

# Wide-Area IP Network Mobility

Xin Hu<sup>†</sup> Li (Erran) Li<sup>§</sup> Z. Morley Mao<sup>†</sup> Yang Richard Yang<sup>‡</sup>

<sup>§</sup>Bell Labs, Alcatel-Lucent, Murray Hill, NJ

<sup>†</sup>University of Michigan, Ann Arbor, MI

<sup>‡</sup>Yale University, New Haven, CT

**Abstract**—IP network mobility is emerging as a major paradigm for providing continuous Internet access while a set of users are on the move in a transportation system. The intense interest on its support has led to the establishment of the NEMO IETF working group and a test-deployment by a major airline equipment vendor – Boeing – on major airline routes. However, the previously proposed solutions are either inefficient or may cause instability to the global Internet. In this paper, we propose WINMO, a simple, systematic, novel solution for wide-area IP network mobility using techniques including route aggregation, scoped update propagation, and packet mobility states. Our solution provides efficient routing when users travel both across autonomous systems (ASes) and within a single AS, generates minimal global routing overhead to prevent global instability, ensures good location privacy, and helps to defend against denial-of-service attacks. Furthermore, our basic scheme (without packet mobility state) is transparent to both clients and servers. Our extensive evaluations demonstrate the effectiveness of our mobility solution.

## I. INTRODUCTION

Seamless mobility is a major challenge facing the Internet. As the Internet becomes a basic infrastructure of our information-based society, the ability to access the Internet anywhere anytime becomes particularly important. Many people depend on the transportation systems (*e.g.*, trains, airplanes, ships, etc.) to achieve physical mobility, and increasingly the common transportation systems are fitted with data connectivity (*e.g.*, hotspot networks). We refer to a network on a mobile transportation device as a mobile network. With increasing commute time and growing mobility, a higher percentage of the population will spend more time on the move. Since using mobile networks provided by the transportation systems presents minimal safety hazard (*i.e.*, users are not drivers or in charge of the vehicle) and can significantly increase productivity, their popularity can only increase.

However, there are significant challenges in implementing mobile networks. As a mobile network moves (*i.e.*, a train, an airplane or a ship), its attachment point to the Internet also moves, changing to different base stations or even to base stations of different service providers. Given the current routing architecture of the Internet, this leads to disruption to applications. However, people expect that their Internet data sessions (*e.g.*, streaming video, or instant messaging) continue seamlessly, just as people expect that a cellular phone conversation continues uninterrupted while they are in transit.

The importance of the problem has motivated proposals to extend the Internet architecture to provide seamless session mobility for mobile networks. The predominant one is to directly use or extend the mobile IP protocol [1], which is designed for host mobility, to support network mobility. In mobile IP, each mobile host (MH) has a home address. If the host has moved out of its home network, traffic destined

to the mobile host will be intercepted by its home agent and tunneled to its foreign agent which in turn forwards the traffic to the mobile host. However, mobile IP and its variants have several problems which make them unsuitable for network mobility. First, mobile IP depends on public home agents, but many users of transportation systems may not have static home addresses or home agents deployed at home. Second, the redirection by the home agent in Mobile IP introduces triangular routing. The extra delay caused by triangular routing can cause serious performance degradation, and can be unacceptable to some important applications [2].

To address the preceding issues, Boeing developed Connexion [2], a commercial service to use BGP to support network mobility. This service removes inefficient routing and leads to positive user experiences (*e.g.*, [3]). Despite its successful real-field technical test by major airlines on major flight routes, there are serious concerns on the scalability of the Connexion approach. This could be one of the reasons contributing to the discontinuation of this service. With increasing deployment of mobile networks, BGP announcements by these roaming mobile networks will generate a large number of BGP updates to the whole Internet. This may cause global instability. Furthermore, Connexion handles only mobility when a mobile network moves across ASes (autonomous systems). The common case, when a mobile network moves within an AS, is not addressed.

In this paper, we design WINMO, an efficient mobility protocol to support wide-area Internet network mobility. Our protocol addresses IP network mobility both when a mobile network moves across ASes and when it moves within an AS.

Our basic scheme handles network mobility across ASes. In this client and server transparent scheme, we exploit the hierarchical structure of the Internet to design an efficient mobility solution to forward packets to a mobile network to its current AS. Using techniques including route aggregation and scoped update propagation, our solution removes routing inefficiency while generating minimal global routing overhead to prevent global instability. We further extend the basic scheme to optimize performance when a mobile network roams within an AS. We design a novel technique based on *packet mobility states* to allow efficient routing without generating routing overhead. Furthermore, the packet states serve as capability and alleviate potential denial-of-service problems.

We conduct extensive evaluations to demonstrate the effectiveness of WINMO. Our results show that, in terms of average path length, WINMO is only 11% more than Connexion which uses BGP as it is. On the other hand, mobile IP based network mobility solution has an average path length which is 89% more than Connexion. In terms of BGP update overhead, WINMO is orders of magnitude smaller than Connexion.

The rest of this paper is organized as follows. In Section II, we discuss related work. In Section III, we present our design rationale and overall architecture. Our interdomain and intradomain techniques are discussed in Sections IV and V respectively. In Section VI, we analyze the properties of our solutions. In Section VII, we discuss some implementation issues. In Section VIII, we present evaluation results. Our conclusion and future work are in Section IX.

## II. RELATED WORK

The importance of supporting mobility in the Internet has motivated many studies lately. The previous work on mobility spans all layers of the Internet layer hierarchy: the link layer, the network layer (e.g., [4], [5], [1], [6], [7]), the transport layer (e.g., [8]), and/or the application layer. Furthermore, there are recent proposals based on P2P (e.g., [9], [10]) or VPN (e.g., [11]). However, the focus of the aforementioned studies is on host mobility, while our focus is on the mobility support of a mobile network, where a set of hosts belonging to the same network move collectively as a unit. Since most host mobility protocols depend on link layer handoffs to trigger mobility support, but such handoffs may not be seen by all nodes moving as part of a mobile network, they may not be applicable for network mobility support. Furthermore, although it is possible to apply some host mobility protocols to each host individually, this leads to significant inefficiency and requires individual infrastructure support for each host.

There are recent studies extending beyond host mobility to network mobility (e.g., [12]). For a survey, please see [13]. In particular, given the importance of the problem and the interests of the industry, IETF has commissioned the Network Mobility (NEMO) working group [14] to extend the existing host mobility protocols or develop new ones to support network mobility in an IPv6 network. As a first step, the IETF NEMO Working Group has developed a basic protocol [15] to ensure uninterrupted connectivity to the mobile network nodes. However, this protocol does not address important issues such as route optimization and handoff.

Since network mobility support without route optimization can cause unacceptable delay performance, several techniques are recently proposed (e.g., [16]). In [2], Dul presents the implementation of Connexion, a commercial service offered by Boeing to use BGP to support network mobility. Although the paper shows that Connexion removes inefficient routing, a serious concern is that the approach may cause globally visible BGP updates. On the other hand, our approach reduces BGP updates and at the same time resolves the issue of route optimization. In [17], the authors propose an SIP based technique for route optimization for network mobility. However, it applies only to applications using SIP. In [18], [19], the authors propose methods to address route optimization. However, the protocols are based on the NEMO basic protocol and do not handle interdomain mobility well. Another major issue of network mobility is handoff. Various handoff improvements for network mobility have been proposed [20], [21].

## III. OVERVIEW

We first present our high-level design and basic architecture. The detailed design will be presented in Sections IV and V.

### A. Design Decisions

We make the following functional design decisions.

- *Global Network Architecture:* A key scalability feature of the Internet is its decentralized architecture consisting of a large number of interconnected ASes. We design our mobility solution for this architecture. As a result, the mobility pattern we handle is that a mobile network roams most of the time within a single AS. However, it may switch to connect to another AS (e.g., when it moves to a different region). This design choice allows the maximum flexibility for the mobile networks and users.
- *Infrastructure Support:* If an AS provides direct attachment points to mobile networks, we refer to the AS as a mobile ISP or MISP for short. We require that only base stations and routers in an MISP contribute to mobility infrastructure support (e.g., tunneling and forwarding). The service providers of an MISP may also contribute limited support (in BGP routing). As a contrast, mobile IP with host addresses requires that home agents be deployed globally to be effective, and this has been a major barrier for its deployment.
- *Addressing Scheme for Mobile Networks:* We assign a mobile network with a fixed network prefix. Each mobile host obtains, for example, using DHCP, a specific IP address (home address) from the prefix, for the duration when the mobile host is part of the mobile network. We focus on IPv4 networks for better deployment possibility, although extensions to IPv6 is also possible.
- *End-host Support:* We require that our solution be incrementally deployable in the Internet. Thus, there should be a clear upgrade path for the mobile hosts and the correspondent hosts (e.g., application servers). To support mobile devices with varying capabilities, we require that our solution be transparent (i.e., no need for modifications) to mobile hosts acting as information access clients in data sessions. Operating systems support of correspondent hosts as data servers should improve performance, but communications with legacy servers should always be possible.
- *Security Association:* We make the explicit design decision that there is no security association between a mobile host and its correspondent host. The existence of such associations would simplify network design, but establishing such associations faces substantial security challenges or requires fundamental change to the Internet architecture (e.g., HIP [5]). Each MISP has at least one AAA server, which has a security association with each mobile network who has signed up the mobility service in the MISP. The AAA server also has a security association with each base station and router inside the specific MISP. The AAA also distributes a group key to the routers.

### B. Performance Requirements

We impose the following performance requirements. For the network infrastructure, we require:

- There should be no or minimal routing overhead as a mobile network moves across attachment points.
- The impact of denial of service from outside the network should be substantially reduced, as such wireless networks typically have relatively low internal bandwidth.

For the end hosts, we require:

- Path inflation be minimal to support low-latency applications.
- The location privacy of a mobile host be protected.

There are tradeoffs in satisfying the preceding performance requirements. In particular, to completely avoid path inflation, the routing system would need to know the exact location of each mobile network at all time. That is, each handoff, regardless of within the same AS or across AS, would trigger a global routing update. This would allow computing shortest paths to each mobile network, but it would certainly cripple the control plane in terms of routing updates. It also conflicts with hierarchical routing to make the routing system scalable.

### C. Architecture and Technique Overview

Given the preceding design decisions and performance requirements, our basic architecture (Fig. 1) is as follows. Our global network architecture is the current Internet structure: a set of ASes interconnected at peering points by BGP gateway routers; a large number of mobile networks distributed in multiple ASes. Each mobile network has a fixed network prefix allocated from the address space of its home mobility service provider. There is a mobility router (MR) inside each mobile network. Each mobile host inside the network obtains an address within a prefix from the MR of the mobile network.




Fig. 1. WINMO's basic architecture.

The objective of our architecture is to design a scalable routing update scheme with minimal routing overhead and at the same time providing good data path performance. We use the following novel techniques to achieve our objectives.

*Mobility community to reduce state kept by ASes:* A major challenge caused by mobile networks is global routing updates when mobile networks move across ASes. Thus, we limit the number of affected ASes by introducing a mobility community attribute in BGP updates to control the propagation of update messages only to providers and stop the update at the first AS which is a common provider of both new and previous AS. Note that this modification does not introduce routing policy conflicts but may cause slightly longer paths.

*Mobile prefix and aggregation routers to aggregate routing state of multiple mobile networks:* Even with mobility community which limits the scope of ASes involved, an AS may still have to keep routing state for a large number of mobile networks. Thus we need the capacity to aggregate routing state for a large number mobile networks. We introduce the notion of *mobile prefixes*, which allows effective aggregation of a large number of individual prefixes of mobile networks that otherwise may be difficult to aggregate. That is, mobile

networks will be assigned prefixes from a small number of large mobile prefixes. This makes it feasible to achieve effective routing and at the same time only a small fraction of the routers in the Internet need to contain detailed routing information on each individual mobile network. We further use techniques including default routes, aggregation routers and tunneling to aggregate mobile prefixes.

*Authenticated packet state to eliminate border router's routing state:* The essence of our technique is common case optimization. In particular, since a mobile network is likely to move within a single mobile AS, we should substantially reduce intradomain routing updates. Furthermore, a common case is that a mobile host communicates with a correspondent host (e.g., a server) from another AS. Thus, when the packets from the correspondent host enter the MISP at one of the border gateway routers, the border gateway router needs to know the current care-of-address of the mobile host to avoid triangular routing. Requiring each border gateway router to store the current location of each mobile network would require significant control overhead and is not scalable for large MISPs. To resolve this issue, our key technique is to use authenticated packet state to avoid triangular routing inside an AS. The mobility state is a secure token that can be decrypted and verified by all BGP routers inside the AS. It carries the encrypted care-of-address and is sent by the mobile host and echoed back from the correspondent host. By decrypting the mobility state, the gateway router can forward the packet directly to the care-of-address, avoiding triangular routing. Note that if a correspondent does not understand the packet state, our protocol still performs correctly, but with slightly worse performance gain.

### IV. BASIC SCHEME: INTER-DOMAIN MOBILITY

In this section, we present the details of our basic scheme: mobility support when a mobile network moves across ASes (e.g., switching to a different mobility service provider when it enters a different country or region). To correctly deliver traffic to its new location, BGP requires that the new provider announce the newly arrived IP prefix, and the previous provider withdraw its announcement of the prefix.

There are two issues if we use standard BGP. The first is that it may increase BGP routing table size significantly since there can be a large number of mobile networks. The second, which could be more serious, is that standard BGP may generate a large number of updates when mobile networks move around, possibly resulting in global routing instability. Furthermore, due to the withdrawal and re-announcement of a mobile IP prefix, some routers could temporarily lose their routes to the prefix, leading to packet losses and performance degradation to applications. Thus, for both global stability and application performance, we need to limit the propagation of BGP updates without causing incorrect forwarding decisions at routers that do not receive those updates. We design and adopt several techniques to address the preceding issues: *mobile prefixes*, *aggregation routers*, *mobility community*, *scoped BGP updates*, and *tunnel mapping*. The overall objectives of these techniques are actually pretty simple: to reduce the number of ASes which need to keep routing state for a mobile network; to allow an AS to aggregate the routing states for multiple mobile networks; and to reduce the number of routers in an AS which need to keep routing state for a mobile network.

Some of the techniques have been also proposed to improve the scalability of BGP routing (e.g., CRIO[22]).

#### A. Mobile Prefixes

The key to reducing routing table size is to introduce certain structure on the prefixes assigned to mobile networks. To this purpose, we introduce the notion of mobile prefixes. We assume that each tier-1 ISP designates a set of prefixes as its mobile prefixes. To simplify presentation, we assume that each tier-1 ISP  $i$  allocates a single large prefix  $m_i$  as its mobile prefix. We refer to  $m_i$  as the *root mobile prefix* of ISP  $i$ . Let  $M$  be the union of all root mobile prefixes of all ISPs. There should be a small number of root mobile prefixes on the Internet. Note that our discussion uses tier-1 ISP to mean more generally large ISPs providing network mobility services.

An ISP can allocate sub-prefixes from its root mobile prefix to its customers, who may further divide the sub-prefixes to their mobile network customers. When a mobile network becomes the customer of an ISP, the ISP allocates a sub-prefix (e.g., a /24 prefix) from its root mobile prefix to this mobile network as its home network prefix. Since mobile networks change attachment points, the sub-prefixes of an ISP's root mobile prefix may be scattered all across the Internet.

#### B. Aggregation Routers and Mobility Community

To reduce the number of routers keeping explicit routing state for mobile networks, a tier-1 ISP configures that only a subset of its routers advertise its root mobile prefix and know how to reach each sub-prefix of its root mobile prefix. We refer to these routers as *aggregation routers (ARs)*. Note that it is not necessary for an aggregation router to maintain routing states for all mobile networks. Aggregation routers can partition the address space so that each is responsible for only a subset of mobile networks. Such partition can approximate geographic distribution of the home location of mobile networks to minimize suboptimal routing.

To allow non-aggregation routers to install routes to mobile networks, each aggregation router of a given tier-1 ISP will create a BGP UPDATE message, for each root mobile prefix, with the next hop set to its own address. This UPDATE message is propagated to the non-aggregation routers in the ISP. Thus, each non-aggregation router should have a routing entry for each root mobile prefix default to its closest aggregation router. To reduce excessive path inflation, we require that each POP of a tier-1 ISP have an aggregation router.

The partition of routers in a tier-1 ISP into aggregation routers and non-aggregation routers requires that the ISP limit the propagation of BGP update messages only among aggregation routers, as non-aggregation routers should not be aware of such routing states. To achieve this, we design a new BGP community attribute called *mobility community*. Specifically, a BGP UPDATE message with the mobility community should not be propagated to non-aggregation routers. To guarantee correct routing, we assume that aggregation routers of a tier-1 ISP form a connected topology (not necessarily a complete topology). This can be achieved through either direct connectivity (physical) or tunneling (logical).

#### C. Scoped Interdomain BGP Updates and Tunnel Mapping

The mobility community attribute not only limits routing states to aggregation routers in tier-1 ISPs, but also controls the

propagation of BGP UPDATE and WITHDRAWAL messages, and the creation of tunnel mapping.

**BGP UPDATE:** When a mobile network with prefix  $p$  switches to a new AS, the new base station will inject a BGP announcement on the prefix  $p$  with a mobility community attribute. The mobility community attribute controls that a BGP UPDATE is not propagated to customers or non-tier-1 peers. Thus, a BGP UPDATE message may propagate up along the AS hierarchy and reach a tier-1 ISP. If it reaches a non-aggregation router first, that non-aggregation router will forward it to its closest aggregation router. This will trigger an update for  $p$  that may propagate across all aggregation routers in all tier-1 ISPs. However, the UPDATE message may arrive at a provider AS with a previous route to  $p$ . This suggests that the AS is a common provider to both the previous AS and the current AS which the mobile network attaches to. In this case, the AS suppresses the update message so that a change of base station by a mobile network does not trigger updates among any tier-1 ISPs. Note that it may not always be possible to identify the common provider as the previous route may already have been removed by a BGP WITHDRAWAL message.

**BGP WITHDRAWAL:** When a mobile network leaves an AS, the designated border router will announce a BGP WITHDRAWAL message for that prefix  $p$  with the mobility community attribute. Again, this message propagates only to providers. The message will stop at the common provider which has a new route (different from the one in the WITHDRAWAL message). It is possible that some routers at the common provider have not received the new route. In this case, some WITHDRAWAL messages can go all the way up to tier-1 ISP, triggering updates among tier-1 ISPs.

**Tunnel Mapping:** The existence of mobility community in a BGP message may also trigger the creation of a *tunnel mapping*. Specifically, assume that a tier-1 ISP's border router, referred to as PE (Provider Edge), receives a BGP UPDATE message for a specific mobile prefix  $p$  from its customer border router, referred to as CE (Customer Edge). When PE propagates the BGP UPDATE to aggregation routers in the same ISP or aggregation routers in other tier-1 ISP, the BGP UPDATE message should keep CE's IP address, and each router should create a tunnel using CE's IP address as the tunnel endpoint. This allows any aggregation router in any tier-1 ISP to tunnel packets destined to  $p$  to CE. This ensures packets can reach the CE, given that non-aggregation routers have only a default route to its closest aggregation router.

In summary, with scoped BGP updates, our scheme reduces the routing table size by requiring only ASes in the propagation path from the current attachment point of a mobile network to the top tier to have regular (not tunneled) routing entries for the specific prefix of the mobile network. The tunnel information for mobile prefixes is maintained at all aggregation routers inside tier-1 ISPs to ensure that each aggregation router knows how to reach the mobile network. With this hierarchical architecture, all other non-tier-1 ISPs need not maintain detailed routes to the mobile prefixes. Instead, they can set up default routes and forward packets (destined to the mobile prefixes) to its provider, so that the packets can still be correctly delivered via aggregation routers to the destination.

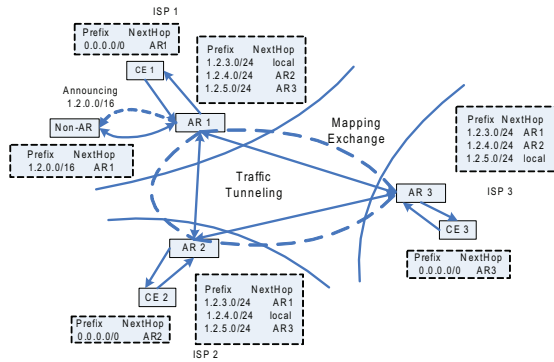


Fig. 2. An example for interdomain mobility.

#### D. An Example

Let us illustrate the whole process using an example shown in Fig. 2. Assume that 1.2.0.0/16 is the allocated mobility prefix for ISP 1. A subprefix of 1.2.0.0/16 *i.e.* 1.2.5.0/24 is currently attached to the CE 3 which belongs to another tier-1 ISP (ISP 3). The base station where the prefix is attached will inject a BGP update for the prefix with mobility community that causes the update to be propagated to an aggregation router inside ISP 3 (*i.e.* AR 3). AR 3 installs the route into its routing table and advertises this mobile prefix to other aggregation routers (AR 1 and AR 2) in other tier-1 ISPs. Assuming some customer inside ISP 1 wants to send packets to a destination within 1.2.5.0/24 and the packet is received by CE 1. Because CE 1 does not have any entry for 1.2.5.0/24 in its routing table (the mobile prefixes are only advertised among aggregation routers) except for a default route to its provider ISP 1, it forwards the packet toward the aggregation router AR 1. When AR 1 receives the packet, it looks up its routing table and finds that the next hop for the packet is AR 3. Thus it tunnels the packet to AR 3 in ISP 3 which has a local route to the mobile prefix 1.2.5.0/24. Therefore AR 3 detunnels the packet and locally forwards it to the final destination.

### V. INTRADOMAIN MOBILITY

In this section, we present the detailed design of mobility support when a mobile network moves within an MISIP. This will be a common case optimization. The components of a MISIP are shown in Fig. 1.

#### A. Infrastructure Support

To prevent iBGP routing changes due to roaming within an AS, we require that, only a designated BGP speaking router (DBR) act as the origin of the prefix of the mobile network and announces the mobile network prefix. We require that a mobile network always update this router of its care-of-address. For redundancy, multiple DBRs are selected. We leave out the detailed protocols on selecting DBR and registering care-of-address with DBR, as these protocols can be proprietary.

We assume that routers can be configured to carry out specific processing based on the flags set in the routing table entry. We require a flag denoting whether a prefix is originating within an AS (`insideAS()` in Fig. 3), a flag denoting whether a given router originated a prefix (`origin()`), and a flag denoting whether a destination prefix is a mobile network (`mobilePrefix()`). We assume that routers have different priority queues, as is the case for most routers today. We require packets to mobile networks without capability be sent on a low priority queue.

```

Process(p) – On receiving a packet p
1. if insideAS(dest(p))
2.   if presentMOS(p)
3.     if hasKey()
4.       if (validMOS(p))
5.         decrypt to get COA and stripMOS(p)
6.         tunnel p to COA
7.       else drop(p)
8.     else forward(p)
9.   if mobilePrefix(dest(p))
10.    if origin(dest(p))
11.      tunnel p to nexthop
12.    else sendInlowPriQ(p)
13.  else forward(p)
14. else forward(p)
    
```

Fig. 3. Packet processing procedure in routers.

#### B. Packet Mobility State

The key to our intradomain mobility support is packet mobility states. These states serve three purposes: removal of triangular routing, guarantee of location privacy, and also prevention of denial-of-service to the mobile networks.

The packet mobility states of a mobile network encode the binding of the home prefix and the visiting care-of-address of the mobile network. Thus, the state will be updated when the mobile network changes its attachment point. Specifically, when a mobile network switches to a new base station, it needs to authenticate itself before a care-of-address can be allocated to it. Furthermore, the mobile network needs to be sure that it is not attaching to a bogus base station. Specifically, after a successful authentication, the AAA server returns to the base station a token  $t$  which encodes the binding of the mobile network's home network prefix (HoP) and care-of-address(COA). For location privacy and integrity, the binding is encrypted by the current mobility router group key  $K_{mrg}$ ,

$$t = K_{mrg}(HoP, COA).$$

The mobility router group includes all BGP (iBGP and eBGP included) speaking routers and some additional internal routers for performance improvement. The mobility router group key should be managed by a scalable group key management system [23] and be refreshed periodically to avoid replay attacks. On the data path, the token will be stamped by the base station into the IP packets originated from the mobile network. A correspondent host should *bounce* the opaque token back to the mobile host, if its OS is updated with our proposed mobility support. Thus, we refer to the token as the packet mobility state (MOS). Note that the MOS field can be implemented as either an IP option or a shim layer between the IP layer and the transport layer. If IPsec is used, the IP option or shim layer will not be used in integrity checks.

For incremental deployment, the correspondent host always uses the home address of the mobile host for data packets from a correspondent host to the mobile network. The MOS field becomes effective when a correspondent host bounces the field back (*i.e.*, attached to the reply packets). We consider three cases. The pseudo-code of the complete router processing is given in Fig. 3.

The first case is that the correspondent host is outside the AS and the correspondent host bounces back the packet state. When the packet enters the AS, since the packet has



been stamped with the token (obtained from the mobile host), regardless of through which border gateway router the reply comes back to the MISP, the border gateway router can use the mobility router group key to verify whether the secure token is valid (the home address is within the decrypted network prefix). If it is valid, the destination of the packet will be changed to the care-of-address (from the decrypted token); the router will strip the MOS field and tunnel the packet to the care-of-address of the mobile network (line 4-6 in Fig. 3). When the BS receives the tunneled packet, it will de-tunnel and hand over the packet to the gateway of the mobile network. The packet will be dropped, if its mobility state is not valid.

If the packet does not have the MOS field (*e.g.* it initiated the connection first or the correspondent host is a legacy host which can not echo the MOS field back), the packet will be routed to the designated BGP speaking router (DBR) which announced the prefix. Because the mobile network always updates this router of its care-of-address, this router knows the care-of-address. When it receives a packet destined to the mobile network's IP prefix, it will tunnel the packet to its care-of-address (line 10-11 in Fig. 3).

The third case is that the correspondent host is inside the AS. The initial routers along the path then may not have the mobility router group key. In this case, the packet will be routed toward the DBR based on the routing information. If the packet encounters any router in the mobility router group, it will be tunneled to the current care-of-address before reaching DBR.

## VI. WINMO PROPERTIES

In this section, we analyze the correctness, optimality, and security/privacy properties of WINMO.

### A. Global Reachability

Our reduction of the set of ASes to receive the BGP announcement about a specific mobile prefix  $p$  is based on three observations. First, each tier-1 ISP maintains routing information on how to reach  $p$ . Second, a non-tier-1 AS has a default routing configuration. That is, instead of dropping an arrival packet when it does not have an explicit routing entry for the IP address of the packet, it forwards the packet to a provider. This is achieved by using a default prefix  $0.0.0.0/0$ , with the next-hop being the chosen provider. Third, if a non-tier-1 AS announces a prefix  $q$  such that  $p \subset q$ , then it must have routing information to reach every mobile network  $p \subset q$ .

With these observations, it is sufficient to guarantee that packets originated from any AS in the Internet can reach the AS  $A$  that the mobile network  $p$  is currently visiting, if the providers of  $A$  receive the BGP announcement, the providers of each preceding provider receive the BGP announcement, and this propagation continues until each update branch reaches a tier-1 ISP. Below we show that this is enough to guarantee the correctness of BGP routing.

*Proposition 1:* If one tier-1 AS  $i$  receives a route to a mobile prefix  $p$  from its customers, all tier-1 ASes will know how to reach  $p$ .

*Proof:* See appendix for the proof. ■

*Theorem 1:* The scoped BGP propagation and tunnel endpoint information propagation among tier-1 ASes guarantee that every AS on the Internet has a route to the mobile prefix  $p$ .

*Proof:* See appendix for the proof. ■

Given the preceding results, since the average hop count from any customer AS to tier-1 ISP is around 4 as shown empirically, the BGP updates can be efficiently limited to within only a few ASes, thus leading to low update overhead.

### B. Routing Optimality

Our inter-domain solution potentially introduces the following path segments. (1) An AS using default route has to go through a tier-1 AS. In normal BGP routing, the AS may not traverse tier-1. There could be routes through lower tier providers. (2) A non-AR receiving packets has to tunnel to the closest AR. (3) Routing from tier-1 to a mobile prefix  $p$  assigned to a mobile network is via tunneling. The tunnel endpoints then detunnel and route the packet natively. For (1), due to commonly adopted hot-potato routing, the path inflation to a tier-1 router should not be too large. That is, if the correspondent host is in the east coast, it should reach a router of tier-1 in the east coast. For (2), its closest AR is within the same POP. Therefore, this segment introduces very little path inflation. For (3), the tunneling part is using normal Internet routing, albeit the destination is a tunnel endpoint. The native part is done using hot-potato routing. Therefore, the total path inflation should be qualitatively small compared with normal BGP routes.

Our intra-domain solution does not introduce any iBGP or eBGP routing change. If the correspondent host is outside the visiting AS  $i$  of a mobile network, its packets to the mobile host will reach one of  $i$ 's border router  $R$  using hot-potato routing.  $R$  then routes optimally within  $i$ . Therefore, there is no path inflation introduced in this case.

However, if the correspondent host is within the same AS as the mobile host, when non-border gateway routers route the packet inside the AS, they may not understand the MOS field. In that case, the packet will be routed towards the DBR. When it encounters an iBGP speaking router in the same POP (as routing among POP are done using iBGP, this must happen) it will get tunneled to the COA. Therefore, the triangular routing is very limited. Note that only the first packet from correspondent host to mobile host may get routed to DBR.

### C. Security and Privacy

A major challenge in removing triangular routing is to avoid security attacks. In order to remove triangular routing, mobile IPv4 depends on a security authentication between the correspondent host and the mobile host in the network layer. This is challenging to bootstrap without a PKI infrastructure. Mobile IPv6 relies on the out-of-band *return routability test*. