Preserving Privacy in Mobile Environments With Virtual Network Stacks

Alfredo Matos
Instituto de Telecomunicações,
Universidade de Aveiro
Aveiro, Portugal
alfredo.matos@av.it.pt

João Girão
NEC Europe Ltd.
Heidelberg, Germany
joao.girao@netlab.nec.de

Susana Sargento and Rui Aguiar
Instituto de Telecomunicações,
Universidade de Aveiro
Aveiro, Portugal
{ssargento|ruilaa@det.ua.pt}

***Abstract*—User privacy is a growing requirement in the evolution of communication networks. In this sense, the concept of *virtual personae*, which corresponds to different identities of the same user, starts getting much attention. However, to provide privacy and non-linkage between these virtual users, a cross-layer approach to identity needs to be supported. This paper proposes a solution to preserve the application layer privacy models by applying the *virtual personae* concept throughout the network stack. It also proposes mechanisms for non-correlation between identities in 4G mobile environments, and addresses the benefits of the evolving multi-homing characteristics of 4G networks to enrich the non-linkage between identities support of our privacy solution.**

Keywords: privacy, mobility, network, persona, identity.

I. Introduction

Nowadays, users travel through several networks with no concern for their privacy or even awareness of privacy threats. Wherever they go and connect their devices, they inform the world of their presence and movement. As shown in Fig. 1, in current and evolving network scenarios, the user will have a large range of interactions with services and applications, where he plays different roles, while performing different activities. Most of the user's actions under each role are usually unrelated among each other; also, it may be the user's interest that these actions remain unrelated. Privacy concerns arise when actions under different roles are easily linked: the user does not wish that his employer can have the means to track what type of products he usually buys, or what he does on his spare time.

This type of privacy issues are already present in today's networks, but are even more hazardous in Next Generation Networks (NGN). NGN, or 4G networks, are entirely IP based, and provide a very large range of services such as VoIP or IPTV. This type of environments enables the paradigm of being always on and connected. Since the users are constantly on-line, they are easily traceable, and the aforementioned linking is easier. Therefore, user privacy is an increasing requirement in 4G networking. On one hand, users may want to hide their location, or movement, from other peers and users in the Internet (there are already some location privacy proposals in the literature [1][2]). On the other hand, as digital lifestyles are adopted, users want to access information depending on their status or context, such as work or leisure, without disclosing that it belongs to the same user. Moreover, in a 4G environment, users want to access the information everywhere and while moving keeping this privacy and non-disclosure requirements.

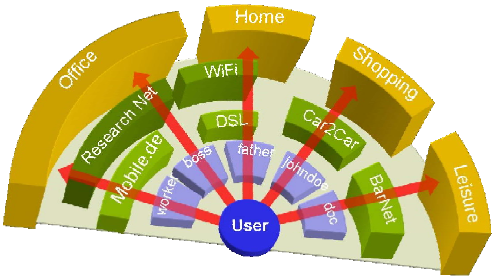

Figure 1. User's interaction with the digital world

With the advent of Web 2.0, users are already moving a great deal of their interactions and services to a digital, internet based, support. Within such environments, there are already emerging protocols, internet based, taking the first steps towards user privacy. These protocols, such as OpenID [3] and Cardspace [4], leverage identity models to support strong authentication between user and services. In the process of providing identity to services, a user may choose to use one of several owned identities, depending on the service. This paradigm consists of representing different perceived views of the user, known as *virtual personae*. Ideally, different *personae*, or identities from the same user, can be provided to the same service without any correlation among them.

While these steps towards identity and privacy are noteworthy, their domain is the application layer. It is unacceptably easy to unwillingly forfeit privacy by disclosing linkable network identifiers. We believe that in order to provide privacy in mobile environments, the application layer models should have repercussions throughout the entire network stack. We present a solution to preserve the application layer models by applying the *virtual personae* concept throughout the network stack. We also propose mechanisms for non-correlation between identities in 4G mobile environments.

