

# Improving Air Interface User Privacy in Mobile Telephony

Mohammed Shafiul Alam Khan\* and Chris J Mitchell

Information Security Group, Royal Holloway, University of London  
Egham, Surrey TW20 0EX, United Kingdom  
shafiulalam@gmail.com, me@chrismitchell.net

**Abstract.** Although the security properties of 3G and 4G mobile networks have significantly improved by comparison with 2G (GSM), significant shortcomings remain with respect to user privacy. A number of possible modifications to 2G, 3G and 4G protocols have been proposed designed to provide greater user privacy; however, they all require significant alterations to the existing deployed infrastructures, which are almost certainly impractical to achieve in practice. In this article we propose an approach which does not require any changes to the existing deployed network infrastructures, i.e. to the serving networks or the mobile devices, but offers improved user identity protection over the air interface. The proposed scheme makes use of multiple IMSIs for an individual USIM to offer a degree of pseudonymity for a user. The only changes required are to the operation of the authentication centre in the home network and to the USIM, both owned by a single entity in the mobile telephony system. The scheme could be deployed immediately since it is completely transparent to the existing mobile telephony infrastructure. We present two different approaches to the use and management of multiple IMSIs, and report on experiments to validate its deployability.

**Keywords:** Multiple IMSIs USIM, pseudonymity, mobile telephony, user privacy

## 1 Introduction

While the first generation (1G) mobile telephony systems did not provide any security features, security has been an integral part of such systems since the second generation (2G). For example, GSM, perhaps the best known 2G system, provides a range of security features, including authentication of the mobile user to the network, data confidentiality across the air interface, and a degree of user pseudonymity through the use of temporary identities. Third and fourth generation (3G and 4G) systems, such as UMTS/3GPP and Long-Term Evolution (LTE), have enhanced these security features, notably by providing mutual authentication between network and phone, and integrity protection for signalling

---

\* The author is a Commonwealth Scholar, funded by the UK government.

commands sent across the air interface. However, user privacy protection has remained largely unchanged, relying in all cases on the use of temporary identities [18, 29], and it has long been known that the existing measures do not provide complete protection for the user identity [11, 17]. The discussion below applies equally to 2G, 3G and 4G systems, although we use 3G terminology throughout.

The problem of user identity privacy in mobile networks is more than two decades old. Samfat et al. [25] first addressed the conflicting requirements of untraceability and disclosure of identity during authentication in a mobile network. The user privacy issue has been discussed extensively in the literature [6–9, 20–23, 31, 32], and many modifications to existing protocols have been proposed to avoid the problem [6, 8, 9, 20, 22, 32]. All these proposals involve making major modifications to the air interface protocol, which would require changes to the operation of all the serving networks as well as all the deployed phones. It seems likely that making the necessary major modifications to the operation of the air interface after deployment is essentially infeasible. Many of the proposed schemes also involve the use of public key cryptography [6, 9, 32], which has a high computational cost, although there do exist schemes which only use symmetric cryptography [8, 20]. It would therefore be extremely valuable if a scheme offering greater user privacy could be devised which did not involve making significant changes to the existing mobile telecommunications infrastructures, and had minimal computational cost. This motivates the work described in this paper. In sum, the contributions of the paper are as follows.

- We propose a new approach to the use and management of multiple IMSIs in a USIM to enhance user pseudonymity in mobile telephony systems.
- We have implemented key parts of the scheme, verifying its feasibility.
- We provide a privacy and functional analysis of the scheme.

The remainder of the paper is structured as follows. In section 2 key terminology for and features of mobile telephony systems are briefly reviewed. Threats to user privacy addressed in this paper are then summarised in section 3. In section 4 our threat model is presented. Section 5 outlines a novel approach to improving air interface user privacy using multiple IMSIs. Sections 6 and 7 provide descriptions of two proposed approaches to the use and management of multiple IMSIs in a USIM. Results from our experimental evaluation are presented in section 8. An analysis of the proposed approaches is presented in section 9. Section 10 provides a brief discussion of related work. Finally, conclusions are drawn and possible directions for future work are considered in section 11.

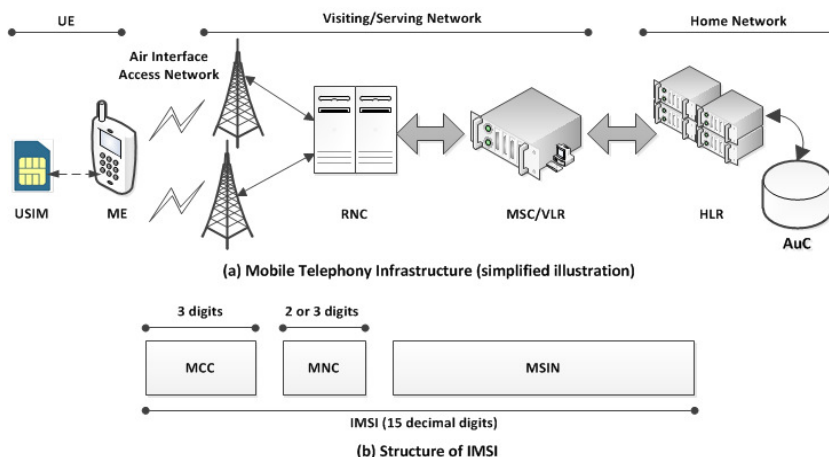
## 2 Background

### 2.1 Mobile Telephony Systems

We start by providing a brief overview of key terminology for mobile systems. A complete mobile phone is referred to as a *user equipment (UE)*, where the term encapsulates not only the *mobile equipment (ME)*, i.e. the phone, but also the

*user subscriber identity module (USIM)* within it [4], where the USIM takes the form of a cut-down smart card. The USIM embodies the relationship between the human user and the issuing *home network*, including the *International Mobile Subscriber Identity (IMSI)*, the telephone number of the UE, and other user (subscriber) data, together with a secret key shared with the issuing network which forms the basis for all the air interface security features. The USIM data storage capabilities are specified in section 10.1 of 3GPP TS 121 111 [13]. Information held within the USIM is stored in files, which can be divided into the following categories: *application dedicated files (ADFs)*, *dedicated files (DFs)* and *elementary files (EFs)* [14]. Most of the subscriber information is stored in EFs, which are the files we focus on in this paper.

