Enabling Seamless Internet Mobility

Anja Feldmann
DT Labs / TU Berlin
anja@net.t-labs.tu-berlin.de

Gregor Maier
DT Labs / TU Berlin
gregor@net.t-labs.tu-berlin.de

Wolfgang Mühlbauer
DT Labs / TU Berlin
wolfgang@net.t-labs.tu-berlin.de

Yevgen Rogoza
TU München

***Abstract*—Mobility is a requirement not appropriately addressed by the original design of the Internet since an IP address has two fundamentally different tasks. It specifies a network location (for routing) and serves as an application identifier. A plethora of suggestions have been made to overcome this, e.g., Mobile IP and HIP. Yet, each of the proposed solutions has drawbacks such as requiring fundamental changes to the Internet architecture or relying on triangular routing.**

We propose the *Seamless Internet Mobility System* (SIMS) for enabling seamless IP network layer mobility. Our goals are (1) to enable mobility even for users that do not have a permanent IP address and therefore cannot rely on a Mobile IP home agent; (2) to impose no overhead for applications initiating network traffic in the current network; (3) to preserve sessions that started in any previously visited network location; (4) to be robust, scalable, and easily deployable in the current Internet; (5) address the economics of roaming between different administrative domains.

The key ideas are to allow any new connection to use the current IP address and to take advantage of the heavy-tailed nature of connections. This implies that after a network change only a small number of connections need to be retained.

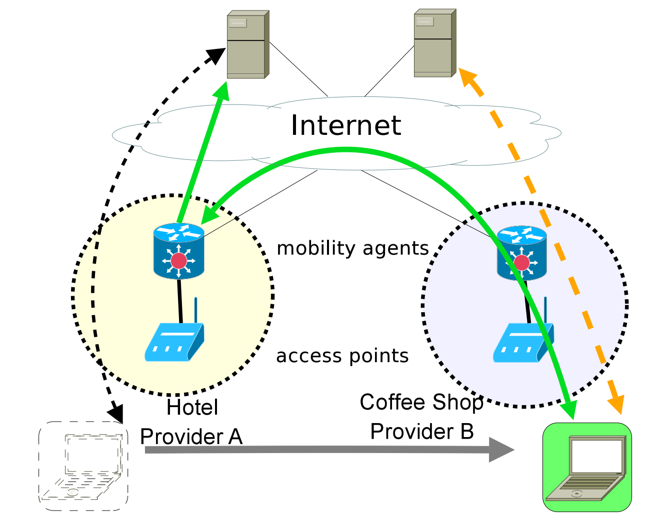

Fig. 1. Scenario addressed by SIMS— new sessions (dashed lines) are routed directly — existing session are maintained by relaying them via the previous network (solid lines).

I. Introduction

Due to increasing miniaturization and decreasing operational cost of mobile devices their use has reached the point where we can hardly imagine life without them. For example, according to an EU report [1] the mobile phone subscriptions in the EU outnumber the citizens – in 2006 the average penetration rate was greater than 103%.

Yet, while seamless use of mobile devices is no problem with cellular technologies such as GSM, this is not the case when relying on the Internet protocol suite. To overcome this we, in this paper, propose *Seamless Internet Mobility System* (SIMS) [2] for enabling seamless IP network layer mobility to everyone without making changes to the IP protocol suite. The example in Fig. 1 illustrates the scenario addressed by SIMS. Initially the user of our mobile node accesses the Internet at a hotel via the network of provider A, e.g., via a wireless access point. Then he moves to a coffee shop across the road and reconnects to the Internet but this time via the network of provider B. Ideally, such mobility should be *seamless* in the sense that any user can use it and maintain his workspace, including all existing network connections without manual configuration and with minimal network overhead. It should allow for roaming between hotspots that are operated by different service providers in airports, train stations or on a university campus.