The paper is organized as follows: Section II presents some of the related work, while section III addresses the virtual identities concept and cross-layer correlation issues. Section IV briefly describes our multi-homed-based network model, and

The work described in this paper is based on results of IST FP6 Integrated Project DAIDALOS. DAIDALOS receives research funding from the European Community's Sixth Framework Programme. Apart from this, the European Commission has no responsibility for the content of this paper. The information in this document is provided as is and no guarantee or warranty is given that the information is fit for any particular purpose. The user thereof uses the information at its sole risk and liability.

section V proposes the concept of virtual network stacks along with the network stack repercussions. Section VI presents how mobility can be handled without correlation between identities, and Section VII briefly discusses the benefits and issues of the proposed architecture. Finally, our conclusions and future work plans are presented in Section VIII.

II. Related Work

Several proposals in the literature already address privacy issues. Here we consider some of the proposals that address location privacy and user privacy at different layers.

One of the attempts to address user privacy issues in mobility environments is addressed in [5]. This work efficiently describes the problem, but does not provide a cross-layer solution or application layer integration. In an attempt to address privacy at the lower layers, [6] includes the concept of link and network layer pseudonyms for each user application. Although this proposal is able to avoid correlation up to the network layer, it fails to provide application integration. Also, different pseudonym sets for each connection, regardless of *persona*, introduces a significant cost on the network. In [11], the ground work for privacy in mobile environments is tackled, providing a complete system aimed at anonymity, emphasizing location protection. It provides several useful requirements and a solution that shares common points with our view, but it does not mention mobility correlation issues, such as mobile identifiers or handovers, nor does it bridge application layer identity schemes or NGN networks.

In a different scope, location privacy has already gained attention. A framework supporting location privacy is proposed in [2]. It is able to prevent network layer attacks and a small degree of privacy. However, it focuses on location and fails to provide a full cross-layer solution. A proposal for location privacy using the Host Identity Protocol (HIP) [7] is described in [1]. This proposal brings together identity concepts and location privacy, but does so at the network layer only. Other layers are not addressed.

A mechanism supporting privacy support at the link layer is proposed in [8] effectively mitigating the major attacks at only the proposed layer. Upper layers are not discussed and there is no strong concern towards identity.

Finally, application layer identity schemes exist such as OpenID [3], where the identity is in fact a URL. Even though a URL can point to different identities, the protocol makes no assumption or consideration about the lower network layers.

All these proposals provide one step further towards the support of integrated solutions, but fail to realize a global cross-layer view that effectively provides privacy and integrates with nowadays identity schemes.

III. Virtual Personae Aren't Enough: Threat Model

We have already discussed the validity of the application layer models, stating that they are not enough: even if a user presents different identities to different services, the network stack shows the linkage between identities breaking the application layer privacy model. As shown in Fig 2, the lower network stack layers use identifiers that are common to the different *virtual personae*, rendering them useless in terms of privacy. At the transport layer, the used endpoint identifier, usually the IP address, allows linking different *virtual personae* because every application uses the same identifier. At the network layer the same IP address is used for every transport and application connection, providing the cross-linkage of the *virtual personae* through the common identifier. Following the same reasoning, link layer addresses can bind higher layer addressing structures, whether they are IP addresses, transport identifiers or *virtual personae*.

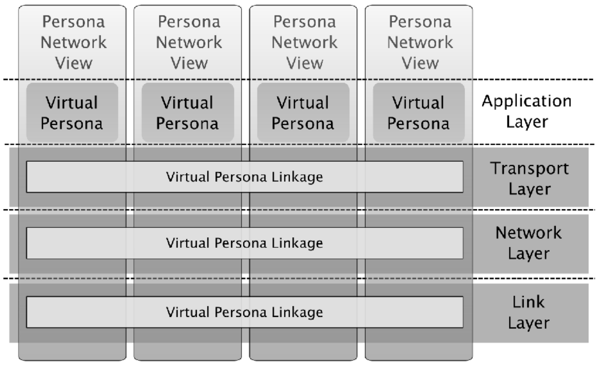

Figure 2. Using cross layer information to link identities.

Our objective is to support the concept of cross-layering *virtual personae*, making their identifiers distinct and non-linkable for different virtual users. Formalizing the threat model, in a privacy-enabled network an attacker cannot be able to:

- Correlate two identities by mapping them to the same endpoint identifier;
- Correlate two identities by mapping them to the same locator, or different endpoint identifiers to the same locator;
- Correlate two identities by mapping them to the same link layer identifiers, or mapping endpoint identifiers (or locators) to the same link layer identifiers.