To attach to a mobile network, a UE connects via its radio interface to a radio tower. Several radio towers are controlled by a single *radio network controller (RNC)* which is connected to one *mobile switching center/visitor location register (MSC/VLR)*. The MSC/VLR is responsible for controlling call setup and routing. Each MSC/VLR is also connected to the carrier network's *home location register (HLR)* where corresponding subscriber details can be found. The HLR is associated with an *authentication center (AuC)* that stores cryptographic credentials required for communicating with the USIM. The RNC and the MSC/VLR are part of the *visiting/serving network* whereas the HLR and the AuC are the *home network* component (see Fig. 1(a)).



**Fig. 1.** Mobile telephony systems

To access mobile network services, a UE needs to complete mutual authentication as soon as it attaches to a network. Mutual authentication is performed using the *Authentication and Key Agreement (AKA)* protocol, described in detail below. If mutual authentication is successful, the MSC informs the HLR, which associates the UE's IMSI with the address of the MSC. The MSC also

assigns a *temporary mobile subscriber identity (TMSI)* and sends the TMSI to the UE in encrypted form. The TMSI is unique to the location area in which the subscriber is currently located. Accordingly, whenever the subscriber visits a new location area, the MSC must update the TMSI value.

An IMSI is a 15-digit decimal number (see Fig. 1(b)). Of the 15 digits, the first three form the *mobile country code (MCC)*. The next two or three digits identify the network operator, and are known as the *mobile network code (MNC)*. The length of the MNC, i.e. whether it contains two or three digits, is a national matter. The remaining nine or ten digits, known as the *mobile subscriber identification number (MSIN)*, are administered by the relevant operator in accordance with the national policy [3, 28]. IMSIs therefore have geographical significance, and their use is typically managed by the network operator in blocks. The combination of the MCC and the MNC can be used to uniquely identify the home network of the IMSI. The MSIN is used by the operator to identify the subscriber for billing and other operational purposes. Each IMSI uniquely identifies the mobile user, as well as the user's home network and home country. The IMSI is stored in the USIM and is normally fixed. The elementary file  $EF_{IMSI}$  contains the value of the IMSI.

## 2.2 Proactive UICC

*Proactive UICC* is a service operating across the USIM-ME interface that provides a mechanism for a USIM to initiate an action to be taken by the ME. It forms part of the USIM application toolkit [15]. The 2G predecessor of the USIM, known as the SIM, supports a similar feature, known as *proactive SIM*, part of the SIM application toolkit.

ETSI TS 102 221 [14] specifies that the ME must communicate with the USIM using either the T=0 or T=1 protocol, specified in ISO/IEC 7816-3 [19]. In both cases the ME is always the *master* and thus initiates commands to the USIM; as a result there is no mechanism for the USIM to initiate communications with the ME. This limits the possibility of introducing new USIM features requiring the support of the ME, as the ME needs to know in advance what actions it should take. The proactive UICC service provides a mechanism that allows the USIM to indicate to the ME, using a response to an ME-issued command, that it has some information to send. The USIM achieves this by including a special status byte in the response application protocol data unit. The ME is then required to issue the *FETCH* command to find out what the information is [16]. To avoid cross-phase compatibility problems, this service is only permitted to be used between a proactive UICC and an ME that supports the proactive UICC feature. The fact that an ME supports proactive UICC is revealed when it sends a *TERMINAL PROFILE* command during UICC initialization.

The USIM can make a variety of requests using the proactive UICC service. Examples include: requesting the ME to display USIM-provided text, notifying the ME of changes to EF(s), and providing local information from the ME to the USIM [16]. The command of interest here is *REFRESH*. The *REFRESH* command requests the ME to carry out an initialisation procedure, or advises

the ME that the contents of EF(s) have been changed. The command also makes it possible to restart the session by performing a reset [16].

### 2.3 The AKA Protocol

The AKA protocol is at the core of mobile telephony air interface security, and is regularly performed between the visited network and the UE. The involved parties are the home network (that issued the USIM), the serving network, and the UE. The AuC of the home network generates authentication vectors (used by the serving network in AKA) and sends them to the serving network. In the schemes proposed in sections 6 and 7, we use the *RAND* value in an authentication vector in a novel way for management of multiple IMSIs.

The AKA protocol starts with the serving network sending a *user authentication request* to the UE. The UE checks the validity of this request (thereby authenticating the network), and then sends a *user authentication response*. The serving network checks this response to authenticate the UE. As a result, if successful, the UE and the network have authenticated each other, and at the same time they establish two shared secret keys.

In order to participate in the protocol, the UE, in fact the USIM installed inside the UE, must possess two values:

- a long term secret key  $K$ , known only to the USIM and to the USIM's home network, and
- a sequence number  $SQN$ , maintained by both the USIM and the home network.

The key  $K$  never leaves the USIM, and the values of  $K$  and  $SQN$  are protected by the USIM's physical security features.

The 48-bit sequence number  $SQN$  enables the UE to verify the 'freshness' of the user authentication request. More specifically, the request message contains two values:  $RAND$  and  $AUTN$ , where  $RAND$  is a 128-bit random number generated by the home network, and the 128-bit  $AUTN$  consists of the concatenation of three values:  $SQN \oplus AK$  (48 bits),  $AMF$  (16 bits), and  $MAC$  (64 bits). The  $MAC$  is a *message authentication code* (or *tag*) computed as a function of  $RAND$ ,  $SQN$ ,  $AMF$ , and the long term secret key  $K$ , using a MAC algorithm known as  $f1$ . The value  $AK$  is computed as a function of  $K$  and  $RAND$ , using a cipher mask generating function known as  $f5$ . The  $AK$  functions as a means of encrypting  $SQN$ ; this is necessary since, if sent in cleartext, the  $SQN$  value would potentially compromise user identity confidentiality, given that the value of  $SQN$  is USIM-specific.

On receipt of these two values, the USIM uses the received  $RAND$ , along with its stored value of  $K$ , to regenerate the value of  $AK$ , which it can then use to recover  $SQN$ . It next uses its stored key  $K$ , together with the received values of  $RAND$ ,  $SQN$ , and  $AMF$ , in function  $f1$  to regenerate  $MAC$  (see Fig. 2); if the newly computed value agrees with the value received in  $AUTN$  then the first stage of authentication has succeeded. The USIM next checks that  $SQN$  is




Fig. 2. Computation of AKA key values at the USIM

a ‘new’ value; if so it updates its stored  $SQN$  value and the network has been authenticated.

If authentication succeeds, the USIM computes another message authentication code, called  $RES$ , from  $K$  and  $RAND$  using a distinct MAC function  $f2$ , and sends it to the network as part of the user authentication response. If this  $RES$  agrees with the value expected by the network then the UE is deemed authenticated.