While an immense number of approaches, including Mobile IPv4, Mobile IPv6, and HIP [3]–[5] have been suggested to extend the Internet with support for mobility, none of them enables the seamless mobility needed for the above scenario. This lack of support is either due to the limited deployment of approaches that require fundamental changes to the Internet architecture (HIP and IPv6), or due to the limited integration in todays Internet (e.g., Mobile IPv4). Mobile IPv4 assumes that a mobile node has a permanent IP address and access to a home agent that can track the current network location of the mobile node. But today most hosts have to use an IP address that is dynamically assigned to them by their connectivity provider, typically via Radius or DHCP.

The fundamental problem with adding mobility to the current Internet architecture is the mangling of two fundamentally different tasks in one entity — the IP address. The first task is to serve as an identifier for addressing an application running on the host. As such the IPv4 address is part of the socket[1] data structure and hence part of any connection identifier. Therefore, an IPv4 address change closes all active connections, making seamless mobility impossible unless using extensions such as Mobile IP [3], [4] or HIP [5]. The second task of the IP address is to specify the location

[1]The socket API is the API of the Internet.

of the network interface within the Internet routing system. Currently, seamless mobility within a single IP network is possible when supported by the layer-2 technology, e.g., within the WLAN network of an organization, but not across different IP subnetworks of the same network access provider or even between different providers.

We tackle the problem of seamless mobility without changes to the Internet architecture even for the case when a user does not have a permanent IP address. Our solution leverages the fact that most users do not care about ubiquitous reachability. Most of those users who do care are using solutions like dynamic DNS [6]. Therefore, the mobility problem reduces to maintaining one's workspace, including all existing network connections, when moving between networks with minimal overhead.

We rely on the observation that most of today's network stacks are able to use multiple IP addresses per interface. Therefore, it is possible, after a move from one subnetwork to another, to add the newly assigned IP address from the new subnetwork to the interface. This allows connections initiated after the move to use the IP address that is native to the current subnetwork without imposing any overhead. But what about existing connections? Here we borrow some concepts from Mobile IP. Every subnetwork offering mobility services runs a mobility agent that acts as home agent for any connection initiated while the mobile node was in the subnetwork. When a mobile node enters a new subnetwork it contacts the local mobility agent. The local mobility agent then contacts the mobility agents of any previously visited subnetworks and requests that they forward any further packets, e.g., via a tunnel, to the mobile node. Fig. 1 shows the corresponding data flow. If the mobile node moves back to any previously visited network the tunneling is stopped and direct communication is reestablished.

In summary, the key idea is to allow any new connection to use the current IP address, taking advantage of the heavy-tailed nature of connections [7]. With the majority of sessions being short-lived, only a small number of connections need to be retained after a move.

We have implemented these ideas within the *Seamless Internet Mobility System* (SIMS) and show that SIMS (1) enables mobility even for users that do not have a permanent IP address and therefore cannot rely on a Mobile IP home agent; (2) imposes no overhead for applications initiating network traffic in the current network; (3) preserves sessions that started in any previously visited network location; (4) is robust, scalable, and easily deployable in the current Internet; (5) addresses economic issues of roaming between different providers.

The remainder of this paper is structured as follows. In Sec. II we give an overview of the Mobile IP standard, while Sec. III reviews related work. In Sec. IV we present the architecture of our system. Next, in Sec. V we discuss the benefits of SIMS and compare our architecture with other mobility solutions. Finally, in Sec. VI we conclude and outline future work.

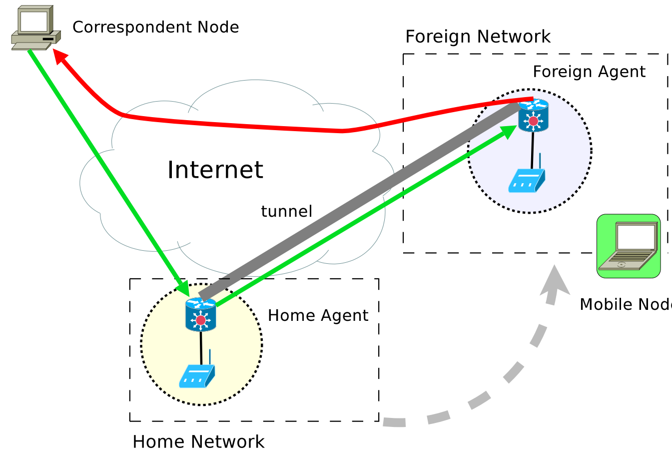

Fig. 2. Mobile IP.

