On Travelling *Incognito*

A. Herzberg H. Krawczyk G. Tsudik

IBM T.J. Watson Research Center
N.Y. 10598, USA
{*amir,hugo*}*@watson.ibm.com*

IBM Zürich Research Laboratory
CH-8803 Rüschlikon, Switzerland
gts@zurich.ibm.com

Abstract

User mobility is rapidly becoming an important and popular feature in today's networks. This is especially evident in wireless/cellular environments. While useful and desirable, user mobility raises a number of important security-related issues and concerns. One of them is the issue of tracking mobile user's movements and current whereabouts. Ideally, no entity other than the user himself and a responsible authority in the user's home domain should know either the real identity or the current location of the mobile user. At present, environments supporting user mobility either do not address the problem at all or base their solutions on the specific hardware capabilities of the user's personal device, e.g., a cellular telephone.

This paper discusses a wide range of issues related to anonymity in mobile environments, reviews current state-of-the-art approaches and proposes several potential solutions. Solutions vary in complexity, degree of protection and assumptions about the underlying environment.

1 Introduction

A typical situation arising in any mobile environment is when an user (or a device that corresponds to a user) registered in one domain appears in a different, foreign domain. The user's goal is to obtain certain services from or through the local (foreign) domain. At the same time, the local domain's goal is either (or both) i) to be paid, or ii) to authorize the user. In the latter case, although the user is not known locally, he must be somehow authenticated and his "solvency" or good standing must be confirmed. Typically, the only entity able to comment on the user's identity and current standing is his home authority. This problem has received some attention as evidenced by several solutions in the recent literature.[1, 4, 5, 7]

An orthogonal security issue arising as a result of mobility is the confidentiality of mobile user's identity and movements. For obvious reasons, it is desirable to keep this information secret. In other words, passive eavesdroppers and active intruders should not be able to identify or track the user. In fact, it can be argued that even the visited locations should not be privy to user's real identity. For the most part, only the home domain authority should be informed as to the mobile user's real identity, itinerary and current whereabouts. Moreover, in some cases, it is even necessary to hide user's movements and whereabouts from the home as well. We refer to these issues collectively as the *anonymity* problem.

The major assumption made so far is that, every user has a **home** location where this user is registered as an *indigent.* This is true of most, if not all, current mobile user environments. However, this assumption may not hold in the future. Increasing popularity of the electronic cash concept (pioneered by David Chaum [8, 9]) may eventually result in *homeless* mobile user environments. For example, a cellular telephone user, instead of subscribing with a particular service provider, may simply obtain some electronic cash at, say, a vending machine or a kiosk, and spend it piecemeal on making calls from anywhere he pleases.

Secure, untraceable and anonymous electronic cash is a well-researched concept and there is little doubt that it will at some point become practical and widespread. However, most current mobile environments subscribe to the *credit* model whereby a user consumes service and pays/accounts for it in bulk while individual charges are funneled through an entity which serves as the user's creditor, i.e., his home location. Those are the environments that already now need a set of security mechanisms to ensure anonymity of mobile users. The goal of this paper is to shed some light onto various aspects of the anonymity problem in the credit-based mobile environments and propose a range of possible solutions.

2 Problem Scope

To-date, the anonymity problem has received the most attention in the wireless/cellular context. Both GSM in Europe and CDPD in the North America recognized the need to protect the identities of mobile users. (We review their respective approaches in the next section.)

Somewhat surprisingly, anonymity has not been addressed in more traditional, wired network and internetwork environments. A possible explanation is that user mobility is not yet widespread in today's wired networks. One exception is the so-called *anonymous remailers* which have been gaining popularity on the Internet. (An anonymous remailer is, essentially, a

[0] Names appear in alphabetical order.

[1] In Proceedings of IEEE Workshop on Mobile Computing Systems and Applications, December 8-9 1994; Santa Cruz, Ca.

clearinghouse for anonymous newsgroup postings and regular electronic mail.)

Even though it is most pressing in computerized mobile environments, anonymity is an issue wherever user mobility is supported.

