind-cca2-secure by indistinguishability under an adaptative chosen ciphertext attack if any adversary \mathcal{A} running in time t and making at most q queries to its challenger $\mathcal{C}_f^{\text{ind-cca2}}$, given an encryption of a message randomly chosen from a two-element message space, cannot distinguish efficiently the encryption of one of both messages. We can evaluate the ind-cca2-advantage of a such adversary, denoted $\text{Adv}_f^{\text{ind-cca2}}(\mathcal{A})$:

$$\begin{aligned} \text{Adv}(\mathcal{A}_f^{\text{ind-cca2}}) &= \Pr[\text{Adv}(\mathcal{A}_f^{\text{ind-cca2}}) \rightarrow 1 | b \leftarrow \{0, 1\}, \\ &(\text{pke}, \text{ske}) \leftarrow \{0, 1\}^k, (x_0, x_1) \in \{0, 1\}^d \times \{0, 1\}^d, \\ &C \rightarrow f(\text{pke}, x_b)]. \end{aligned}$$

AE security. An AEAD-scheme $\Pi = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ is *ae-secure* if every probabilistic polynomial time adversary has only a negligible advantage of the following indistinguishability game using the all-in-one formulation from [7]:

Real $_{\Pi}$	Ideal $_{\Pi}$
Initialization $K \leftarrow \mathcal{K}$	Initialization $K \leftarrow \mathcal{K}$
<hr/> Oracle AE.Enc Input: $(N, A, M) \in (\mathcal{N}, \mathcal{A}, \mathcal{M})$ $C \leftarrow \mathcal{E}(K, N, A, M)$ return C	<hr/> Oracle AE.Enc Input: $(N, A, M) \in (\mathcal{N}, \mathcal{A}, \mathcal{M})$ $C' \leftarrow \mathcal{E}(K, N, A, M)$ $C \leftarrow \{0, 1\}^{ C' }$ return C
<hr/> Oracle AE.Dec Input: $(N, A, C) \in (\mathcal{N}, \mathcal{A}, \mathcal{C})$ $M \leftarrow \mathcal{D}(K, N, A, C)$ return M	<hr/> Oracle AE.Dec Input: $(N, A, C) \in (\mathcal{N}, \mathcal{A}, \mathcal{C})$ return \perp

Fig. 9. Real and Ideal security game for AEAD-schemes.

1. The challenger $\mathcal{C}_{\Pi}^{\text{ae}}$ chooses a key $K \leftarrow \mathcal{K}$ and a bit $b \in \{0, 1\}$. If $b = 0$, it uses the encryption and decryption functions of the model Real_{Π} . Otherwise, it uses the functions of the model Ideal_{Π} . These models are described in the figure 9.
2. The adversary \mathcal{A} may perform q queries \mathcal{E} and \mathcal{D} such as for any tuples (N, A, M, A', M') , if the query $C = \mathcal{E}(K, N, A, M)$ has been required, then the two queries $\mathcal{E}(K, N, A', M')$ and $\mathcal{D}(N, A, C)$ cannot be required.
3. The adversary \mathcal{A} outputs its guess d of the bit b used by the challenger.

Definition 15. (*[General security of ae-scheme.]*) An AEAD-scheme Π is a triple of algorithms $\Pi = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ where \mathcal{K} is a randomized algorithm that samples a key $K \in \{0, 1\}^*$, \mathcal{E} is a deterministic algorithm that maps the key K , a nonce N , additional data A , and a message M to a ciphertext C , and \mathcal{D} is a deterministic algorithm that (K, N, A, C) to M . We assume that for all $K, N, A, M \in \{0, 1\}^*$, we have $\mathcal{D}(K, N, A, \mathcal{E}(K, N, A, M)) = M$. We consider that a such scheme is *ae-secure* if any adversary \mathcal{A} running in time t and making at most q queries to its challenger $\mathcal{C}_{\Pi}^{\text{ae}}$ cannot success with a non-negligible advantage. A such advantage, denoted $\text{Adv}_{\Pi}^{\text{ae}}(\mathcal{A})$, is defined as follows:

$$\text{Adv}_{\Pi}^{\text{ae}}(\mathcal{A}) = |\Pr(\mathcal{A} \rightarrow 1 | \text{Real}_{\Pi}) - \Pr(\mathcal{A} \rightarrow 1 | \text{Ideal}_{\Pi})|.$$

B Full protocol description

In the AKA protocol [3], mutual client-backend authentication is provided using Message Authentication Codes (MAC) computed by three of the TUAK algorithms, while the secret keys are derived from a random value and a shared secret key with a key derivation function (KDF), by means of the rest of the TUAK functions.