II. Background: Mobile IP Review

Mobile IP (MIP) comes in two flavors: Mobile IPv4 [3] and Mobile IPv6 [4]. While the basic principles are the same, Mobile IPv6 (MIPv6) offers additional capabilities as IPv6 is inherently designed to offer support for mobility.

MIP relies on three architectural components (see Fig. 2): mobile nodes, home agents, and foreign agents. A *mobile node* is any IP-capable host that may change its point of network attachment. Each mobile node has a home network which provides it with a permanent IP address (its home address) from the address space of the home network. Moreover, the home network offers a *home agent* that is responsible for tracking the current location of any mobile from this network. If a mobile node is not within its home network, datagrams are tunneled by the home agent to the remote network for delivery to the mobile node. When a mobile node visits another network it obtains a so-called "care-of address". It is the task of the *foreign agent* of the visited network to assign the care-of IP address to the mobile node. This address is from the address space of the visited network and the mobile node registers this address with its home agent. The foreign agent serves as default gateway for the visiting mobile node. As such it can serve as end-point of the tunnel to the mobile node's home agent and decapsulates and forwards packets received via the tunnel to the mobile node.

Now the mobile node is reachable via its permanent IP address. The hosts communicating with the mobile node, called *correspondent nodes*, do not even know that the mobile node is currently in another network. Datagrams from the correspondent nodes are sent to the home network, where they are intercepted by the home agent. There they are encapsulated and forwarded over a tunnel to either the mobile node or its current foreign agent. In the reverse direction, Mobile IP can route the packets directly from the mobile node to the correspondent node, using the home IP as source address. This is referred to as *triangular routing* and only works if the foreign network and its provider does not use ingress filtering [8]. Note that the use of ingress filtering is not only highly recommended to avoid address spoofing but also part of "best common practice" of providers. Mobile IPv4's packet flow and rounting are depicted in Fig. 2.

Mobile IPv6 [4], [9] offers some additional functionality. MIPv6 offers two communication modes: With *bidirectional tunneling* all packets are tunneled to the home agent, even packets sent by the mobile node to the correspondent node. This ensures that ingress filtering does not hinder communication but it introduces delays for both paths. *Route optimization* eliminates the need to tunnel packets and therefore enables faster and more reliable transmission. A mobile node registers its care-of address and home address (*binding*) with the correspondent node, which can then send the packets directly to the mobile node, using the care-of address. The binding process requires an exchange of IPv6 Binding Updates between the correspondent node and the mobile node (see [9]). Signalling a hand-over requires contacting the MIP home agent, which imposes a delay that for some scenarios is considered to be too large.

III. Related work

The requirement of supporting mobility on the Internet is rather old — discussion has already started in 1992 and in 1996 Mobile IP has been standardized as RFC [3]. Yet, all proposed solutions require either fundamental changes to the Internet architecture or inadequately address simple scenarios. Related work can be roughly classified into three categories: network layer solutions, shim layers between the network and transport layer, and application layer solutions.

Network layer solutions: One can distinguish between proposals that introduce an alternative network layer and those that augment the existing Mobile IP standards [3], [4] (see Sec. II).

Examples of the former generally require major changes to the Internet architecture. For example, ROFL [10] suggests to route using a flat label space, or role based routing [11] which proposes a new non-stack network architecture. One of the reasons for even considering such approaches is the need for separating the locator and identifier functionality of IP addresses. Numerous new network architectures and research initiatives, including NewArch [12], Daidalos [13] and FIND [14], therefore address this problem.