In particular, anonymity can be an important factor in the all-too-familiar electronic banking. Most banks offer some form of electronic services and (at least in the US) participate in inter-bank networks such as CIRRUS or STAR. Typically, a customer is issued an ATM card by his "home" bank. Armed with and ATM card, a customer can (among other things) withdraw cash from a multitude of ATM-s located throughout the world. The ability to obtain instant cash is very useful, however, the customer's identity is revealed every time an ATM is used. This makes is possible for foreign banks or ATM providers to track the movements and whereabouts of the customer.

Another familiar environment is the ubiquitous credit card payment system. A typical consumer shopping and paying for goods or services with his credit card (e.g., Master Card, VISA, American Express, etc.) discloses his identity to the payee (i.e., the retailer or service provider.) Moreover, the identity of the payee is revealed to the central clearinghouse and, subsequently, to the organization that issued the credit card. All this is, strictly speaking, not necessary. Ideally, a consumer should not have to reveal anything to the payee other than the confirmation of his good standing with respect to the credit card. Conversely, the identity of the payee should not necessarily be made known to the consumer's credit card issuer.

3 Issues

There are several aspects underlying the problem of anonymity in mobile user environments.

3.1 General Identity Confidentiality

The central issue in maintaining a secret identity is to prevent anyone from discovering a correspondence between a mobile *user* and a particular user registered in a particular home domain. The intuitive first-step solution is to assign a travelling alias to every mobile user or device when away from the home domain. As described in the subsequent sections, this alias can be fixed or ever-changing.

3.2 Anonymity in Foreign Domains

A mobile user who appears in a foreign domain typically does so because he would like to obtain certain services, e.g., make telephone calls or send/read mail. Such services are, for the most part, *not free*[2] and the foreign domain needs to make sure that it can be paid for services consumed. However, in doing that, it does not necessarily need to know the real identity of the mobile user.

In some circumstances, the foreign domain may, as a matter of policy, demand to know the real identity of the user. Then, the user himself or his home authority can communicate the real identity in secret (assuming there exist means for secure communication) in order to prevent any unauthorized party from *unmasking* the user. On the other hand, if it is not imperative for a foreign domain to know the real identity, an alias may suffice. Of course, an alias may still have to be corroborated by the user's home authority.

[2] Even if certain services are free, the user's identity may still have to be established and/or confirmed.

3.3 Preventing Tracking/Correlation

Even if a mobile user adopts a travelling alias, his movements can still be tracked by a hostile intruder. This is possible if the alias is fairly static, e.g., fixed for a given trip or semi-permanently fixed for a given user. An alias of this type is similar to a long-term password; once cracked, the identity and the movements of the user can be compromised on a long-term basis. Consequently, it is desirable to have *frequently-changing* aliases.

However, the actual frequency of alias change depends on the particulars of the environment. For example, a user not equipped with some kind of a trusted device can not be expected to generate aliases that change every single time a network is accessed. On the other hand, a modern cellular telephone is a sufficiently sophisticated device that can generate one-time aliases within a tamper-proof module.[3]

3.4 Anonymizing Foreign Domains

Finally, a relatively non-obvious anonymity-related issue is keeping the identities of the (visited) foreign domains secret from one's own home domain. This is probably something that has no appeal to cellular telephony since incoming calls for a mobile user are typically routed (at least initially) through the home domain. Nonetheless, customers carrying credit or ATM cards may benefit from this feature for the reason that it keeps the individual's travelling and spending habits private. As an example, it can be envisioned that, upon making a purchase, a credit card customer authorizes a certain amount and the store then *anonymously* verifies that the same amount is authorized by the credit card company.

3.5 Limitations

There remain some basic limitations that are difficult to circumvent. One such limitation, for example, occurs if the foreign domain authority needs to to know the identity of the user's home domain. This is likely to be the case in many mobile-user environments since charges incurred "abroad" must be eventually propagated to the home domain. Furthermore, as mentioned before, only the home domain can comment on the user's current standing. (As a side note, one can envisage an environment where communication between domain authorities is "anonymized" by a central *clearinghouse*. In this case, it is necessary to assign aliases to domains so that a travelling user can reference his home domain by an alias; it is then left up to the central clearinghouse to resolve domain aliases.)