The basic framework is a challenge-response stateful protocol between two main actors: the HLR (Home Network Register) and the ME/USIM (Mobile Equipment/User Subscriber Identity Module). This protocol needs an intermediate entity, the VLR (Visited Network Register), as specified in Section 3.2. Both the ME/USIM and the HLR keep track of counters, denoted respectively Sqn_C and Sqn_{HLR} ; these sequence numbers are meant to provide entropy and enable network authentication (from HLR to ME/USIM). Technically, one can view the user's sequence number as an increasing counter, while the latter keeps track of the highest authenticated counter the user has accepted.

The AKA protocol uses a set of seven functions: $\mathcal{F}_1, \mathcal{F}_2, \mathcal{F}_3, \mathcal{F}_4, \mathcal{F}_5, \mathcal{F}_1^*, \mathcal{F}_5^*$. The first two are used to authenticate a MAC answer, proving that both participants know the same subscriber key sk_C and the same operator key sk_{op} . Algorithm \mathcal{F}_1 is called the network authentication function. As its name implies, it allows the subscriber to authenticate the network. Furthermore, this function provides the data integrity used to derive keys (in particular authenticating the random, session-specific value R). Algorithm \mathcal{F}_2 is called the subscriber authentication function, and it allows the network to authenticate the subscriber C by proving that the entity owns the subscriber key sk_C and the operator key sk_{op} .

The following three algorithms, $\mathcal{F}_3, \dots, \mathcal{F}_5$, are used as key derivation functions, outputting respectively a cipher key (CK), an integrity key (IK), and an anonymity key (AK), all derived on input the subscriber key sk_C , the operator key sk_{op} , and the session-specific random value R. Notice that the master key sk_C is only known by HLR and ME/USIM, but not by the intermediate entity VLR.

The last key, AK, is used to mask the sequence number Sqn , but it is not part of the session keys. Its function is to blind the value of Sqn since the latter may leak some information about the subscriber. In order to ensure that no long-term desynchronization occurs, the AKA protocol provides a re-synchronization procedure between the two participants, in which the user forces a new sequence number on the back-

end home network, using the \mathcal{F}_1^* and \mathcal{F}_5^* to authenticate this value much in the same way that the server has authenticated its own sequence number and random value. Figure 10 details the challenge-response of AKA procedure.

The operator key.. Subscribers to the same operator all share the operator's own secret key, in practice a 256-bit integer. This value is not directly stored on the phone, but rather an intermediate value, obtained by running the internal Keccak permutation on input sk_{op} and several constants, is embedded in the SIM card. Thus, whereas this value enters in all future runs of the cryptographic algorithms, it is never stored in clear on the user's mobile.

Identification. Globally, the procedure starts with the identification of the ME/USIM to VLR when the user equipment switches on. At first, the mobile equipment receives a user request from the VLR and then responds in cleartext, with a temporary identity, called TMSI, which is known by the ME/USIM and VLR. A TMSI is a local number and its construction is specified in the technical specification 23.003 [ref]. This value can be used only in a specific given area: the TMSI is always accompanied by the Location Area Identification (LAI) to avoid ambiguities. The VLR (globally the network) manages suitable data based which keeps the relation between the IMSIs and the TMSIs. A new TMSI must be allocated at least in each location updating procedure, i.e when you use a TMSI we need to replace it by a fresh value. So the mobile station de-allocates the old value and allocates the fresh temporary identity from the VLR. We note that this value need to be store in a non volatile memory with its LAI in order to they are not lost when the ME/USIM switches off. When the ME/USIM receives a *user identity request* from the VLR, it sends its $TMSI_0 || LAI$. When it receives its value, the ME/USIM verifies if the LAI matches the current ME/USIM. If it is not the case, it starts a *Local TMSI Unknown Procedure*. Otherwise, it tries to recover the corresponding permanent identity using its suitable data based to accept the identification. If a such value cannot be recover, it sends a *Permanent Identity Request* to the ME/USIM which answers with its IMSI. All these flows are exchanged in cleartext.

Local TMSI Unknown Procedure: As we said previously, if the LAI does not match the VLR, a *Local TMSI Unknown Procedure* has been engaged. Globally, it asks to the previous VLR₀ the relation with its TMSI. Then, either its received the corresponding IMSI or a "error message" implying a *Permanent Identity Request* to the mobile station.

Permanent Identity Request: In some cases, notably when the user cannot be identified with any temporary identity, the identification of a user on the radio path by means of the permanent subscriber identity, called IMSI. This basic procedure is sum up in the following figure:

The request and its answer are sent in cleartext. The procedure is used too when the user registers for the first time in a serving network.