### 3 User Privacy Threats

In a mobile telephony context, a user identity can be any of the mobile number or the IMSI of a USIM, or the international mobile station equipment identity (IMEI) of an ME. Of these various identities, the IMSI is used to identify the subscriber for authentication and access provision; limiting the degree to which its use compromises user privacy is the main focus of this paper. When a subscriber is roaming, i.e. accessing service from a network other than its home network, the IMSI is sent from the UE via the visited network to the home network. Since the IMSI is a permanent user identity, the air interface protocols are designed to minimise the number of circumstances in which it is sent across the air interface.

Clearly, providing user privacy requires that the permanent user identity cannot be intercepted when sent across the radio link. A level of identity confidentiality is provided by use of the TMSI instead of the IMSI. However, on certain occasions a UE needs to send its IMSI across the air interface in clear-text. One such case is when a UE is switched on and wishes to connect to a new network, and hence will not have an assigned TMSI [23]. Another case is where the serving network is unable to identify the IMSI from the TMSI [17].

An active adversary can intentionally simulate one of these scenarios to force a UE to transfer its IMSI in cleartext. Moreover, several further scenarios have been identified [6, 7, 11, 31] in which user identity privacy is at significant risk. In this paper, we address this privacy threat.

## 4 Threat Model

The schemes we propose in this paper are designed to address real-life threats to user privacy in 3G networks. In particular we have already observed that there are circumstances in which an adversary can cause a UE to send its IMSI across the network in plaintext. This is the threat we aim to mitigate by reducing the impact of IMSI compromise. That is, although the possibility of IMSI compromise remains unchanged, we propose making the IMSI a short term identity and hence prevent the compromise of a long-term user identity. In doing so we must also ensure that two different IMSIs for the same UE are not linkable, at least via the network protocol. This issue is examined further in section 9.1.

In designing our schemes we make the underlying assumption that AKA is sound, and provides mutually authenticated key establishment. We also implicitly assume that the USIM and the network have not been compromised by other means. Of course, if these assumptions are false, then very serious threats exist to both user privacy and security. Since we assume that AKA is secure, and no changes are made to its operation, we do not need to re-examine its security for the schemes discussed here. The main risk introduced by use of the multiple-IMSI schemes we propose is the possibility of loss of IMSI synchronisation between UE and home network, and this issue is addressed in section 9.2.

It is important to note that two of the three schemes we describe rely on using the *RAND* sent as part of the AKA protocol as a means of signalling from the home network to the USIM. From our assumption regarding the security of AKA, we can assume that this provides an authenticated channel with replay detection. This is fundamental to the schemes presented in sections 6.2 and 7.

## 5 A Pseudonymity Approach

We consider here the possible use of multiple IMSIs for a single account to provide a form of pseudonymity on the air interface, even when it is necessary to send the IMSI in cleartext. The use of multiple IMSIs is described here using 3G terminology, but a precisely analogous approach would apply equally to both GSM and LTE systems. However, while all the techniques for IMSI distribution specified in sections 6 and 7 would also work for LTE, only the scheme described in section 6.1 would work for GSM, since the other two schemes rely on UE authentication of the network which is not provided in GSM.

At present, a USIM holds one IMSI along with other subscription and network parameters. We propose that a USIM and the home network support the use of varying IMSIs for a single user account, in such a way that no modification is required to the operation of any intermediate entities, notably the visited

(serving) network and the ME itself. This allows the provision of a more robust form of pseudonymity without making any changes to the air interface protocol. In this section we consider how a change of IMSI can be made.

The following issues need to be addressed to allow use of multiple IMSIs.

- *Transferring IMSIs.* Clearly, before a USIM switches to a new IMSI, it must be present in the USIM and in the database of the home network. Also, new IMSIs must always be chosen by the home network to avoid the same IMSI being assigned to two different USIMs. This requires a direct means of communication between the home network and the USIM (which must be transparent to the serving network and the ME, since our objective is to enable changing of an IMSI without making any changes to existing deployed equipments). In sections 6 and 7 we describe in detail two different strategies for transferring IMSIs from the home network to a USIM.
- *Initiating an IMSI change.* Clearly the IMSI needs to be changed in such a way that both the home network and the USIM know at all times which IMSI is being used, and the home network always knows the correspondence between the IMSI being used by the USIM and the user account. An IMSI change can be triggered either by the USIM or by the home network, as we describe below. However, use of a new IMSI is always implemented by the USIM, since it is the appearance of a mobile device in a network using a particular IMSI which causes a request to be sent by the serving network for authentication information for use in the AKA protocol. That is, when the ME sends an IMSI to the serving network, it is forwarded to the home network. Once the home network sees the ‘new’ IMSI it knows that an IMSI change has occurred and can act accordingly.

This requires that the home network knows that both the previously used IMSI and the ‘new’ IMSI belong to the same account. This will require some minor changes to the operation of the home network’s account database, i.e. to allow more than one IMSI to point to a single account. However, this does not seem likely to be a major problem in practice.

- *Triggering an IMSI change.* Whether the USIM or the home network is responsible for initiating a change of IMSI, logic needs to be implemented to cause such a change to take place. Regardless of whether the USIM or the home network makes the decision, logic needs to be in place in the USIM either to make the decision or to receive the instruction to make the change from the home network; for convenience we refer to this logic as an application, although this is not intended to constrain how it is implemented. The decision-making logic could take account of external factors, including, for example, the elapsed time or the number of AKA interactions since the last change; indeed, if the ME included an appropriate user-facing application, then it might also be possible to allow user-initiated changes. Of course, if the home network is responsible for triggering the change of IMSI, then it needs a means of communicating its decision to the USIM that is transparent to the existing infrastructure, including the serving network and the ME. This issue is addressed in sections 6 and 7.



- *Rate of change of IMSI.* The rate of change of IMSI will clearly be decided by the USIM-issuing network (which equips the USIM with the IMSI-changing application). We observe in this context that section 4.2.2 of ETSI TS 131 102 [12] recommends that IMSI updates should not occur frequently. The rate of change of an IMSI could be determined by the customer contract with the issuing network; for example, a USIM which changes its IMSI frequently might cost more than a fixed-IMSI USIM (or one that only changes its IMSI occasionally), and could be marketed as a special ‘high-privacy’ service.
- *Implementing an IMSI change.* A mechanism will be required for the USIM to indicate to the ME that the IMSI has changed. We propose that this should be achieved by the following steps.
  1. As noted in section 2.1, the IMSI is contained in the elementary file  $EF_{IMSI}$ . When the USIM wishes to change the IMSI, it first updates this file accordingly.
  2. At the first opportunity, the USIM uses the proactive UICC status byte to indicate to the ME that it wishes to issue a command.
  3. When the ME responds with a *FETCH* command, the USIM sends a *REFRESH* command to the ME.
  4. The *REFRESH* command causes the ME to read the  $EF_{IMSI}$  file, allowing it to discover the new IMSI. The next time that the ME needs to send its IMSI to the serving network, it will send the new value.