The mathematical model that defines the previous conditions is based on Naive Set Theory [10], where a *persona* *k* is characterized by the set of its identifiers (V_k) used across the OSI layers, as shown in expression (i).

$$V_k = \{L2_{ID}, L3_{ID}, L4_{ID}, L7_{ID}\} \qquad \text{(i)}$$

$$A = \{V_k\}, 0 < k < n \qquad \text{(ii)}$$

Set *A*, in expression (ii), represents the user's identity set, which is in fact the complete set of *n* different *personae* belonging to the user. To guarantee that a set of *n* identity *personae* are not correlated, they must obey by expression (iii), which represents the pair wise correlation, *i* and *j*, of all the user's *personae*, resulting in a void set of information. Every intersection of identifiers must not yield any match, i.e. the identifier in two different identities needs to be distinct.

$$V_i \cap V_j = \varnothing, i \neq j \wedge (0 < i, j \leq n) \qquad \text{(iii)}$$

$$V'_s = V_i \cup V_j, i \neq j \wedge (0 < i, j \leq n) \qquad \text{(iv)}$$

$$A' = A - V_i - V_j + V'_s \qquad \text{(v)}$$

However, if expression (iii) is not true, it indicates that there is a correlation between the sets of identifiers. An attacker must aggregate the *personae* that yield positive results according to (iv). The result is a set V'_s which contains all the identifiers present in both identities and joins them under the same *persona s*. V'_s becomes the identity set that supersedes the 'now' partial identities containing both V_i and V_j. The vision of the attacker now becomes more accurate and set *A'* can be narrowed by expression (v).

IV. Network Model

The network model, which serves as reference for our proposed solution, is based on 4G scenarios. These environments contain heterogeneous access technologies, such as WiFi, Wimax or DVB, seamlessly integrated into the global architecture. Similarly, user terminals are evolving into multi-technology devices, capable of sustaining several connections across different interfaces. This multi-technology availability enables a better user experience, serving the 4G paradigm of "always on, always best connected" and creates a real possibility for multi-homing and development of multi-homed based solutions. In heterogeneous multi-homed capable networks, mobility is no longer governed by signal availability, but by user preferences, possibly identity based – each *persona* governs its mobility patterns and selections. In this context, flow [1] based mobility is a well suited candidate for the minimum granularity available to mobility mechanisms. The way flows are distributed through the available interfaces should depend on network availability, provider information, cost and more importantly, on preferences set by the user or, in our privacy model, by the *virtual persona*. This allows the user to take full advantage of the concepts of identity management, *personae* and multi-homing.

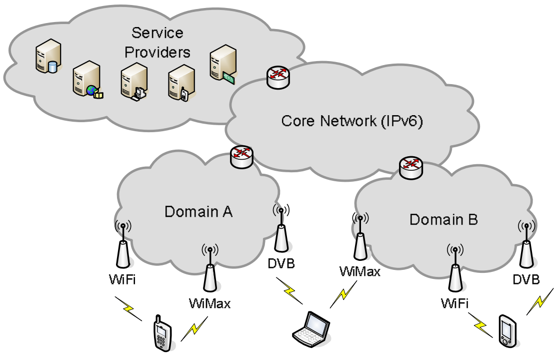

Figure 3. Network Architecture Overview

While the proposed solution to tackle privacy can be applied to any mobility scenario, it becomes particularly interesting in such volatile environments, where a user can distribute flows belonging to different identities across the same interfaces. It would lead, in the current models, to the sharing of similar network addresses, and consequently, to the linkage of identity *personae*, which would increase the threats previously presented. However, multi-homing capabilities can be used to diminish the effects caused by mobility on identity correlation. Considering an example of a terminal with only one interface, when a new network with better signal is sensed, the terminal performs handover and all flows, even belonging to different *virtual personae*, need to handover simultaneously as well, which gives information on the correlation of the different flows and different *personae*.

V. Virtual Network Stacks

The standard use of common identifiers at different layers per user enables the correlation between different *personae*. As presented in [6], a terminal can disguise itself under several layers of pseudonyms. Our proposal to support user privacy considers the concept of multi-layer pseudonyms on an Identity basis, leading to the concept of a Virtual Network Stack (VNS) per identity. Each identity is assigned different L2, L3 and L4 pseudonyms, making it impossible to use these identifiers for *personae* correlation. We assume that, at the application layer, the identity scheme in use provides different pseudonyms for the user, presented at different services.