Challenge-Response. After receiving the IMSI the HLR generates a fresh sequence number Sqn and an unpredictable variable R. By using the subscriber's key sk_C and the corresponding operator key sk_{op} , it then generates a list of n unique authentication vectors AV composed of five strings: R, Mac_C, CK, IK, Autn. For every authentication vector, the sequence number is updated. The update procedure depends on the chosen method. The specifications feature a first method which does not take into account the notion of time, and which basically increments by 1 the most significant 32-first value of the sequence number. A second and third subsequent methods feature a time-based sequence number update based on a clock giving universal time [3]. The authentication vector is generated as follows:

$$\begin{aligned} Mac_S &\leftarrow \mathcal{F}_1(sk_C, sk_{op}, R, Sqn, AMF), \\ Mac_C &\leftarrow \mathcal{F}_2(sk_C, sk_{op}, R), \\ CK &\leftarrow \mathcal{F}_3(sk_C, sk_{op}, R), \\ IK &\leftarrow \mathcal{F}_4(sk_C, sk_{op}, R), \\ AK &\leftarrow \mathcal{F}_5(sk_C, sk_{op}, R), \\ Autn &\leftarrow (Sqn \oplus AK) || AMF || Mac_S, \end{aligned}$$

where Mac_S is the message authentication code of the network by the subscriber, Mac_C is the message authentication code of the subscriber by the network and AMF the authentication and key management field (which is a known, public constant).

The HLR sends the list of the authentication vectors AV to the VLR. This list may also contain only a single authentication vector. Upon the reception and storage of these vectors, when the VLR initiates an authentication and key agreement, it selects the next authentication vector from the ordered array and stores Mac_C and the session keys CK and IK. Then, it forwards (R, Autn) to ME/USIM.

The ME/USIM verifies the freshness of the received authentication token. To this end, it recovers the sequence number by computing the anonymity key AK which in its own turn depends on three values: sk_C , sk_{op} , and the received R. Then, the user verifies the received Mac_S computing $\mathcal{F}_1(sk_C, sk_{op}, R, Sqn, AMF)$ with the received value R and the Sqn. If they are different, the user sends *authentication failure* message back to the VLR and the user abandons the procedure. In case the execution is not aborted, the ME/USIM verifies if the received Sqn value is in a correct range relatively to a stored