As noted above, using multiple IMSIs requires a direct and transparent means of communication between the home network and the USIM. The *Unstructured Supplementary Service Data (USSD)* protocol appears at first sight to be a possible channel for such communications. However, the protocol end points are the home network and the ME, rather than the USIM. As such, it could only be used for our purposes if the ME was aware of the multiple IMSI scheme, i.e. the ME would need to be modified — contradicting our design objectives. As a result we do not consider the use of USSD further here, although we note that it might be possible to deploy a smart phone application which could provide the necessary additional phone functionality, a possible avenue for future research.

## 6 Predefined Multiple IMSIs

Our first means of deploying multiple IMSIs involves a USIM being pre-equipped with a number of IMSIs. These IMSIs are all associated with a single account in the home network’s account database. Initially, one of the IMSIs is stored in  $EF_{IMSI}$ . We propose below two ways of initiating an IMSI change in this case.

### 6.1 USIM-Initiated IMSI Change

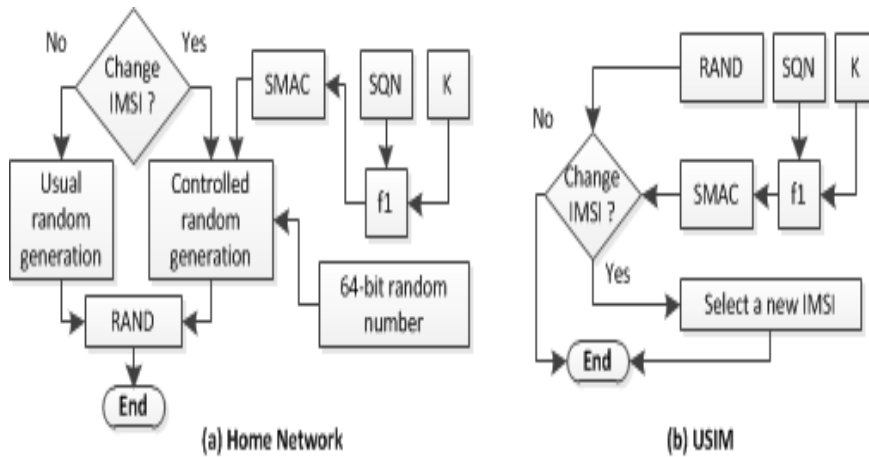
This is the simpler of the two approaches. We suppose the USIM has an application that decides when to initiate an IMSI change. The new IMSI will clearly need to be selected from the predefined list. How the list is used is a matter for

the issuing network. For example, the IMSIs could be used in cyclic order or at random (or, more probably, pseudo-randomly). The USIM changes the IMSI to the ‘new’ IMSI using the procedure described in section 5.

## 6.2 Network-Initiated IMSI Change

In this case, the home network decides when to trigger an IMSI change. The home network will have a richer set of information to use to decide when to change IMSI than the USIM. For example, the home network could change the IMSI whenever the USIM changes serving network or after a fixed number of calls.

As discussed in section 5, when the home network decides to trigger an IMSI change, it must, by some means, send an instruction to the USIM. We propose to use the AKA protocol as the communications channel for this instruction. More specifically, we propose using the value *RAND* of AKA to carry the signal. The IMSI change procedure operates as follows (see also Fig. 3).



**Fig. 3.** IMSI change procedure for predefined multiple IMSIs

1. When the logic in the home network decides that an IMSI change is necessary, a flag is set for the appropriate user account in the AuC database of the home network.
2. Whenever the AuC generates authentication vectors for use in AKA, it checks this flag to see if an IMSI change signal is to be embedded in the *RAND* value. If so it resets the flag and executes the following steps (as in Fig. 3(a)).

- (a) The AuC uses the MAC function  $f1^1$  to generate a 64-bit *MAC* on the subscriber's current sequence number *SQN* using the subscriber's long term key *K*. We refer to this as the *sequence-MAC* or *SMAC*.
- (b) The AuC generates a 64-bit random number *R* using the same process as normally used to generate 128-bit *RAND* values.
- (c) The AuC sets *RAND* to be the concatenation of the *R* and *SMAC*.

If an IMSI change signal is not required, the AuC generates *RAND* in the normal way.

- 3. The AuC follows the standard steps to generate the authentication vector from *RAND*, and sends the vector (including *RAND*) to the serving network.

Whenever the USIM receives an authentication request, it follows the usual AKA steps. If the AKA procedure completes successfully, the USIM checks the *RAND* in the following way (as shown in Fig. 3(b)).

- 1. The USIM uses the received *SQN* and its stored key *K* to regenerate *SMAC*.
- 2. It compares the computed *SMAC* with the appropriate part of *RAND*.
- 3. If they do not agree then the USIM terminates the checking process. However, if they agree then the USIM performs the next step.
- 4. The USIM selects a 'new' IMSI value from the stored list, and later changes the IMSI accordingly using the procedure described in section 5.

We next consider how IMSI changes will work in practice. There are two cases to consider. If the home network is also the serving network then it could potentially force an instance of AKA to occur at will, i.e. making the IMSI change happen almost immediately. However, if the serving network is distinct from the home network, then the home network can only send new authentication vectors when requested by the serving network. Moreover, the serving network may delay before using the supplied authentication vector in AKA. That is, there may be a significant delay between the decision being made to change an IMSI and the signal being sent to the USIM. In either case the phone may be switched off or temporarily out of range of a base station, in which case there will inevitably be some delay. However, regardless of the length of the delay in the signal reaching the USIM (or even if it never reaches the USIM) there is no danger of loss of IMSI synchronisation between the USIM and the home network, since the home network will always keep the complete list of IMSIs allocated to the USIM.

We observe that there is always the chance that a randomly chosen *RAND* will contain the 'correct' *SMAC*, leading to an unscheduled IMSI change by the USIM. However, the probability of this occurring is  $2^{-64}$ , which is vanishingly small. In any case, the occurrence of such an event would not have an adverse impact, since the home network would always be aware of the link between the new IMSI and the particular USIM.