[3] Of course, not all cellular phones are tamper-proof. However, a cellular phone offers some basic security by virtue of being a *personal* device. Also, some phones are password- or PIN-activated.

Another limitation has to do with tracking user's movements. For example, in the cellular milieu, a user can migrate from one domain to the next (adjacent domain) while actively using the phone. A common technique called *hand-over* is used to pass the call state between adjacent domains. In doing so, it inevitable that the two domains – both of which can be foreign – discover at least a portion of the user's path.

4 Existing Approaches

Current state-of-the-art mobile systems are exemplified by *Groupe Spécial Mobile* (GSM) in Europe and *Cellular Digital Packet Data* in North America. In this section we briefly describe their respective anonymity services.

4.1 Anonymity in GSM

An active mobile unit (cellular phone) in GSM is always under control of the local base station (BS). Whenever a mobile unit crosses the *cell* boundary, a different BS takes over the handling of the unit. The transfer of state is sometimes referred to as the hand-off process. If the two adjacent cells belong to different domains, a somewhat more involved process takes place. Both are discussed below.

Over the airlink, GSM protects the identity of the mobile unit and its home location by transmitting them to the local BS from the previous visited BS. Each mobile unit, when registering with a BS, gets a temporary identity (TMSI)[4] The mobile uses the TMSI instead of its actual identity, whenever it "talks" with this BS. When moving to a new BS, the mobile unit sends the previous TMSI and LAI (Location Area Identifier) of the previous BS. The new BS receives the actual identity of the mobile unit and the home from the previous BS.

If the previous BS is unreachable, current BS can give up and ask the mobile to reveal its actual home and identity. This fall-back process can be exploited by an active attacker who, masquerading as a BS, can ask the mobile to reveal its identity claiming to have no contact to the previous base. Most implementations, in this case, would reveal the true identity of the mobile unit (IMSI)[5].

Even greater opportunities for hostile tracking exist on the inter-domain level. When a mobile unit crosses the domain boundary or is simply activated in a new domain, a registration process takes place. The purpose of the latter is to establish (at the domain level) the necessary state for the mobile unit. In the course of registration, the mobile unit is authenticated with the direct aid of its home location (see [1, 5]) and a TMSI is assigned. However, authentication of the mobile involves its *real* identity (IMSI) communicated in the **clear** over the air link.

In summary, the GSM design focuses on mobile unit's anonymity over the air link. The design offers no anonymity against the base stations. In fact, every base discover not only the identity of the mobile but also his previous and next base stations.

[4] Temporary Mobile System Identifier.

[5] International Mobile System Identifier.

Furthermore, all information between BS-s flows in the clear. Hence, an attacker can easily discover identities and locations by passive eavesdropping on the inter-BS communication. An active attacker may, in fact, use a base as an oracle that reveals the identities, therefore it does not even need to intercept messages but can issue them at its convenience. For this, the attacker would claim to be another BS to which the user have connected.

A final note is that the GSM procedure relies very heavily on synchronized state in the mobile unit. If this state is lost in either the mobile or previous BS, anonymity is compromised. This requires bases to keep state of mobiles even after they left. The scheme is really designed to support a mobile moving between adjacent BS-s.

4.2 Anonymity in CDPD

In many aspects, CDPD is very different from GSM. The equivalent of a base station in CDPD is *Mobile Data-Base Station* (MDBS). Unlike a base station in GSM, MDBS is a low-level entity that is not involved in any security-related activity. In fact, an MDBS does not even take part in inter-cell hand-over of mobile unit's state. All of the mobility management as well security, functions are concentrated in the Message Data Intermediate System – MD-IS. Each MD-IS controls an *area* covered by a number of MDBS-s. Therefore, it only makes sense to discuss anonymity with respect to inter-area mobility.