Koodli [15] in "Fast Hand-overs for Mobile IPv6" suggests to enable a mobile node to detect that it has moved to a new subnet by providing the new access point and the associated subnet prefix information while the mobile node is still connected to its old network. Singh [16] in "Reverse Address Translation" proposes to leverage NAT for traffic delivery rather than tunneling which is easier to setup. Castelluccia [17], Mao et al. [18] and Wakikawa et al. [19] propose to rely on multiple home agents to reduce layer-3 hand-off times. "Hierarchical Mobile IPv6" [17] suggests to use additional agents, called mobility anchor points (MAP), to localize binding update messages destined to the home agent. Contrary to these Mobile IP extensions, our solution does not require users to have a permanent IP address or access to a Mobile IP home agent.

Shim layer: Proposals in the second category (e.g., [5], [20]) add a level of indirection between the network layer and the transport layer. The Host Identity Protocol [5] introduces a new layer in the TCP/IP stack. This layer uses identifiers based on public keys and hides IP addresses from the layer above. However, this means that upper layers and sockets need to be modified to support such identifiers.

Application layer solutions: Application-layer solutions make up the third category and provide mobility *only* for a *specific* application. The Session Initiation Protocol (SIP) [21] is an application-layer solution for establishing multimedia sessions like Internet telephony calls or multimedia conferences. "Migrate" [22], [23] is a session-based architecture which leverages application naming services to provide communication between end points.

Other approaches (e.g., [24], [25]) propose to unambiguously associate an identifier with each host that stays the same even if the host moves between networks. To map the identifier to the currently appropriate IP address and vice-versa a level of indirection, via peer-to-peer technology, is used.

IV. Seamless Internet Mobility System

In this section we give an overview of the design requirements for SIMS followed by an overview of the proposed architecture.

A. Design Requirements

The problem of mobility consists of two parts: *reachability* by others and *persistence* of work space. Unlike MIP we focus on persistence. Most users either do not care about reachability or have been forced to address it using existing solutions such as dynamic DNS [6]. We also do not propose a new architecture to tackle the problem of location and identifier functionality. Rather we aim at providing a solution that can be deployed immediately and incrementally even in todays IPv4 Internet, yet provides a solution to the scenario shown in Fig. 1. Our requirements for SIMS are the following:

Mobility without permanent IP address: Today a typical Internet user gets access to the Internet via their local Internet provider. These providers generally do not allow their users to obtain permanent IP addresses. Rather they dynamically assign IP addresses, e.g., via DHCP, to the users. Moreover, almost none of the ISPs currently offers a MIP home agent to their users. Therefore, today MIP is not available to a typical Internet user. Our system should enable anyone to use mobility.

No overhead for new sessions: Existing mobility solutions impose a significant overhead on all sessions, e.g., MIPv4 relies on triangular routing, MIPv6 relies on binding updates or tunneling. While it is impossible to offer mobility without overhead, we propose to differentiate between connections started before moving to a new network and those started after moving to a new network. While it appears impossible to avoid adding some overhead (at least during the move) to already established network connections, new connections should not suffer. Our solution, therefore, aims at adding no overhead to either the signaling or the data path for sessions started in the current network.

Preservation of sessions: To achieve seamless mobility existing network sessions have to be retained. First, this implies that hand-overs have to be transparent to the application layer (e.g., a SSH or FTP session). If TCP is used, the IP address needs to be kept, as it is part of the connection identifier. After moving to another network, the previous IP address continues to be used for connections initiated in the previous network, while "new" connections use the IP address of the current network. Second, preserving existing sessions during a network change requires low hand-over latencies to avoid session termination due to timeouts.