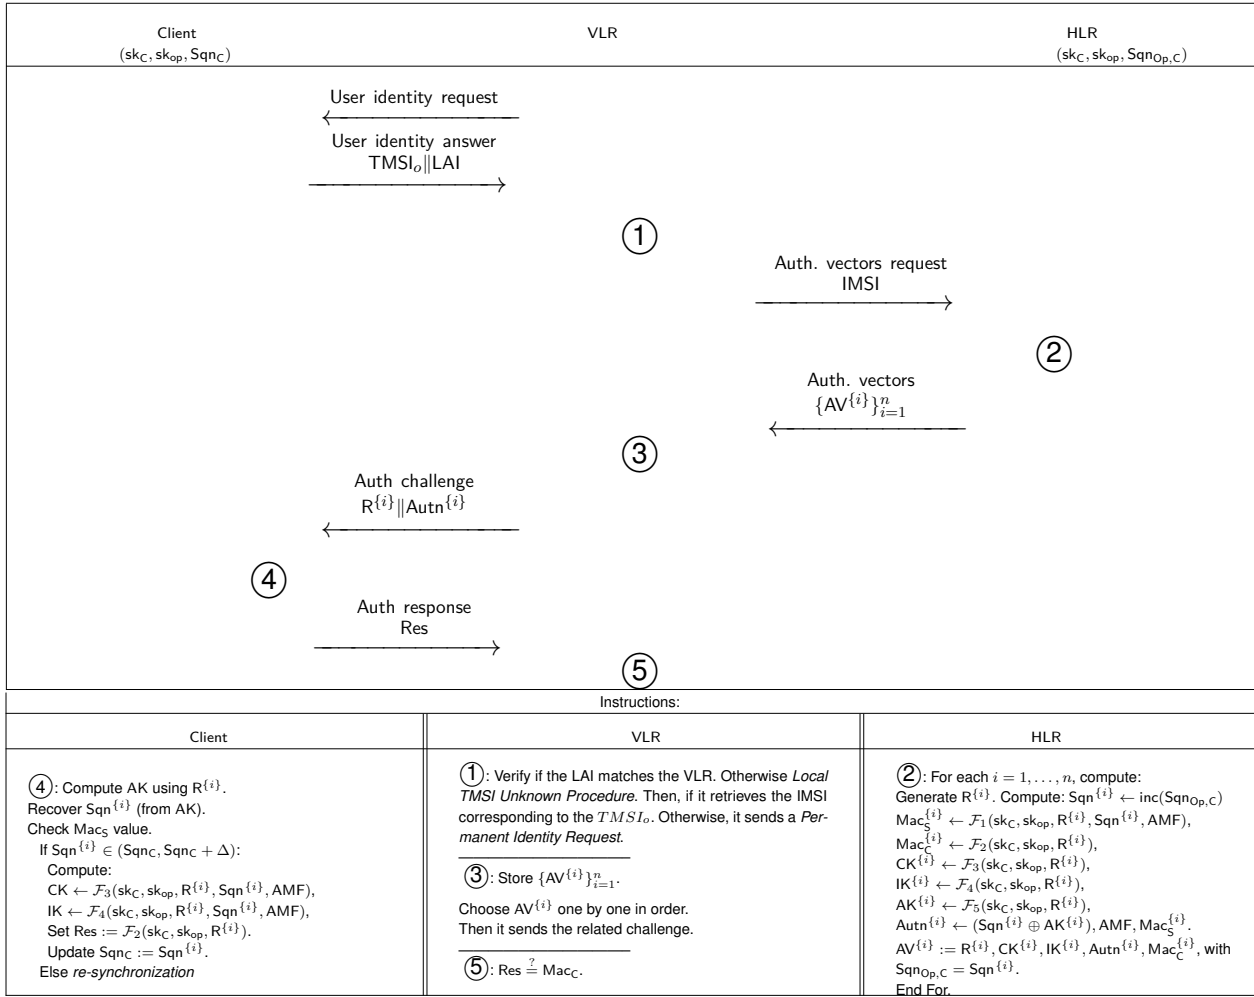


Fig. 10. The AKA Procedure.




Fig. 11. Local TMSI Unknown Procedure




Fig. 12. Identification by the permanent identity.

value Sqn_C^{11} . If the Sqn is out of range, the user sends a *syn-*

¹¹ The sequence number Sqn is considered to be in the correct range relatively to Sqn_C if and only if $Sqn \in (Sqn_C, Sqn_C + \Delta)$, where Δ is defined by the operator.

chronization failure message back to the VLR, which triggers a *re-synchronization procedure*, depicted further in Figure 13.

The Mac_S value does not only ensure integrity, but also the authentication of the network by ME/USIM. If the two previous verifications are successful i.e if the received authentication token is fresh, the network is authenticated by the ME/USIM. Then, the ME/USIM computes CK , IK and $Res \leftarrow \mathcal{F}_2(sk_C, sk_{op}, R)$. To improve efficiency, Res , CK , and IK could also be computed earlier, at the same time that AK is computed. Finally, the user sends Res to VLR. If $Res = Mac_C$, the VLR successfully authenticates the ME/USIM. Otherwise, the VLR will initiate an *authentication failure* report procedure with the HLR. Note that the verification of the sequence number by the ME/USIM will cause the rejection of any attempt to re-use an authentication token more than once.

Re-synchronizing. The re-synchronization procedure is used when the subscriber detects that the received sequence number is not in the correct range, but that it has been correctly authenticated. The single goal of this procedure is the re-

initialization of the sequence number, and does not imply immediately any mutual authentication or key agreement (rather it triggers a new authentication attempt).

Indeed, the ME/USIM sends an synchronization failure message, consisting of a parameter $Auts$, with

$$Auts = (Sqn_C \oplus AK^*) \parallel Mac^*,$$

where the key is computed as $AK^* = \mathcal{F}_5^*(sk_{op}, sk_C, R)$ and $Mac^* = \mathcal{F}_1^*(sk_{op}, sk_C, R, Sqn_C, AMF)$.

The \mathcal{F}_1^* algorithm is a MAC function with the additional property that no valuable information can be inferred from Mac^* (in particular this function acts as a PRF). Though similar to \mathcal{F}_1 , the \mathcal{F}_1^* algorithm is designed so that the value $Auts$ cannot be replayed relying on the output of \mathcal{F}_1 . Furthermore, the anonymity key generated by the client in the resynchronization is obtained via the \mathcal{F}_5^* algorithm rather than by \mathcal{F}_5 , even if the same random value R is used.