Upon arrival to a new area, the mobile unit first engages in a Diffie-Hellman key exchange protocol[2] with the local MD-IS. As a result, both parties obtain a shared secret key. Subsequently, the mobile unit encrypts its real identity (Network Equipment Identifier – NEI) and transmits it to the local MD-IS.

While seemingly more secure than GSM, the CDPD approach has two major drawbacks. First, it allows the local MD-IS to discover the *real* identity of the mobile unit. As we argued above, the identity of the mobile unit should not be revealed to the local authority. It is sufficient for the mobile's identity and current standing to be corroborated by the home domain authority. The second problem is due to the nature of the Diffie-Hellman key exchange protocol. Its purpose is to establish a secret key on-the-fly. This means that an active attacker masquerading as the local domain authority can engage in the key exchange protocol with the mobile unit and obtain a shared key. The mobile unit then transmits its real identity enciphered with the new key and the intruder can simply decipher the entire transmission.

4.3 Summary

Anonymity services provided by GSM and CDPD are summarized in Table 1. This table is presented not so much to illustrate what GSM and CDPD offer, but, rather, what they **do not** offer.

Both GSM and CDPD view their network environment as divided into two parts: air links and fixed network. The former is the "ether" over which subscribers communicate with base stations and the latter is the rest of the network, i.e., the medium over

Identity protection of:	Protection From: Active Attacker on air link	Active Attacker on fixed network	Passive Attacker on air link	Passive Attacker on fixed network	Visited Location	Home Location
Mobile Unit	GSM		GSM, CDPD	GSM		
Home Location						xxxxx
Visited Location					xxxxx	

Table 1: Summary of Anonymity Services in GSM and CDPD

which base stations, message switching centers and other "wired" elements communicate.

The air links are considered wide-open and vulnerable while the fixed network is considered secure. (Exactly what makes the fixed network secure is not specified.) Thus, anonymity and other security services are either not implemented or greatly relaxed over the fixed network. This is a reasonable approach if dedicated, private links make up the fixed network (and/or if strong link security is provided.) However, not all mobile environments have the luxury of secure fixed network.

Another underlying assumption is the benevolence of both home and visited locations. This means that current whereabouts and movements (path) of the mobile user a known to the home location authority. Conversely, both the user identity and the identity of his home location are exposed to the visited location authority.

In the remainder of this paper we consider several environments supporting user mobility. We begin with the most *low-tech* case: a wired network (e.g., the Internet) where mobile users gain network access through a multitude of entry points (workstations, hosts, dial-up terminals, PDAs, etc.) without possessing any trusted personal equipment. In spite of limitations inherent to this low-tech case, a certain degree of anonymity can be provided.

We then proceed to a more sophisticated scenario where users are assumed to possess a personal trusted device (e.g, a smartcard or a token-card.) First we consider anonymity protection scheme based on the conventional, shared-key cryptography model. Finally, we briefly discuss the case with the a device based on public key cryptography.

5 Weak Anonymity

We begin addressing the anonymity problem by developing solutions suitable for relatively low-tech mobile environments where users are not necessarily equipped with smartcards or some such devices. Moreover, even if smartcards are available, we assume that they are unable to perform modular exponentiation – the cornerstone of public key encryption.

5.1 Time-Based Aliasing

In the description below we assume that the "world" is partitioned into administrative domains. Every user has a permanent home in at least one domain and each domain has at least one Authentication Server (AS) – an entity that performs authentication, key distribution, alias resolution and other security-related tasks. In addition, every domain D_X selects a domain-wide time interval δ_X. δ_X is expected to be relatively coarse, e.g., an hour, a day or longer.

When a user U_x whose home is in domain D_X travels to a foreign domain D_Y, he first needs to be authenticated and a temporary record must be created for him in D_Y so as to facilitate subsequent accesses in D_Y. In other words, if the user plans to linger within D_Y for some time, it makes sense to establish some temporary "home" for him instead of having to contact the home domain upon every access.