---

<sup>1</sup> For cryptographic cleanliness it should be ensured that the data string input for this additional use of  $f1$  can never be the same as the data string input to  $f1$  for its other uses; alternatively, a slight variant of  $f1$  could be employed here.

Finally, an active interceptor could introduce its own *RAND* into the channel to try to force an IMSI change. However, given that  $K$  is not compromised and  $f_1$  has the properties required of a good MAC function, then no strategy better than generating a random *RAND* will be available. Replays of old *RAND* values will be detected and rejected as a normal part of AKA, at least for 3G and 4G networks, which enable the USIM to check the freshness of an authentication request. Finally, assuming the *SMAC* value is indistinguishable from a random value, a standard assumption for MAC functions, then an eavesdropper will be unable to determine when an IMSI change is being requested.

## 7 Modifiable Multiple IMSIs

The second proposed means of deploying multiple IMSIs involves distributing new IMSI values from the home network to the USIM after its initial deployment, where the home network will choose each new IMSI from its pool of unused values. Such an approach clearly requires a means of communicating from the home network directly to the USIM, and, analogously to the scheme proposed in section 6.2, we describe how the AKA protocol, and specifically the *RAND* value, can be used for this purpose. Before describing the details of the IMSI transfer procedure, we describe some relatively minor changes which are required to the operation of the home network in order to support the scheme.

- The home network must maintain a pool of unused IMSIs, enabling the AuC to dynamically assign a new IMSI to an existing subscriber.
- For each subscriber account in its database, the home network must maintain an *IMSI-change* flag indicating whether an IMSI change is under way. The database must also hold up to two IMSIs for each subscriber; it will always hold the current IMSI (with status *allocated*) and, if the *IMSI-change* flag is set, it will also hold the new IMSI (with status *in transit*), where the possible status values for an IMSI are discussed below. If use of the new IMSI is observed then IMSI status changes are triggered (see below).
- The home network must manage the use of IMSIs so that no IMSI is assigned to more than one subscriber at any one time. This can be achieved by maintaining the status of each IMSI as one of *allocated*, *free*, or *in transit*. The set of IMSIs with status *free* corresponds to the pool of available IMSIs, as above. The status of an IMSI can be updated in the following ways.
  - When the home network selects an available IMSI from the pool to allocate to a USIM, the status is changed from *free* to *in transit*.
  - When the home network receives implicit acknowledgement (in the form of a request for authentication vectors for that IMSI from a network) of a successful IMSI change, the home network changes the status of the IMSI from *in transit* to *allocated*, and the status of the previously used IMSI for that subscriber from *allocated* to *free*. In addition, the current IMSI for the subscriber will be set equal to the new IMSI, the new IMSI will be set to null, and the *IMSI-change* flag will be reset.

- A third case also needs to be considered, that is when an IMSI change instruction never reaches the USIM. If this case is not addressed then future IMSI changes for that USIM will be blocked. On the other hand, making a decision to abandon an IMSI change could be disastrous, i.e. if a USIM makes an IMSI change after the home network has terminated this change (and changed the status of the ‘new’ IMSI back to *free*), then the USIM could be rendered unusable. As a result we propose never to abandon an IMSI change, and instead to resend the new IMSI as many times as necessary until the change is accepted by the USIM. How this works should be clear from the description below.
- If the home network is required to do so by its regulatory environment, e.g. to support lawful interception, it can maintain a log of all the IMSIs assigned to a particular subscriber for however long is required. It is in any case likely to be necessary to retain this information for a period to enable processing of billing records received from visited networks.

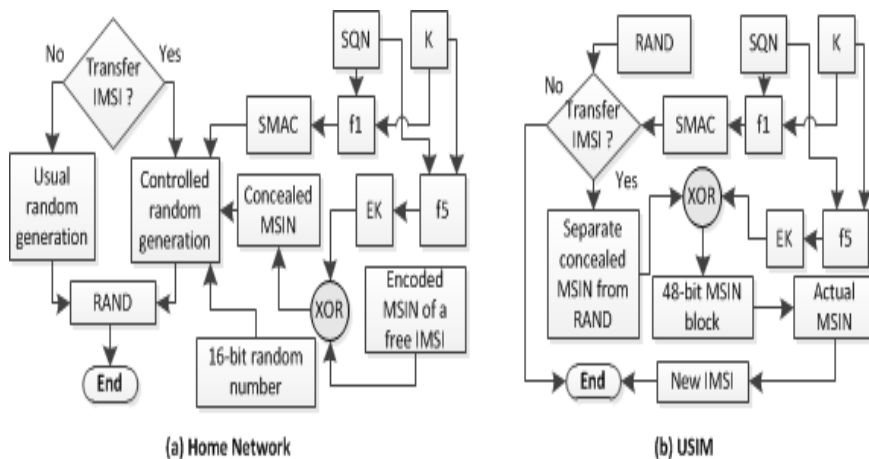


Fig. 4. IMSI change procedure for modifiable multiple IMSIs

The details of the IMSI transfer procedure are as follows (see also Fig. 4).

1. When the logic in the home network decides that an IMSI transfer is necessary for a particular subscriber, it must set the *IMSI-change* flag for that subscriber. Observe that if an IMSI change is already under way then the flag will already be set; in this case the flag is left as it is.
2. Whenever the AuC needs to generate authentication vectors for use in AKA, it checks this flag to see if an IMSI transfer signal and a new IMSI are to be embedded in the *RAND* value. If so it performs the following steps (as shown in Fig. 4(a)). Note that this means that, once an IMSI change has

been initiated, the new IMSI will be embedded in all *RAND* values until evidence of the successful changeover by the USIM has been observed.

- (a) The AuC uses the MAC function  $f1$  to generate a 64-bit *MAC* on the subscriber's current sequence number *SQN* using the subscriber's long term cryptographic key  $K$ . The generated MAC is referred to as the *SMAC*.
- (b) The AuC generates a 48-bit encryption key  $EK$  using the key generation function  $f5$ . The function takes *SQN* as the data input and  $K$  as the key input. Note that observations regarding cryptographic cleanliness and the use here of functions  $f1$  and  $f5$ , analogous to those given in section 6.2 step 2(a), apply here.
- (c) If the new IMSI field in the home network database entry for this subscriber is non-null then a new IMSI has already been assigned, and it is not necessary to choose another new value. Otherwise a new IMSI is selected from the pool of unused IMSIs; the status of this IMSI is changed from *free* to *in transit* and the new IMSI field in the database is given the chosen value. We assume that the MCC and MNC of the IMSI are known to the USIM (since they are fixed for this network operator) and hence only the 9- or 10-digit MSIN needs to be sent embedded in *RAND*. The MSIN is encoded as a 36- or 40-bit value using binary coded decimal, the 'standard' way of encoding IMSIs, and the result is padded to 48 bits by an agreed padding scheme.
- (d) The 48-bit MSIN block is XORed with the encryption key  $EK$ , and we refer to the result as the concealed MSIN.
- (e) The AuC generates a 16-bit random number  $R$  using the same process as normally used to generate 128-bit *RAND* values.
- (f) The AuC sets *RAND* to be the concatenation of the concealed MSIN,  $R$  and *SMAC*.