Upon receiving a re-synchronization failure message, the VLR does not immediately send a new user authentication request to the ME/USIM, but rather notifies the HLR of the re-synchronization failure, sending the parameter $Auts$ and the session-specific R . When the HLR receives this answer, it creates a new batch of authentication vectors. Depending on whether the retrieved, authenticated Sqn indicates that the HLR's sequence number is out of range or not, the back-end home network either starts from the last authenticated sequence number, or updates the latter to the user's sequence number.

More precisely, the HLR retrieves the Sqn_C by computing $\mathcal{F}_5^*(sk_C, sk_{op}, R) \oplus [Auts]_{48}$. Then, it verifies if the incremented Sqn_{HLR} is in the correct range relatively to Sqn_C . If the Sqn_{HLR} verifies this property, it sends a new list of authentication data vectors initiated with Sqn_{HLR} else HLR verifies the value of Mac^* . If this step is successful, it resets the value of Sqn_{HLR} to $Sqn_{HLR} := Sqn_C$ and sends a new list of authentication data vectors initiated with this updated Sqn_{HLR} . This list may also contain only a single authentication vector. Figure 13 details this re-synchronization procedure.

Re-allocation of the TMSI. At the end of the key derivation, the both entities (ME/USIM and VLR) need to be update with a fresh value. To allocate a fresh value, denoted $TMSI_n$, the VLR generates this value in the same LAI¹². Only after a successful identification based on the old $TMSI_o$ and a successful key setting permitting to share the ciphering key CK (by the AKA protocol). The VLR forges the new $TMSI_n$ and sends

it to the ME/USIM in a ciphered mode by A5/3 algorithm (globally an encryption with KASUMI more details in annex C and section 4.2 in TS 43.020) based the derived key CK . The ME/USIM recovers with the key CK the new $TMSI_n$ and stored it and de-allocate $TMSI_o$. Then the ME/USIM sends an "acknowledge message" in cleartext to prevent its allocation. After receiving this message, the VLR de-allocates the $TMSI_o$ and stores $TMSI_n$. If the VLR does not receive a such message, the network shall maintain the association between the old TMSI and the IMSI and between the new TMSI and the IMSI. For the next identification, the mobile station can used the both TMSI ($TMSI_o$ and $TMSI_n$). This will allow the network to determine the TMSI stored in the ME/USIM; the association between the other TMSI and the IMSI shall then be deleted, to allow the unused TMSI to be allocated to another ME/USIM.

Milenage algorithms: MILENAGE [1] is a set of algorithms which aims to achieve authentication and key generation properties. As opposed to TUAK which is based on Keccak's internal permutation, the MILENAGE algorithms are based on the Advanced Standard Protocol (AES).

The functions \mathcal{F}_1^* and \mathcal{F}_2^* must provide authentication while the functions \mathcal{F}_3^* , \mathcal{F}_4^* and \mathcal{F}_5^* are used to derive key material in order to achieve confidentiality, integrity and anonymity. The different parameters of these functions are:

- Inputs: sk_{op} a 128-bit long term credential key that is fixed by the operator, a 128-bit random value R , a 48-bit sequence number Sqn and a 16-bit authentication field management AMF chosen by the operator (the last two values are only used for the MAC generation). We denote that the subscriber key sk_{op} is a private key shared by all the subscriber of the same operator. Consequently, we do not consider sk_{op} as a private key.
- A 128-bit subscriber key sk_C shared out of band between the HLR and ME/USIM.
- Five 128-bit constants c_1, c_2, c_3, c_4, c_5 which are Xored onto intermediate variables and are defined as follows:
 - $c_1[i] = 0, \forall i \in \{0, 127\}$.
 - $c_2[i] = 0, \forall i \in \{0, 127\}$, except that $c_2[127] = 1$.
 - $c_3[i] = 0, \forall i \in \{0, 127\}$, except that $c_3[126] = 1$.
 - $c_4[i] = 0, \forall i \in \{0, 127\}$, except that $c_4[125] = 1$.
 - $c_5[i] = 0, \forall i \in \{0, 127\}$, except that $c_5[124] = 1$.
- Five integers r_1, r_2, r_3, r_4, r_5 in the range $\{0, 127\}$ which define amounts by which intermediate variables are cyclically rotated and are defined as follows: $r_1 = 64; r_2 = 0; r_3 = 32; r_4 = 64; r_5 = 96$.