The authentication and temporary record establishment procedure can be abstracted as shown in Figure 1.[6] (See [5] for detailed protocol description.) The exact format of the authentication flows is **not important** in this context. Regardless of the authentication specifics, the identity of U_x must be somehow communicated to AS_x. Since U_x can not communicate with AS_x directly, all communication has to flow through the local authority, AS_y. The first authentication flow, as shown, must include a user identification field denoted $SUid$. Similarly, the second flow (from AS_y to AS_x) must also include some form of user identification; it is denoted $\overline{SUid}$.[7]

The most important aspect of this protocol, with respect to user identity confidentiality, is the computation of $SUid$. $SUid$ is computed as:

$$SUid = F\ (\ U,\ T_u,\ PW_u\ )$$

where F is a strong one-way function, such as DES [3] or MD5 [6].[8] T_u is the current time *rounded* to the nearest δ_X value. If the user is **not** equipped with a smartcard-like device, PW_u is the user's password which he enters into the public workstation or some

[6] The authentication protocol may be optionally preceded by a two-flow Diffie-Hellman[2] key exchange *a la* CDPD[4].In this case, the entire procedure becomes resistant to passive intruders since all messages can be enciphered under the new key.

[7] The reason for different notation is that $\overline{SUid}$ may be, for example, an encrypted version of $SUid$.

[8] In case of DES or some other encryption-based function, it is important to note that **no additional** secret key is necessary to compute F since PW_u is sufficient for that purpose.

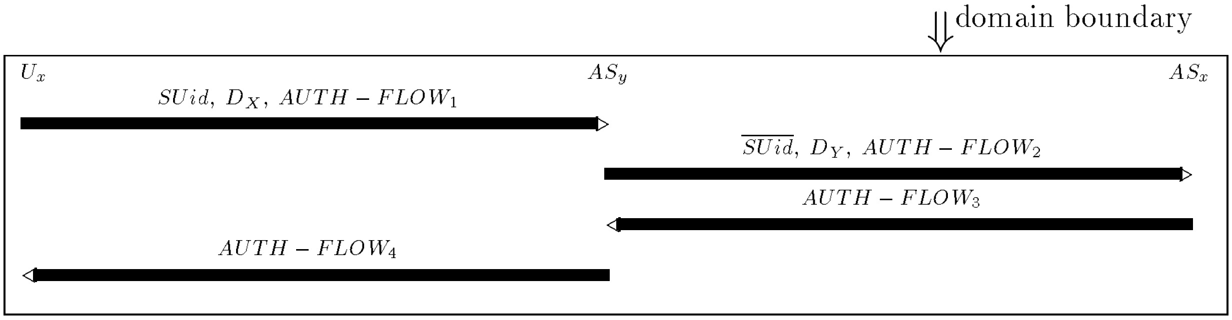

Figure 1: Sample Mobile User Authentication Protocol

such terminal. For a smartcard-bound user, PW_u can be either: 1) a strong key stored within the smartcard (for those smartcards that lack a keypad or other means of input), or 2) a combination of the smartcard's key and the user's password (for smartcards with input capabilities).

As specified, $SUid$ is unintelligible to AS_y. The only information AS_y is able to obtain is that the mobile user is registered in D_X. In the second flow, AS_y transmits $SUid$ (along with other authentication information) to the user's claimed home domain authority AS_x.

The crucial issue is how AS_x determines that $SUid$ corresponds to the locally-registered user U_x. It does so by maintaining an up-to-date table which, for each native user, lists the corresponding $SUid$ value. This translation table is computed for every δ_X interval. Since AS_x already stores the values of U and PW_u for every user, it has all the necessary information to compute up-to-date translation tables.

We note that, since $SUid$-s are not dependent on the users' current location, the translation tables can be pre-computed **off-line** and well in **advance**. This is particularly the case when a relatively coarse δ_X value is used, e.g., one hour or one day. Pre-computation is, of course, commensurate with increased space requirements (to store the alias tables) but it makes the protocol more efficient.