If an IMSI transfer is not required, the AuC generates *RAND* in the normal way.

3. The AuC follows the standard steps to generate the authentication vector from the *RAND* value, and sends it (including *RAND*) to the serving network.

On receipt of an authentication request, the USIM proceeds using the standard AKA procedure. After successful completion of the AKA protocol, the USIM checks whether the challenge value contains an embedded IMSI in the following way (as shown in Fig. 4(b)).

1. The USIM uses the received *SQN* and its stored long term key  $K$  to regenerate *SMAC*.
2. It compares the computed *SMAC* with the appropriate part of *RAND*.
3. If they do not agree then the USIM terminates the checking process. However, if they agree then the USIM performs the following steps.
  - (a) The USIM retrieves the concealed MSIN from *RAND*.

- (b) The USIM regenerates the encryption key  $EK$  using  $f_5$  with the value of  $SQN$  retrieved during the AKA processing and its long-term stored key  $K$  as inputs.
- (c) The  $EK$  is XORed with the concealed MSIN to recover the cleartext encoded MSIN.
- (d) The USIM generates the new IMSI by prefixing the decoded MSIN with the MCC and MNC.
- (e) The USIM checks whether the new IMSI is the same as the value it is using already; this is essential since it may receive the change instruction more than once. If they are the same it does nothing. If they are different it keeps a record of the new IMSI and later updates its IMSI using the procedure described in section 5.

To reduce signalling costs, it appears to be standard practice for the AuC to generate a small set of authentication vectors for provision to a serving network. If the procedure specified above is followed to generate this set of vectors, and an IMSI change is scheduled for the subscriber, then all the  $RAND$  values in the set will contain an embedded concealed MSIN. Whilst this will cause minimal additional overhead for the USIM, since  $RAND$  values are always checked for an embedded  $SMAC$  value, it will have the benefit of maximising the chance that the IMSI change will be performed by the USIM.

As discussed in section 6.2, there may be a significant delay in the IMSI change signal embedded in  $RAND$  reaching the USIM. However, this will not affect IMSI synchronisation between the home network and the USIM since the home network will not update the current IMSI entry in the subscriber database until it receives a request for authentication vectors from a visited network using this new IMSI. As discussed above, once a new IMSI has been assigned to a subscriber (with the *in transit* status), every  $RAND$  generated for that USIM will contain the embedded IMSI value until the success of the change has been observed. Finally, as in section 6.2, there is the chance that a randomly chosen  $RAND$  could contain a ‘correct’  $SMAC$ , triggering an unauthorised IMSI change. However, the probability of such an event is vanishingly small, and certainly orders of magnitude smaller than the probability of a USIM failure. We therefore do not consider it further here.

## 8 Experimental Validation

Testing the schemes is challenging due to the unavailability of a test UMTS network. However, the availability of the *SIMtrace* [1] hardware and the software implementations of the *Osmocom project* [30] support testing of the necessary modifications to the USIM. We used SIMtrace to trace USIM-ME communications, together with *SysmoUSIM-SJS1* [2], a standards-compliant test UMTS UICC card.

To validate the proposed modification to the USIM, we first ran an experiment to update the IMSI value in the test USIM. Using a standard contact

smart card reader, smart card scripting tool, and a custom script we were able to modify the IMSI value. The process is similar to updating any file in the USIM file structure given the appropriate access grant. We subsequently developed a *SIM toolkit applet* using the Java card framework and the packages included with the 3GPP technical specification covering USIM API for Java card [5]. The SIM toolkit applet was designed to use the *REFRESH* proactive command to arrange for the ME to fetch the new IMSI. We chose to use the *REFRESH* command, as the command is understood by all MEs which support proactive commands. We tested the full range of modes of the *REFRESH* command [15], i.e. USIM initialization and file change notification, file change notification, and UICC reset.

We loaded the applet in the USIM to carry out further tests. We connected the USIM and the ME to the SIMtrace device, which was connected to a laptop to record the APDU commands exchanged between the USIM and the ME. We observed that, when a *REFRESH* is executed in the USIM initialization and file change notification mode, a series of read commands are issued by the ME. Although the record of commands exchanged showed that the IMSI file is read, because of our lack of a test UMTS environment we were unable to confirm that the read operation actually updated the IMSI value stored in the ME. When the mode is changed to UICC reset, the ME simply restarted its session, as expected. These observations confirmed that, if the *REFRESH* command is used in the UICC reset mode, the ME is made aware of the new IMSI. As a result, all future authentication procedures performed by the UE will use the new IMSI, which will have the effect of notifying the home network of use of the new IMSI.

During the experiments, we tested a range of standard MEs, all of which support the proactive command. As mentioned earlier, due to the unavailability of a test UMTS network we were unable to implement the modified home network. However the changes required at the home network are purely minor software changes to the operation of the AuC database, which should not require significant additional computing resource.

## 9 Analysis

The above proposals raise privacy and availability issues, which we now discuss.

### 9.1 User Privacy

The use of multiple IMSIs does not provide a complete solution to user identity confidentiality. While in use, the IMSI still functions as a pseudonym, potentially enabling the interactions of a single phone to be tracked for a period; of course, this is always true for any mobile network when a subscriber resides in a single location area, even where only a privacy-preserving TMSI is used. Of course, the more frequently IMSIs are changed the less the impact of possible tracking, but frequent IMSI changes have an overhead in terms of database management. The use of a predefined set of IMSIs further restricts the degree of user identity



confidentiality protection. In this case, over a period of time it might be possible for an eavesdropper to link at least some of the fixed IMSIs.

The design of the schemes ensures an eavesdropper is unable to infer any confidential information from the value of *RAND*. As discussed in sections 6.2 and 7, in schemes where the *RAND* is used to signal to the USIM, the *RAND* is constructed so that it is indistinguishable from a truly random value; this is based on the assumption that a MAC generated using  $f_1$  and a data string encrypted using the output from  $f_5$  are indistinguishable from random data. Moreover, in the scheme described in section 7 where the IMSI is sent embedded in *RAND*, the IMSI (actually the MSIN) is encrypted to prevent an eavesdropper observing it.