Robust, scalable, easy to deploy: In theory the most simplistic solution to the persistence problem is to offload it to the routing system by asking it to use host routes. However the routing system cannot handle it by itself as it is already reaching its scalability limits (see e.g., [26]). As the system should be incrementally deployable it also is not possible to change the fundamental network architecture, the control plane, or the networking stacks of all servers. We are thus limited to using the protocols as they currently exist. Nevertheless, as soon as two providers introduce SIMS as a service, mobile nodes should be able to switch between these networks.

Roaming: More and more WLAN hotspots can be found in public places, such as hotels, coffee shops, and airports. Unfortunately, they are frequently not administered by the same authority. Yet, it is very convenient for end users to "roam" between such networks. Therefore, we envision an architecture which inherently enables network authorities to implement such roaming services.

B. Architecture

The architectural design of SIMS is based on two key observations: First, the majority of mobile users solely want to maintain their workspace when moving to another network. Generally, they do not have to be reachable via a permanent IP address. Second, the vast majority of connections in the Internet is very short-lived [7], [27], [28]. Therefore, only few sessions need to be retained when moving between different networks. Accordingly, SIMS works with dynamic IP addresses. It does not require permanent IP addresses or a home agent. To describe the architectural design of SIMS we use the following terminology:

- **Mobile Node (MN):** A host that can change its point of network attachment from one subnetwork to another. Mobile nodes can include all devices that implement the IP protocol (e.g., laptops, cell phones).
- **Correspondent Node (CN):** A peer with which a mobile node is communicating. Frequently, a CN provides some service, e.g., a Web, SSH, VPN server which is accessed by roaming mobile nodes.
- **Mobility Agent (MA):** A MA is a router within a subnetwork which provides the SIMS routing services to any mobile node currently registered in the subnetwork. To enable seamless mobility every subnetwork that offers the SIMS service needs to have a MA. When a MN moves, the MA can in cooperation with a remote MA use tunneling and/or network address translation to preserve the connections of the MN.

SIMS does not rely on permanent IP addresses and as a consequence gives up the notion of centralized home agents. Our solution solely requires one MA in each subnetwork that wants to offer mobility services to its users. It presumes co-operating MAs that exchange routing information to preserve existing sessions.

On first glance this looks like a lot of work and a lot of overhead. But this is where the second key observation comes into play. We can take advantage of the heavy-tailed nature of connections [7], [27], [28]. For example, Miller et al. [7] found that the average flow duration of TCP connections is less than 19 seconds. Hence, we can safely assume that there are not that many sessions lasting longer than a few minutes.

Our approach utilizes this observation by differentiating between "new" sessions, initiated in the current network and "old" sessions that have been started in a previous network. Whenever a "new" session is established, an IP address from the address space of the new network is used. Packets are directly forwarded based on the routes computed by standard IP routing protocols. No overhead is imposed for these. On the other hand, there is going to be a small number of ongoing sessions For these sessions we use a similar mechanism as in Mobile IP: Packets from the MN are encapsulated by the MA of the current network and sent, e.g., over a tunnel, to the MA of the "old" network. From there they are forwarded to the CN (see Fig. 1). To provide IP-layer transparency to the application layer of the CN and of the MN, we need to continue using the IP address assigned by the previous network. This design ensures that we do not introduce any overhead for "new" sessions and only minimal overhead for "old" sessions.

Having explained the basic design of our architecture, we need to discuss how a mobile node can discover the MA, how traffic forwarding is actually handled and how we maintain state in a scalable way.