Finally, establishing the "real" identity of the mobile user is only half the work; AS_x must then verify the authentication information supplied in the second flow. However, this is unrelated to the problem at hand. (See, for example, [5] for a treatment of this subject.)

6 Stronger Anonymity Solutions

Although suitable for personal devices (smartcards), the time-based anonymity approach described in the preceding section is truly appealing to an *unarmed* user, i.e., a user without any personal gadgets. We now turn to environments where users are assumed to be more *sophisticated* insofar as personal devices.

A user equipped with a personal device, e.g., a laptop computer or an intelligent cellular phone, can rely on the device to perform complex computations as well as provide secure and non-volatile storage of strong keys and other state information. Naturally, this type of environment is much more amenable to a wide range of anonymity solutions.

In particular, much stronger mechanisms can be implemented such as one-time aliases that change with each mobile user's registration or network access. One-time aliases can be computed by the home authority and communicated to the mobile user/device (in secret) for use in the next registration. Alternatively, one-time aliases can be computed by the mobile device itself by (probabilistically) encrypting its name under a private or public key shared with the home domain.

As an illustration, we briefly describe one such technique which can be quiete easily integrated with other security mechanisms in mobile communications, e.g., into authentication, registration and authorization protocols implemented by different mobility-aware applications and protocols.

The main idea is that in each registration (or network access) the mobile identifies itself by an alias that was generated and communicated to it by the home domain in the previous registration. These aliases need to look unintelligible to a hostile (and, perhaps, active) observer of the communication. The home authority generates each one-time alias by encrypting the name of the mobile using a strong secret key **known only** to the home authority *itself*.

Encryption is done probabilistically, e.g., by adding some random *salt* to the name before encryption. This results in several benefits:

- aliases for the same mobile are always different
- an alias gives away no information about the true identity of the mobile
- compromise of one mobile unit does not compromise aliases of others and does not reveal *previous* aliases for the same mobile unit

Furthermore, the home authority only needs to perform one decryption operation in order to recover the true identity of a mobile user. Since this can be done using a strong symmetric encryption function (e.g., triple-DES), the resultant solution is very efficient.[9] (We note that a solution with totally random aliases assigned by the home authority would require more memory and additional data search in order to identify a user).

Another important property of the present approach is that the strength of the underlying encryption function as well as the size of the encryption key is entirely upto to the individual home authority. In other words, a home domain authority choose any encryption function with any size key. Moreover, it is free to change keys and *even* encryption functions periodically without any impact on the mobile users and with no significant performance degradation. This is because any change in the alias computation procedure is transparent to the mobile users as long as the home authority remains *backward-compatible*, i.e., it continues to "recognize" old aliases computed under older keys or encryption functions.

One important issue in the present solution is the traceability of a mobile user by an intruder. If we assume (naively) that one-time aliases are communicated in the clear then an intruder (even without knowing the true identity of the user) can track the user's movements by correlating the alias supplied by the home authority in one session with the one used by the mobile user in the next session. Therefore it is necessary to encrypt the new alias as it is communicated by the home domain to the mobile user. However, this has no impact on other communication between the mobile user and the home authority, i.e., it is only the new alias that must be encrypted.

This method is illustrated in figure 2. We assume that the network access is initiated by the mobile user U_x. It provides the current alias $ALIAS_i$ in the clear along with some protocol-dependent authentication information. The home authority – AS_x first decrypts $ALIAS_i$, identifies U_x and verifies the authentication information. Next, AS_x generates a new session key K, computes the new alias – $ALIAS_{i+1}$ – and encrypts it under K. Finally, AS_x sends to U_x: i) its own authentication (if applicable), ii) key distribution expression containing K, and iii) new, encrypted alias. The same procedure is repeated upon the next network access.

Our example in figure 2 assumes that a new session key (K_i) is distributed by the home authority to the mobile user upon every network access. This is not mandatory. Instead, a mobile uses (device) can share a long-term strong key with the home authority. In this case, the one-time alias (in the flow from AS_x to U_x) would be encrypted using this long-term key.