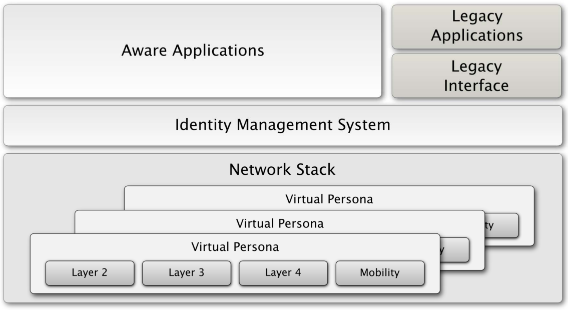

Figure 4. Terminal Cotrol Plane

The support and management of a virtual stack requires terminal architecture modifications. The current legacy model is connection oriented: identifiers are used or generated at the pace they are required to connect to the network at different layers (e.g., an IP address is generated or assigned at the time the terminal connects to an access router, and is normally used by all upper layer protocols). The proposed approach turns the focus to identity, generating different identifiers when an identity wishes to connect to the point of service.

The first step towards this paradigm is to introduce a control layer, that directly instantiates the identity layer functionality, pulling the *persona* concept further down the OSI stack. This control plane interacts with applications, which are used as input for network stack management. As can be seen in Fig. 4, in the terminal control plane, applications might be identity aware and provide specific inputs to the management plane, deciding independently on the usage of identity *personae*; or applications might be legacy, where the management decisions will be extrapolated by a legacy interface component, that analyses the application requirements and selects an appropriate *persona*.

[1] A flow is defined as an application level association between endpoints, which is an aggregation of source address and port, destination address and port, and protocol.

To gain control of the network stack, identity needs to have direct influence on the control plane. However, the data plane is entirely sustained by a VNS: identity operations are not required on data packets (integration of identity into OSI layer is out of the scope of this work). The objective is to provide control over the network stack without linking identifiers that can be correlated. While the VNS concept is fairly simple, it needs to be extended to all mentioned layers, providing the means to instantiate pseudonyms at every layer, taking privacy into account. The next sub-sections address this instantiation over the network stack.

A. Link Layer Pseudonyms

On the Link layer, independent addresses need to be generated for each VNS. These addresses do not correspond to the physical address present in every NIC. This is in fact a virtual address that is created for every *persona*. While it is straightforward to generate addresses, the real interfaces need to be abstracted on the user's terminal, so that each virtual stack is supplied with the necessary device information, such as signal strength or provider availability. This is accomplished through a Virtual Interface Manager (VIM) that controls the real interfaces, and supplies the necessary primitives and information to Virtual Interfaces (VIF) with a Virtual MAC (VMAC) address. It is important to notice that a VNS might be extended to more than one physical device, therefore, creating several VIFs under the control of a particular identity, as show in Fig. 5. From the network point of view, each VIF/VMAC pair represents a different device, competing among each other for network access (otherwise, it would be easy to see at the link layer, especially with IEEE 802.11 protocols, that at each DIFS interval assigned to a particular device, spoofed frames were being received). While the aforementioned abstraction bloats the network stack, it is necessary to cope with the privacy issues, namely to have uncorrelated physical devices and addresses. Moreover, it provides the added benefit of enabling terminals to deal with future access technologies, such as multi-head radios and heterogeneous technologies, coping with the previously mentioned 4G scenarios.

B. Network Layer Pseudonyms

Using pseudonyms at the network layer is simpler than at the link layer. Network identifiers are used as locators to determine the routing path. If we consider IPv6 as our target protocol, the address is usually auto-configured. The support of a Virtual IP address can be achieved by running independent instances of the Neighbor Discovery Protocol (NDP) performing Stateless Address Auto-configuration (SLAC) for each virtual interface. This leads to independent addressing for each identity *persona*. The support of stateful address acquisition methods is also straightforward and requires several instances of a DHCPv6 daemon. While running several protocol instances poses a strain on the device and network, its impact is much lower than in [6], due to the fact that the instances run per *persona*, instead of on a per application basis.

C. Mobility Pseudonyms

Mobility solutions at the network layer tend to create an indirection between the locator, the IP address, and the identifier used by the transport layer, which is the actual endpoint of the transport connections. We discuss two mobility protocols that are easily extended to support the VNS concepts, MIPv6 [9] and HIP [7].