- **Agent discovery:** To establish contact with the MA of the current subnetwork and to obtain an IP address, the MA can either broadcast advertisements at regular intervals or the MN can explicitly search for MAs via broadcast or multicast messages.[2] Note that layer-2 connectivity (e.g., association with a wireless access point) is required before the layer-3 hand-over can be initiated.
- **Traffic forwarding for existing sessions:** Any traffic originated by a mobile node that belongs to an "old"

[2]Stateless address autoconfiguration and neighbor discovery in MIPv6 simplifies the agent discovery process. Still, the basic mechanism stays the same.

connection is intercepted by the current MA and passed to the previous MA, e.g., over a tunnel. Packets of "old" connections use the IP address of the "old" network as source address — packets of "new" connections use the new IP address. At the "old" MA packets are decapsulated and forwarded to the CN. Communication in the opposite direction from the CN to the MN is tunneled between the "old" and the "new" MA in a similar fashion. In general, there are multiple IP addresses associated with a MN: addresses assigned by the current network and by previous networks. Based on the source IP addresses the MA can decide whether to forward packets directly to the CN or whether to pass them over some tunnel to the "old" moblity agent. Note that the MA does not have to establish too many tunnels as it only has to communicate with MA's of networks with which its provider has a roaming agreement.
- **Keeping state:** For correct forwarding, routing state needs to be maintained. In our architecture each mobile node is in charge of keeping enough information to enable its own mobility. It stores information about all MAs, with which it has been associated and for which an ongoing connection still exists. Whenever a MN changes its network, it provides the new MA with the relevant information to set up the tunnels for the retention of the "old" sessions. Implementing this responsibility on the client is not a big burden for the client but ensures scalability and incremental deployment. After all the client can be expected to install a small program before it can use the SIMS service.

V. Discussion

Even after Mobile IP has been standardized in 1996, a plethora of suggestions have been made to extend the Internet with mobility. This shows that none of the proposed solutions has been satisfactory. We believe that the design goals of Sec. IV-A accurately summarize the essential requirements for mobility in the Internet. In this section we discuss how SIMS meets the design goals.

The benefits of our system can be best understood by comparing SIMS with other approaches. For this purpose we compare SIMS against Mobile IP [3], [4] and HIP [5], which proposes an alternative to the dual use of IP addresses as locators and identifiers. In a nutshell: With HIP, sockets are bound to host identities rather than IP addresses. However, each host must at least know one IP address at which its peer is reachable. The required mapping can be performed by DNS or by dedicated rendezvous-servers (RVS).

In the following, we briefly discuss each of our five design goals from Sec. IV-A separately, while Table I summarizes the comparison of Mobile IP, HIP, and SIMS. A "yes" indicates that an approach satisfies the design goal, "?" indicates a partial match, and a "no" indicates no match.

	MIP	HIP	SIMS
No permanent IP needed	no	yes	yes
New sessions: no overhead	?	yes	yes
Short layer-3 hand-over	?	?	yes
Easy to deploy	no	no	yes
Support for "roaming"	no	yes	yes

TABLE I
Comparison of Mobile IP, HIP and SIMS.

1) *Mobility without permanent IP address*: Obviously Mobile IP requires a permanent IP address from the address space of the home network. HIP distinguishes between routing locators (IP address) and end-point identifiers which are then used in transport-layer sessions. Therefore, only SIMS and HIP allow for mobility without assigning permanent IP addresses to mobile nodes.
2) *No overhead for new sessions*: Mobile IPv6 route optimization enables direct data flow between the correspondent node and the mobile node as soon as the binding update has been sent. As such it imposes no additional delay. But not all Mobile IP implementations support binding updates which results in a "?". HIP maps identifiers to IP addresses and then uses IP addresses as routing locators. Hence, HIP and SIMS do not impose overhead for sessions initiated after a network change.
3) *Short layer-3 hand-over times*: First we note that all three approaches are transparent to the transport-layer. Still retaining existing sessions when moving (seamless mobility) requires short layer-3 hand-over times. The time required for signaling depends on the round trip time between a mobile node and the home agent (Mobile IP) or the DNS/RVS (HIP) as such they can vary and at times be fairly large. SIMS has to inform MAs of previously visited networks for which it still has active connections. For most application scenarios we can expect the previous MAs to be geographically close to the current location of the mobile node. Hence, we expect layer-3 hand-over times to be short.
4) *Robust, scalable, easy to deploy*: Neither Mobile IP nor HIP are easy to deploy. The triangular routing schema, used by traditional Mobile IP, is not compatible with ingress filtering, frequently performed by ISPs to prevent spoofing. IPv6 is not yet widely deployed and enhanced Mobile IPv6 features such as route optimization have to be supported by *all* potential CNs to get their full benefit. In particular for servers, we cannot expect this to be the case in the near future. The main drawback of HIP is the need for a rendezvous-mechanism (extending DNS or setting up a dedicated RVS) which provides the (initial) mapping between IP addresses and identifiers. Contrary to HIP and Mobile IP, our proposed solution does not require changes to the existing Internet architecture and therefore can be deployed easily and incrementally. Furthermore it is robust and scalable, as it does not rely on an central infrastructure.
5) *Roaming*: The architecture of SIMS inherently supports roaming between networks of different administrative domains. In contrast, Mobile IP needs to undergo sig-