The generation of MAC's or derived key starts similarly by initializing a value Top_C . To do so, one applies a first called of the well-known function AES on inputs the operator and

¹² They are not recommended methods to generate the TMSI as the spec 23.003 ("the structure and coding of it can be chosen by agreement between operator and manufacturer in order to meet local needs"). They provide only one advice (some parts of the TMSI may be related to the time)

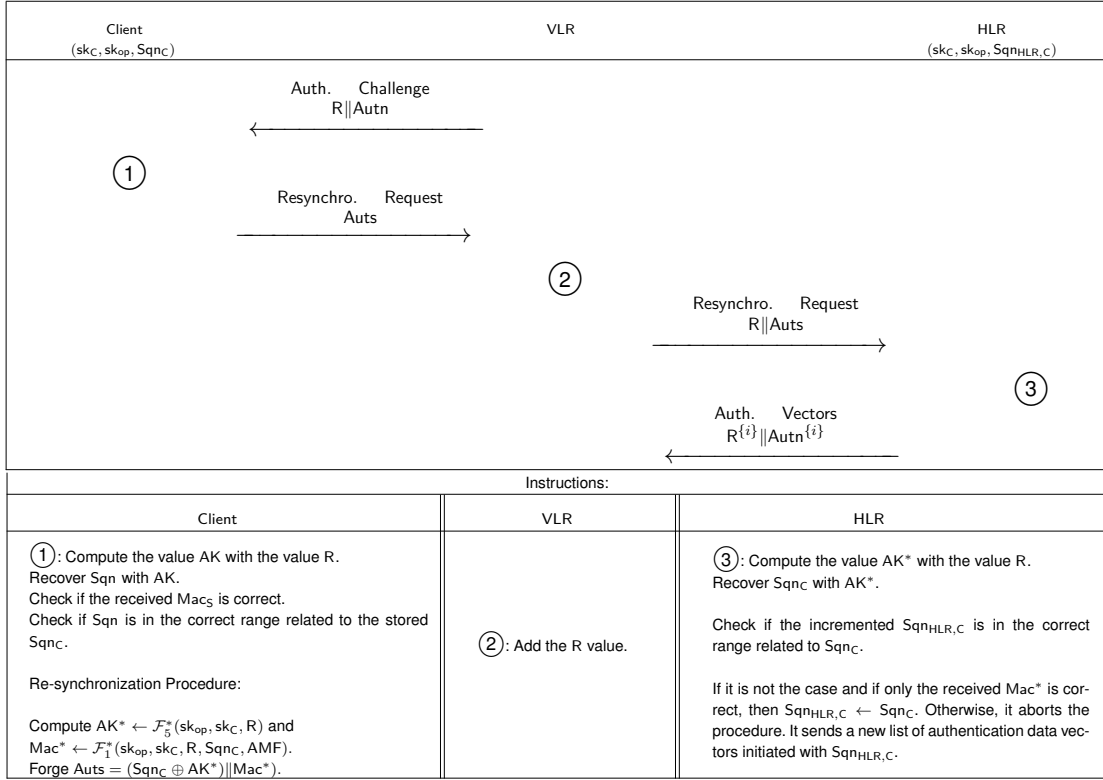


Fig. 13. The re-synchronization procedure of AKA protocol.

subscriber keys such as:

$$Top_C = sk_{op} \oplus AES_{sk_C}(sk_{op})$$

. We recall that, $AES_K(M)$ denotes the result of applying the Advanced Encryption Standard encryption algorithm to the 128-bit value M under the 128-bit key K . Then, we compute the following values taking as input Sqn, R, AMF and others constants:

- $Temp = AES_{sk_C}(R \oplus Top_C)$,
- $Out_1 = AES_{sk_C}(Temp \oplus Rot_{r_1}(Sqn || AMF || Sqn || AMF) \oplus c_1) \oplus Top_C$,
- $Out_2 = AES_{sk_C}(Rot_{r_2}(Temp \oplus Top_C) \oplus c_2) \oplus Top_C$,
- $Out_3 = AES_{sk_C}(Rot_{r_3}(Temp \oplus Top_C) \oplus c_3) \oplus Top_C$,
- $Out_4 = AES_{sk_C}(Rot_{r_4}(Temp \oplus Top_C) \oplus c_4) \oplus Top_C$,
- $Out_5 = AES_{sk_C}(Rot_{r_5}(Temp \oplus Top_C, r_5) \oplus c_5) \oplus Top_C$.