MIPv6: The Layer 3 pseudonyms mentioned in the previous section provide pseudonimity for the Care-of Address (CoA), which is in fact the aforementioned locator. However, for each VNS, a different Home Address (HoA) should be used. This means that each VNS has a different HoA and a different set of CoAs. Each HoA needs to be generated and independently registered at the Home Agent. Again, added signaling is required to support privacy.

HIP: To use HIP each VNS has to generate its own Host Identity, and therefore, generate different Host Identity Tags (HIT) for each identity, that are passed onto the transport layer. HIP has the same problems as MIPv6, in the sense that more signaling is required to support the VNS concept.

D. Transport Layer Pseudonyms

When mobility is in place, using pseudonyms for mobility already makes the transport layer VNS ready, since transport protocols, such as TCP and UDP, establish their bindings with the mobility identifiers, i.e. HoA or HIT. If no mobility solution is used, one can argue that the L3 pseudonyms already provide the necessary VNS support for the transport layer, because the bindings are to the L3 identifiers, and the IP address (IPv6 in our case) is used both as locator and identifier.

E. Application Layer Modifications

We argued that most existent Identity Management solutions work at the application level, such as OpenID [3] and Cardspace [4]. Nonetheless, some modifications are required to keep the user's privacy intact. The common use case for identity is the connection to a service presenting the necessary credentials, or identity claims, to that service (this acts as the identity selection period). The operation mechanism of these identity models considers that, upon initial connection to a service, the user is presented with a graphical user interface, known as an identity selector, from which the user selects the desired *persona*. While this applies to application layer privacy, it breaks lower layer privacy because the user will have a set of identifiers already in use when connecting to the desired service. In fact, the user would be bound to the present VNS, since it would be hard to change the underlying identifiers without having to start a new connection to that service. Furthermore, the change would provide an attacker with the means to link past and present identifier sets.

We argue that the user should be presented with the identity selector as part of the application startup. An example of this model is the case of a user browsing a web site and being prompted with the selection of an identity for that website. This model should be replaced by the identity selector being shown when the browser starts, leading to the paradigm that the identity is already selected upon connection to the website.

F. VNS Application Example

In this section we present a detailed example on an instantiation of a VNS across multiple layers and interfaces.

Our example scenario, illustrated in Fig. 5, consists of two personae that run independent applications spanning over two real interfaces. In this case, two virtual stacks are instantiated to cope with the desired privacy levels.

Each network stack can use one or more application identifiers bound to the specific identity. Going through the layers, for transport, each VNS is bound to one MIPv6 HoA, providing the binding for mobility. The network and link layers are dependant on the number of active interfaces, which spawn into virtual interfaces. From this example, we observe that two real interfaces spawn into four virtual interfaces (two per identity), leading to the requirement of one MAC address and CoA for each virtual interface. We then have two MAC and IP addresses per identity: there will be no linkage between identifiers across virtual stacks, and consequently, across identities.

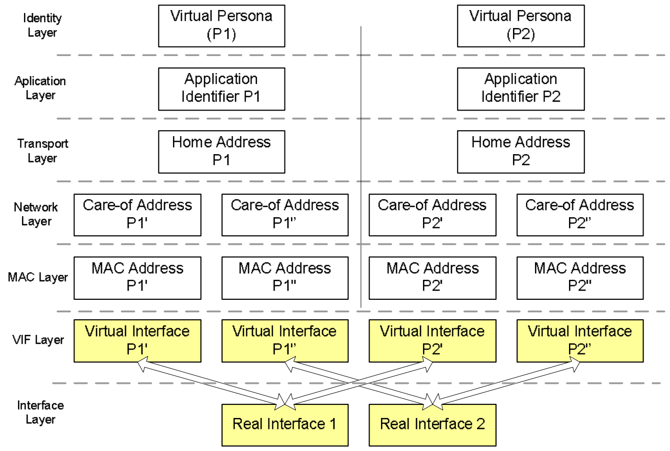

Figure 5. Virtual Network Stacks for two personae

VI. Mobility & Privacy

While the aforementioned concept of virtual network stacks enhances the user's privacy by avoiding the correlation of identifiers across identities, there are other means by which an attacker can correlate identities, e.g., through the observation of mobility patterns. If two identities travel consistently through the same networks, within the same time periods, an attacker can assume that both identities are using the same terminal; and therefore belong to the same user. We present a threat model which extends the previous one (Section III), and discusses how to minimize the inference from mobility patterns and timing attacks on handover.