nificant changes to allow for roaming, e.g., the design of a federation of home networks in order to exchange information between different service providers. As HIP does not have a notion of service providers, roaming is naturally supported without further modifications to the architecture.

SIMS satisfies all design requirements of Sec. IV-A. Numerous applications are conceivable. For example, SIMS enables a network administrator of any major cooperation or university campus to split its wireless network into multiple subnetworks (e.g., one for each department or one for each building) while retaining mobility. Furthermore, airports or other public places may profit by allowing roaming between hotspots, operated by different service providers.

Adding security and accounting mechanisms to SIMS is straightforward. In terms of security, it is necessary to secure the access network (e.g., via WPA) and to protect tunnels between MAs. Since anyone could pretend to be the originator of a session, the architecture of SIMS needs to prevent sessions from being hijacked. This can be achieved by introducing credentials that are generated by the MA of the network where a session is initiated. Accounting requires tracking of intra-provider and of inter-provider traffic. While the volume of intra-domain traffic can be measured by the current MA, inter-provider traffic can be measured at the tunnel endpoints.

VI. Conclusion

Although a plethora of solutions have been suggested, none of them has established itself as the ultimate answer to the mobility problem in the Internet. In this paper we propose *Seamless Internet Mobility System* (SIMS) which can add mobility support to the Internet and is incrementally deployable.

The architecture we propose is based on two key observations. First, the majority of connections in the Internet is short-lived. Therefore, only few sessions need to be retained when moving between different networks. Second, the majority of mobile users only want to maintain their workspace when changing to another network. Generally, they do not have to be reachable via a permanent IP address.

Relying on these two key observations we designed SIMS. SIMS (1) enables mobility even for users that do not have a permanent IP address and therefore cannot rely on a Mobile IP home agent; (2) imposes no overhead for applications initiating network traffic in the current network; (3) preserves sessions that started in any previously visited network location; (4) is robust, scalable, and easily deployable in the current Internet; (5) addresses economic issues of roaming between different providers.

First experiences with a prototype implementation of SIMS are promising. Regarding future work we plan to extend SIMS with security and accounting mechanisms. We also plan to test our system in a real operational environment. Given the simplicity of SIMS, we feel confident that it is an enticing alternative to the numerous solutions for IP network mobility, suggested in the past.

Acknowledgments

The research results presented herein are partly funded by a grant from Deutsche Telekom Laboratories.

References

[1] European Commision, "European Electronic Communications Regulation and Markets 2006 (12th Report)," *Commision Staff Working Document*, vol. 1, pp. 13–15, 2006.

[2] G. Maier, W. Mühlbauer, Y. Rogoza, and A. Feldmann, "Enabling seamless internet mobility," in *CoNEXT'07 Student Workshop*, 2007.

[3] C. Perkins, "IP Mobility Support for IPv4," RFC 3344, 2002.

[4] D. Johnson, C. Perkins, and J. Arkko, "Mobility Support in IPv6," RFC 3775, 2004.

[5] R. Moskowitz and P. Nikander, "Host Identity Protocol (HIP) Architecture," RFC 4423 (Proposed Standard), 2006.