All the outputs of the MILENAGE algorithms are computed as follows:

- **Output \mathcal{F}_1 :** $Mac_C = \lfloor Out_1 \rfloor_{0..63}$,
- **Output \mathcal{F}_1^* :** $Mac^* = \lfloor Out_1 \rfloor_{64..127}$,
- **Output \mathcal{F}_2 :** $Mac_S = \lfloor Out_2 \rfloor_{64..127}$,
- **Output \mathcal{F}_3 :** $CK = Out_3$,
- **Output \mathcal{F}_4 :** $IK = Out_4$,
- **Output \mathcal{F}_5 :** $AK = \lfloor Out_5 \rfloor_{0..47}$,
- **Output \mathcal{F}_5^* :** $AK^* = \lfloor Out_5 \rfloor_{0..47}$,

This is also described in figure 14

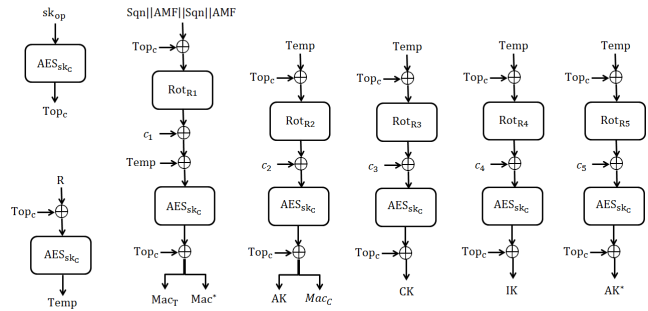


Fig. 14. MILENAGE diagram.

TUAK algorithms:

TUAK [2] is a set of algorithms which aims to achieve secure mutual authentication and key generation properties. The TUAK algorithms are based on the TUAK permutation, which in turn relies on a truncation of the internal permutation function of Keccak. Moreover, for efficiency reasons,

only one or two iterations of the internal TUAK permutation, $f_{\text{Keccak}}[1600]$, is used.

The functions \mathcal{F}_1 (respectively \mathcal{F}_1^*) and \mathcal{F}_2 must provide authentication while the functions \mathcal{F}_3 , \mathcal{F}_4 and \mathcal{F}_5 (respectively \mathcal{F}_5^*) are used to derive key material in order to achieve confidentiality, integrity and anonymity. The different parameters of these functions are:

- Inputs: sk_{op} a 256-bit long term credential key that is fixed by the operator, a 128-bit random value R , a 48-bit sequence number Sqn , and a 16-bit authentication $fic^{1,4}$ management AMF chosen by the operator (the last t values are only used for the MAC generation). We denote that the subscriber key sk_C is a private key shared by the subscriber of the same operator.
- A subscriber key sk_C shared out of band between the HI and ME/USIM allows to initialize the value Key:
 - If $|sk_C| = 128$ bits, then $Key \leftarrow sk_C[127..0] || 0^{128}$.
 - If $|sk_C| = 256$ bits, then $Key \leftarrow sk_C[255..0]$.
- Several public constants:
 - AN: a fixed 56-bit value $0x5455414B312E30$.
 - Inst and Inst' are fixed binary variables of 8 bits, specified in [2] and different depends on the functions and the output sizes.

The generation of MAC's or derived key starts similarly by initializing a value Top_C . To do so, one applies a first f_{Keccak} permutation on a 1600-bit state Val_1 as follows:

$$Val_1 = sk_{op} || Inst || AN || 0^{192} || Key || Pad || 1 || 0^{512},$$

where Pad is a bitstring output by a padding function. The value Top_C corresponds at the first 256 bits of this output.

At this point, the behavior of the functions \mathcal{F}_1 and \mathcal{F}_1^* diverges from that of the other functions. To generate the MAC value of \mathcal{F}_1 and \mathcal{F}_1^* , we take as input Sqn , AMF and R , three values chosen by the HLR and some constants. After the generation of Top_C , we initialize a second state, namely,

$$Val_2 = Top_C || Inst' || AN || R || AMF || Sqn || Key || Pad || 1 || 0^{512}.$$

Then, one applies the TUAK permutation on Val_2 . Next only the first 64 bits are used to compute the Mac_C value. To generate derived keys and the response of \mathcal{F}_2 , one initializes a second state for this function, too, namely,

$$Val_2 = Top_C || Inst' || AN || R || 0^{64} || Key || Pad || 1 || 0^{512}.$$

Then, the TUAK permutation is applied on Val_2 and one obtains the value Out. Finally, one derives the response Mac_C and the derived keys from the resulting Out:

$$\begin{aligned} Mac_C &= [Out]_{|\ell|-1..0}, \ell \in \{16, 32, 64, 128\}, \\ CK &= [Out]_{256..384} \text{ and } |CK| = 128, \\ IK &= [Out]_{512..640} \text{ and } |IK| = 128, \\ AK &= [Out]_{768..816} \text{ and } |AK| = 48. \end{aligned}$$

This is also depicted in Figure 15.




Fig. 15. TUAK diagram.