Overall the IMSI-changing proposal can be seen as allowing a trade-off between user privacy and the cost of implementing frequent IMSI changes.

## 9.2 IMSI Synchronisation

If, in the modifiable multiple IMSIs case, an active adversary is able to persuade the USIM to change its IMSI to an unauthorised value, then the USIM (and the UE) will cease to be able to access the network. It is therefore essential that robust cryptographic (and other) means are used to guarantee the correctness and timeliness of the new IMSI.

In the predefined IMSIs schemes described in section 6, loss of synchronisation cannot arise, as even if the USIM is persuaded to make an unauthorised change the new IMSI will be known to the home network. In any event, as argued in section 6.2, the probability of such an event is vanishingly small.

Loss of synchronisation appears to be a more significant threat in the case where new IMSIs are sent embedded in the *RAND* value, as in the scheme described in section 7. However, as discussed there, for similar arguments to those given in section 6.2, the probability of a random *RAND* giving a correct *SMAC* is negligible. Also, malicious changes to a valid *RAND*, e.g. involving changing the encrypted MSIN whilst leaving the *SMAC* unchanged, will be detected by the AKA network authentication process.

## 10 Related Work

To the best of our knowledge, no user privacy enhancing scheme for mobile telephony has previously been proposed that does not require changes to the existing network infrastructure, i.e. the serving network and/or the mobile equipment. While other authors observe that significant changes to widely deployed infrastructure are unlikely to be feasible [8, 22], realistic and practical proposals have not been made. Choudhury et al. [8] proposed a scheme to improve user identity confidentiality in the LTE network. Their scheme involves significant changes to the air interface protocol. They propose the use of a frequently changing dynamic mobile subscriber identity (DMSI) instead of the IMSI across the air

interface. The DMSI is constructed by concatenation of the MCC, MNC, a random number chosen by the mobile operator, and a 128-bit encrypted version of the chosen random number. As a result, the structure of the DMSI differs significantly from the structure of the IMSI. The DMSI is updated on every run of the authentication. The DMSIs are managed by the home network and the USIM. However, the use of the DMSI imposes changes in the protocol messages, mobile equipment, and the serving network. Table 1 summarises the change impact of this scheme, where a ‘Yes’ indicates that changes are required to the indicated part of the system.

Køien [22] has recently proposed a privacy enhanced mutual authentication scheme for LTE. Although the author claims to use existing signalling mechanisms, the author introduces identity-based encryption to encrypt the IMSI when sent across the air interface. The scheme does not introduce any new system entities, but does make significant changes in system operation. The scheme suggests that the ME generates a public key ID computed from the home network ID, serving network ID, and an agreed expiry value, where the home network computes the corresponding private key. The serving network broadcasts the key components of the public ID to its subscribers. Moreover, a mobile device needs to generate a random number to be used in encryption. All these functional components require major modification to the mobile device. The home network shares the private key with the serving network as the key is used by the serving network to decrypt the encrypted IMSI. A serving network has to maintain individual private keys for each home network. In the event of updating the public key ID for any serving network, all partner home networks need to recompute the private key and ensure a secure exchange of keys. This will introduce new signalling messages into the core network. As a result it appears that the scheme is not easily deployable as it is not transparent to the mobile equipment or the serving network. Thus the privacy-protecting enhancements to the system come at the cost of significant modifications to all the major elements of the system; it is therefore more appropriate to consider it as a proposal for a future network. Table 1 summarises the change impact of the Køien scheme.

**Table 1.** Comparison of change impact of the proposed scheme with other proposals

Domain/Proposal	Choudhury et al. [8]	Køien [22]	Our proposal
USIM	Yes	Yes	Yes
Mobile equipment	Yes	Yes	No
Signaling message	Yes	Yes	No
Serving network	Yes	Yes	No
Home network	Yes	Yes	Yes

Sung et al. [26] have proposed a scheme to provide location privacy which uses multiple IMSIs for a single (U)SIM in some ways similar to our proposal. However, it involves an additional party in its operation, needs support by the

ME, and requires wireless data connectivity for sending and receiving calls. The threat model is also very different, in that the home network is a potential adversary. The scheme employs phones without a local SIM; instead the phone's software retrieves a virtual SIM offered by an Internet-accessible third party, which is used for a limited period and paid for using an anonymous Internet payment system where messages are sent via Tor.

Tagg and Campbell [27] describe a scheme to use multiple IMSIs for multiple networks with a single USIM. Their scheme involves the use of an update server to provide a suitable IMSI as and when it is required. The objective is to avoid roaming charges by dynamically switching network provider. Marsden and Marshall [24] propose a similar approach. The focus of their work is thus very different to the schemes described above; they also do not provide a means of transferring new IMSIs transparently to a USIM.

Dupré [10] presents a process to control a subscriber identity module (SIM) for mobile phone systems. Generic guidance regarding the transmission of control information from the network to the SIM is provided. The schemes described in this paper extend Dupré's idea in a more concrete way.

## 11 Conclusions

In this paper we propose two general approaches to using multiple IMSIs for a mobile telephony subscriber. The goal of these proposals is to improve user privacy by reducing the impact of IMSI disclosure on the air interface. The approaches do not require any changes to the existing deployed network infrastructures, i.e. to the serving network, air interface protocols or mobile devices. The overhead introduced is modest and should be feasible to manage in real-world networks. One major advantage is that the proposed schemes could be deployed immediately since they are completely transparent to the existing mobile telephony infrastructure.

The proposed schemes provide a form of pseudonymity on the air interface, even when it is necessary to send the IMSI in cleartext. The schemes reduce the impact of user privacy threats arising from IMSI capture.

Future work could include setting reliable rules for triggering an IMSI change. A formal security analysis of the complete proposal could provide additional confidence in its robustness.