A final comment on the present approach is that it requires non-volatile memory in the user's device in order to store the current alias. This is both a blessing and a curse. One one hand, the requirements on the device amount to nothing more than storing the alias (no computation whatsoever.) On the other hand, the alias has to be stored securely and reliably. In the event that the current alias is somehow lost, a recovery scenario can be envisioned whereby the mobile defaults to, say, its serial number or some other "permanent" identity for just one network access.

As mentioned in the beginning of this section, many other similar (i.e., device-based) anonymity approaches are possible. Some would rely on the mobile to generate aliases. This makes sense only with the aid of public key encryption. (It is easy to see that symmetric, conventional encryption would not work unless all mobiles share the same key with the home authority.) Public key encryption solves some problems, e.g., it would no longer require non-volatile storage on the device. However, the computational overhead would naturally increase as would the size of protocol messages. Furthermore, good-quality random number generators would be required on all mobile devices and, finally, a change in a home authority's public key would require additional protocol support.

[9] As opposed to solutions based on Public Key cryptography.

References

[1] M. Rahnema, *Overview of the GSM System and Protocol Architecture*, IEEE Communications Magazine, April 1993.

[2] W. Diffie and M. Hellman, *New Directions in Cryptography*, IEEE Transactions on Information Theory, November 1976.

[3] National Bureau of Standards, *Federal Information Processing Standards*, National Bureau of Standards, Publication 46, 1977.

[4] *Cellular Digital Packet Data (CDPD) System Specification*, Release 1.0, July 19, 1993.

[5] R. Molva, D. Samfat and G. Tsudik, *Authentication of Mobile Users*, IEEE Network, Special Issue on Mobile Communications, Spring 1994.

[6] R. Rivest, *The MD5 Message Digest Algorithm*, Internet DRAFT, July 1991.

[7] M. Beller, L. Chang and Y. Yacobi, *Privacy and Authentication on a Portable Communications System*, IEEE JSAC, Special Issue on Wireless Personal Communications, August 1993.

[8] D. Chaum, A. Fiat and M. Naor, *Untraceable Electronic Cash*, Proceedings of Crypto'88, August 1988.

[9] D. Chaum, *Security Without Identification: Transactions Systems to Make Big Brother Obsolete*, CACM Vol. 28, No. 10, October 1985.

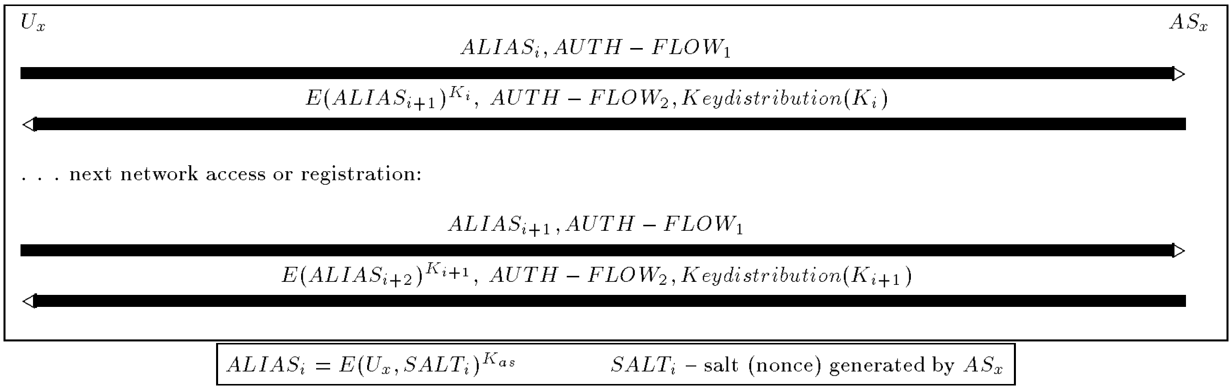

Figure 2: Device-based Mobile User Identification Protocol