The way the output of the functions is truncated and used is the reason why TUAK is called a *multi-output function*. This notion is one of the main differences with MILENAGE and has a no-negligible impact on the efficiency of TUAK, as it saves a few calls of the internal function. However, this multi-output function property can be an issue for the security of the master key. Indeed, during one session we can have four calls of the same function with similar inputs but with a different truncation. Having different chunks of the same global 1600-state (called Out in our description) can permit to recover the long-term key sk_C by the reversibility of the TUAK permutation. The union of all the different chunks provided during one session, gives at most only 432 bits on the 1600 bits. Thus, having multiple outputs may be hazardous in general, the Keccak based construction of TUAK allows this without compromising the long-term parameters.

C Evaluation

In proposing our variant of AKA we explicitly or implicitly addressed several attacks. We discuss these below, referring the reader to Figure 16 for a better overview.

Server Corruptions : The original AKA protocol only offers a degree of key-indistinguishability and impersonation security, only in the absence of server corruptions. Since servers are trusted to run the authenticated key-exchange step, corrupting a server compromises any security of a channel this server establishes; however, in the AKA routine, this flaw is exacerbated, since the corruption results can be re-used later in non-vulnerable areas. This is an

Added countermeasures	Cost	Attacks it Prevents	Attack Impact
Client sends encrypted IMSI	- Needs IND-CCA PKE encryption - Simple PKI (only operators)	Client Confidentiality: (IMSI Catchers)	Trace many users - Parallelizable - Passive/Active
	Large TMSI size	Client unlinkability: Distinguish TMSI/IMSI msg.	Trace 1 user: - Non-parallelizable - Active only
Authenticate TMSI reallocation (see also: index)	New reallocation alg.	Client unlinkability: (Denial-of-Service)	Trace many users - Parallelizable - Active only
		Client unlinkability: Distinguish TMSI/IMSI	Trace 1 user - Non-parallelizable - Active only
Index idx_C , idx_S	New 1-bit state variable	Client unlinkability: Prompt resynch, distinguish	Trace 1 user - Non-parallelizable - Active only
		S.Imp-resistance: Challenge is un-replayable	Impersonate servers - Parallelizable - Active only
Introducing Res_S	- New server identifier - Changed crypto algs.	S.Imp-resistance k.ind-security S.sound-security (Server Corruptions)	Break sec. channel - Parallelizable - Needs corruptions - Active only
Use only current LAI	- Clients must know LAI - Clients store Res_S	Location privacy: (Track past LAI)	Trace 1 user per LAI - Non-Parallelizable - Passive

Fig. 16. Assessment of our AKA variant: cost and effect of countermeasures.

active, and rather complex attack, but it is highly parallelizable and has a great security impact. To mitigate this risk, we added a server-specific, unique, publicly-known identifier Res_S , which is now given as input to all the cryptographic functions.

Client Confidentiality : IMSI catcher attacks compromise the client's identity in a direct way (the adversary learns a static identifier). This attack can be run (with a reduced success probability) even by passive attackers, and it is highly parallelizable. The consequence is that multiple clients can be tracked simultaneously in a mass-surveillance operation. We mitigate such risks by ensuring that no static identifier is leaked through, by using a PKE scheme, in which only the operators have secret and public keys.

Client Unlinkability : Even if the adversary cannot track a user back to a permanent identifier, she can still try to distinguish between two chosen users, e.g. by causing some unusual protocol steps. Attacks like distinguishing between two different failure messages (which are actively triggered by the adversary), injecting a message and then seeing its effect in a protocol run (which is accepted if the chosen client is compatible with the injected message, and rejected otherwise), or distinguishing between messages of distinct lengths allow client linkability. While not as versatile, nor as parallelizable as client confidentiality attacks, these threats nevertheless allow an insidious adversary to track a user that is singled-out for mass surveillance. In our variant, we make protocol executions for different users indistinguishable from one another, at the cost

of larger TMSI values, a new index variable, using IND-CCA PKE encryption, and making the operator intervene in the case of an error.

Denial of Service : Apart from being a mean of breaking client-unlinkability, DoS attacks can also facilitate IMSI catchers, and add to the complexity of the AKA procedure. One way of causing a DoS in the original protocol is to send a random string as a replacement for the TMSI reallocation message. The client will parse this as a different TMSI than the intended one, and thus the server will need to request the user's IMSI in clear. We mitigate DoS attacks by using authenticated encryption for the TMSI reallocation and ensuring that no desynchronizations can occur.

Itinerary tracking : One disadvantage of AKA is that the client's past location is revealed during the protocol, allowing to track up to 1 user per LAI at any one time. We bypass this difficulty by only using current LAI values.