A. Extended Threat Model

Even though each *persona* has different sets of identifiers, satisfying expression (iii) in Section III, mobility, and in particular the time at which handover occurs, can still link identities due to the fact that these identities are bound to the same physical terminal. This attack does not imply correlating identifiers, but instead specific events linked to the identifiers that occur from different identities at approximately the same time. There are possible attacks ranging from the physical layer to the application layer.

At layer 2 there is the possibility to correlate wireless events that occur within a specific time window. These events pertain mainly to IEEE 802.11 association request, which corresponds to a new station attaching to the access point. At layer 3, address acquisition events, such as simultaneous duplicate address detection, when using SLAC, and DHCP requests when using managed infrastructure to obtain the address, can also be used to trace the user. Mobility also offers a broader range of mechanisms and associated events, which are highly dependent on the protocol. Using MIPv6, simultaneous binding updates imply simultaneous movement (update messages in HIP or redirect messages in SIP). In fact, any mobility protocol can be correlated by applying the same strategy – link two independent messages that occur simultaneously. At the transport layer, attacks are less efficient, since the use of identifiers is reduced to TCP or UDP ports. One specialized attack is the probe for TCP connections for which no handshake occurred. Application layer attacks are harder to mitigate, but also harder to achieve, since they require knowledge over several protocols, such as *Zeroconf daemons*, that probe the network for resources.

All previous attacks can take advantage of any layer specific signaling required for mobility procedures. Different *personae* can be easily linked, generating an extended set of identifiers, if the condition present in expression (vi) is true. It shows that, for any event *E* that occurs at layer *l* belonging to different identities *i* and *j*, they can be correlated to the same terminal if the time between them is smaller than a given *Threshold*. Once (vi) is verified, an attacker can again use (iv) and (v) to generate the set with correlated identities.

$$\left|E_l V_i t_i - E_l V_j t_j\right| \le Threshold \qquad \text{(vi)}$$

Another type of specialized attack consists of gathering mobility patterns. If two identifier sets move at constant intervals, then an attacker can infer that they belong to the same physical device and the user is ineffectively trying to conceal their relationship. This can be achieved by gathering several time intervals of different events, represented in expression (vii), and intersecting those intervals. Correlating the gathered values is given by expression (viii), which states that if the difference between two time values is within a 2*t* interval centered on zero, then they are close to each other and may be running on the same terminal. The value *t* is the metric that determines the granularity of the event times: a large value *t* will create highly inaccurate linkages.

$$\Delta E_l t_{ij} = \left|E_l V_i t_i - E_l V_j t_j\right| \qquad \text{(vii)}$$

$$\left|\Delta E_l t_{ij} - E_l \Delta t_{xy}\right| \subset \left[0-t;0+t\right] \qquad \text{(viii)}$$

While the aforementioned attack is feasible, it requires interpolation of several factors. It yields probabilistic results, dependent on *t*, requiring control over several network points to produce any efficient linkage output. Therefore, the performance drop in defending against this attack might not overweight the benefit.

B. *VNS and Mobility Dissimulation*

To effectively mitigate the aforementioned attacks, we propose two separate solutions which depend on whether multi-homing is available or not. If multi-homing is available, we can make use of it to separate the movements of different flows in different identities; if not, we need to move every identity to the new network while avoiding time correlation.

1) *Defferred Handovers*

Using deferred handovers is a synonym for moving one *persona* at time. This is the worst case scenario where no multi-homing is available and the handover is mandatory or imminent. *Personae* should be moved one at a time, generating de-correlated events for each layer. The handover selection mechanism should take into account *personae* priorities: there are only a few applications that require minimum blackout periods during the handover, such as VoIP. The identity associated with such priority applications should be first moved, while other applications, such as file transfers, can sustain longer blackout periods, from which transport protocols (e.g. TCP) can recover. In order to avoid correlation described in (viii), the *personae* should be moved with a delay based on a random backoff period.