[6] P. Vixie, S. Thomson, Y. Rekhter, and J. Bound, "Dynamic Updates in the Domain System (DNS UPDATE)," RFC 2136, 1997.

[7] G. Miller, K. Thompson, and R. Wilder, "Wide-area Internet Traffic Patterns and Characteristics," *IEEE Network Magazine*, pp. 10–23, 1997.

[8] P. Ferguson and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing," RFC 2827, 2000.

[9] C. Perkins and D. Johnson, "Mobility Support in IPv6," in *Proc. MobiCom*, 1996.

[10] M. Caesar, K. Lakshminarayanan, T. Condie, I. Stoica, J. Kannan, and S. Shenker, "ROFL: Routing on Flat Labels," in *Proc. ACM SIGCOMM*, 2006.

[11] R. Braden, T. Faber, and M. Handley, "From Protocol Stack to Protocol Heap – Role-Based Architecture," in *In Proc. of HotNets-I*, 2002.

[12] DARPA research project, "NewArch Project: Future-Generation Internet Architecture," http://www.isi.edu/newarch.

[13] EU Framework Programme 6 Integrated Project, "Designing Advanced network Interfaces for the Delivery and Administration of Location independent, Optimised personal Services (DAIDALOS)," http://www.ist-daidalos.org.

[14] NSF NeTS Initiative, "Future Internet Design," http://www.nets-find.net.

[15] R. Koodli, "Fast Handovers for Mobile IPv6," RFC 4068 (Experimental), 2005.

[16] R. Singh, "RAT: A Quick (And Dirty?) Push for Mobility Support," in *IEEE Workshop on Mobile Computing Systems and Applications*, 1999.

[17] C. Castelluccia, "HMIPv6: A Hierarchical Mobile IPv6 proposal," *SIGMOBILE Mob. Comput. Commun. Rev.*, 2000.

[18] Y. Mao, B. Knutsson, H. Lu, and J. Smith, "DHARMA: Distributed Home Agent for Robust Mobile Access," in *Proc. IEEE INFOCOM*, 2005.

[19] R. Wakikawa, G. Valadon, and J. Murai, "Migrating Home Agents Towards Internet-Scale Mobility Deployments," in *Proc. CONEXT*, 2006.

[20] G. Su and J. Nieh, "Mobile Communication with Virtual Network Address Translation," Tech. Rep., 2002.

[21] J. Rosenberg, H. Schulzrinne, G. Camarillo, A. Peterson, R. Sparks, M. Handley, and E. Schooler, "SIP: Session Initiation Protocol," RFC 3261, 2002.

[22] A. Snoeren and H. Balakrishnan, "An End-to-End Approach to Host Mobility," in *Proc. MobiCom*, 2000.

[23] A. Snoeren, H. Balakrishnan, and M. Kaashoek, "Reconsidering Internet Mobility," in *Proc. USENIX HotOS*, 2001.

[24] I. Stoica, D. Adkins, S. Zhuang, S. Shenker, and S. Surana, "Internet Indirection Infrastructure," in *Proc. ACM SIGCOMM*, 2002.

[25] H. Balakrishnan, S. Shenker, K. Lakshminarayanan, I. Stoica, S. Ratnasamy, and M. Walfish, "A Layered Naming Architecture for the Internet," in *Proc. ACM SIGCOMM*, 2006.

[26] G. Houston and G. Armitage, "Projecting Future IPv4 Router Requirements from Trends in Dynamic BGP Behaviour," in *In Proc. of Australian Telecommunication Networks and Applications Conference (ATNAC)*, 2006.

[27] V. Paxson and S. Floyd, "Wide area traffic: The failure of poisson modeling," *IEEE/ACM Trans. Networking*, 1995.

[28] K. Park and W. Willinger, Eds., *Self-Similar NetworkTraffic and Performance Evaluation*. Wiley-Interscience, 2000.