## References

1. Osmocom SIMtrace. <http://bb.osmocom.org/trac/wiki/SIMtrace>, accessed: 2015-05-20
2. SysmoUSIM-SJS1 SIM + USIM. <http://www.sysmocom.de/products/sysmousim-sjs1-sim-usim>, accessed: 2015-05-20
3. 3rd Generation Partnership Project: 3GPP TS 23.003 Version 3.14.0 (2003-12): 3rd Generation Partnership Project; Technical Specification Group Core Network; (Numbering, addressing and identification) (December 2003)

4. 3rd Generation Partnership Project: 3GPP TR 21.905 Version 10.3.0; 3rd Generation Partnership Project; Technical Specification Group Services and System Aspects; Vocabulary for 3GPP Specifications (2011)
5. 3rd Generation Partnership Project: 3GPP TS 31.130 Version 10.0.0; Technical Specification Group Core Network and Terminals; (U)SIM Application Programming Interface (API); (U)SIM API for Java Card (Release 10) (2011)
6. Arapinis, M., Mancini, L., Ritter, E., Ryan, M., Golde, N., Redon, K., Borgaonkar, R.: New privacy issues in mobile telephony: Fix and verification. In: Yu, T., Danezis, G., Gligor, V.D. (eds.) ACM Conference on Computer and Communications Security, CCS '12, Raleigh, NC, USA, October 16–18, 2012. pp. 205–216. ACM (2012)
7. Arapinis, M., Mancini, L.I., Ritter, E., Ryan, M.: Privacy through pseudonymity in mobile telephony systems. In: 21st Annual Network and Distributed System Security Symposium, NDSS 2014, San Diego, California, USA, February 23–26, 2014 (2014), <http://www.internetsociety.org/doc/privacy-through-pseudonymity-mobile-telephony-systems>
8. Choudhury, H., Roychoudhury, B., Saikia, D.K.: Enhancing user identity privacy in LTE. In: IEEE 11th International Conference on Trust, Security and Privacy in Computing and Communications (TrustCom), 2012. pp. 949–957. IEEE (2012)
9. Deng, Y., Fu, H., Xie, X., Zhou, J., Zhang, Y., Shi, J.: A novel 3GPP/SAE authentication and key agreement protocol. In: IEEE International Conference on Network Infrastructure and Digital Content, 2009 (IC-NIDC 2009). pp. 557–561. IEEE (2009)
10. Dupré, M.: Process to control a Subscriber Identity Module (SIM) in mobile phone system (2004), US Patent 6,690,930
11. European Telecommunications Standards Institute (ETSI): ETSI TS 121 133 Version 4.1.0 (2001-12): Universal Mobile Telecommunications System (UMTS); 3G Security; Security threats and requirements (December 2001)
12. European Telecommunications Standards Institute (ETSI): ETSI TS 131.102 Version 4.15.0 Release 4; Universal Mobile Telecommunications System (UMTS); Characteristics of the USIM application (2005)
13. European Telecommunications Standards Institute (ETSI): ETSI TS 121 111 Version 8.0.1 (2008-01): Universal Mobile Telecommunications System (UMTS), USIM and IC card requirements (January 2008)
14. European Telecommunications Standards Institute (ETSI): ETSI TS 102 221 Version 8.2.0; Smart Cards; UICC—Terminal Interface; Physical and logical characteristics (2009)
15. European Telecommunications Standards Institute (ETSI): ETSI TS 131 111 Version 7.15.0: Digital cellular telecommunications system (Phase 2+); Universal Mobile Telecommunications System (UMTS); LTE; Universal Subscriber Identity Module (USIM) Application Toolkit (USAT) (2010)
16. European Telecommunications Standards Institute (ETSI): ETSI TS 102 223 Version 11.1.0; Smart Cards; Card Application Toolkit (CAT) (2012)
17. European Telecommunications Standards Institute (ETSI): ETSI TS 133 102 Version 11.5.1 (2013-07): Digital cellular telecommunications system (Phase 2+); Universal Mobile Telecommunications System (UMTS); 3G Security; Security architecture (July 2013)
18. Forsberg, D., Horn, G., Moeller, W.D., Niemi, V.: LTE Security. John Wiley & Sons (2010)

19. International Organization for Standardization: ISO/IEC 7816—3; Identification cards — Integrated circuit cards; Part 3: Cards with contacts — Electrical interface and transmission protocols (November 2006)
20. Juang, W.S., Wu, J.L.: Efficient 3GPP authentication and key agreement with robust user privacy protection. In: *Wireless Communications and Networking Conference, 2007. WCNC 2007.* IEEE. pp. 2720–2725. IEEE (2007)
21. Khan, M.S.A., Mitchell, C.J.: Another look at privacy threats in 3G mobile telephony. In: Mu, Y., Susilo, W. (eds.) *19th Australasian Conference on Information Security and Privacy, ACISP 2014, Wollongong, NSW, Australia, July 7-9, 2014.* pp. 386–396. *Lecture notes in Computer Science*, Springer International Publishing (2014)
22. Kjøien, G.M.: Privacy enhanced mutual authentication in LTE. In: *IEEE 9th International Conference on Wireless and Mobile Computing, Networking and Communications (WiMob), 2013.* pp. 614–621. IEEE (2013)
23. Kjøien, G.M., Oleshchuk, V.A.: *Aspects of Personal Privacy in Communications: Problems, Technology and Solutions.* Denmark: River Publishers (2013)
24. Marsden, I., Marshall, P.: Multi IMSI system and method (Feb 20 2014), <http://www.google.com/patents/US20140051423>, US Patent App. 13/966,350
25. Samfat, D., Molva, R., Asokan, N.: Untraceability in mobile networks. In: *Proceedings of the 1st Annual International Conference on Mobile Computing and Networking.* pp. 26–36. *MobiCom '95*, ACM, New York, NY, USA (1995), <http://doi.acm.org/10.1145/215530.215548>
26. Sung, K., Levine, B.N., Liberatore, M.: Location privacy without carrier cooperation. In: *IEEE workshop on Mobile Security Technologies, MOST 2014, San Jose, CA, USA, May 17, 2014* (2014)
27. Tagg, J., Campbell, A.: Identity management for mobile devices (Dec 6 2012), <http://www.google.com/patents/US20120309374>, US Patent App. 13/151,942
28. Telecommunication Standardization Sector of ITU: ITU-T E.212: International operation Maritime mobile service and public land mobile service (The international identification plan for public networks and subscriptions) (May 2008)
29. Valtteri, N., Nyberg, K.: *UMTS Security.* John Willey & Sons Limited (2003)
30. Various Contributors: Osmocom Project. <http://osmocom.org>, accessed: 2015-05-20
31. Vintila, C.E., Patriciu, V.V., Bica, I.: Security analysis of LTE access network. In: *ICN 2011, The 10th International Conference on Networks.* pp. 29–34 (2011)
32. Xiehua, L., Yongjun, W.: Security enhanced authentication and key agreement protocol for LTE/SAE network. In: *7th International Conference on Wireless Communications, Networking and Mobile Computing (WiCOM), 2011.* pp. 1–4. IEEE (2011)