2) *Partial Terminal Handovers*

With available multi-homing technologies, the *personae* can be evenly distributed across interfaces, governed by user preferences; it also means that more destinations for handoff are available. We propose partial terminal handoffs, coupled with deferred handovers. Partial handovers use an identity oriented approach, where a user moves only the necessary identities, while leaving the remaining identities on different interfaces, avoiding unnecessary correlation. Coupling partial handovers with deferred handovers is a simple process: if more than one identity requires a handoff to the same destination, the deferred handover approach can be used. In this case, the blackout period would be minimal, since during the deferred random backoff time, the flows are not in a disconnected state. Therefore, multi-homing greatly reduces the downsides of deferred handovers, at the cost of some mobility optimizations.

VII. Discussion

Introducing Virtual Network Stacks provides the necessary independence between different user identities. At every layer each identity has disjoint identifiers, thus providing the necessary privacy that avoids the correlation. Having different identifiers at each layer mitigates the attacks previously described. But, privacy comes at a cost to the user. In order to provide virtual network stacks, the terminal in practice becomes several terminals. This implies that there is some network overhead, which was previously inexistent. Assuming that the user has *N* distinct *personae*, the network mechanisms increase by a factor *N*. At the link layer, assuming IEEE 802.11, this implies the increase of Associate Request/ Response messages in order to maintain the VNS model. Furthermore, the network occupation increases since, instead of one period to transmit and receive frames, the terminal now requires *N*, which leads to an increase in network allocation (because the carrier sense mechanism is applied independently to each *persona*). At the network layer the effort resides in the address acquisition mechanisms, but the overall impact is not as large as for L2, due to the fact that either DHCP or SLAC can run concurrently in each virtual stack. The enhancement to the mobility scheme incurs on the same cost, multiplying the number of messages by *N*, while the effect is minimized by concurrently handling these messages. Such overhead in signalling increases the terminals' power consumption, since the extra carrier sensing reduces the possibilities of the terminal to enter power save mode, and it also impacts the transmit and receive periods, when compared to the normal stack operation.

VIII. Conclusions and Future Work

In this paper we presented a threat model which is a result of current trends in the way users interact with the digital world. The mechanisms used to enhance user privacy at application layer can easily be thwarted, if not accompanied by strong privacy network support. We have also shown how the effect on linking of different identifiers at different layers can be mitigated by applying techniques from the literature. We further extended the model, and performed the same type of mitigation analysis, in cases where information on events related to the identifiers is used.

Our work focuses on the identifiers and the privacy lost but not on where this linkage is performed. One extension of this work could be to evolve our threat model to include the players and the distance from the user to the place in the network where linking is possible. This may result in better tradeoffs between performance and privacy.

References

[1] Alfredo Matos et al., HIP location privacy framework. In First ACM/IEEE MobiArch 2006, San Francisco, USA, December 2006, held in conjunction with GLOBECOM 2006.

[2] Joao Girao et al., A practical approach to provide communication privacy. In ICC2006, Istanbul, Turkey, June 2006.

[3] D. Recordon et B. Fitzpatrick, OpenID Authentication 1.1, May 2006.

[4] Windows Cardspace, Internet Address: http://cardspace.netfx3.com/content/introduction.aspx.

[5] C. Hauser, Mobility management meets privacy: the failure of existing proposals and a new, future-proof approach. In International Conference on Mobile Computing and Networking, Philadelphia, PA, USA, 2004.

[6] D. Gomes, R. L. Aguiar, Privacy through Virtual Hording, Globecom 2006, San Francisco, CA, USA, Dec. 2006.

[7] R. Moskowitz, "Host Identity Protocol." Internet Draft (Work in Progress), July 2006.

[8] F. Armknecht, J. Girao, A. Matos, R. L. Aguiar. "Who said that? privacy at link layer". In INFOCOM 2007. Minisymposium, Anchorage, Alaska, USA, May 2007.

[9] D. Johnson, C. Perkins, and J. Arkko, "Mobility Support in IPv6." RFC 3775 (Proposed Standard), June 2004.

[10] NJ, 1960. Reprinted, Springer-Verlag, New York, NY, 1974, ISBN 0387-90092-6.

[11] Alf Zugenmaier, "Flasche – A meschanism Providing Anonymity for Mobile Users", in Privacy Enhancing Technologies – 4th International Workshop, pp. 121-141, Toronto, Canada, May, 2004.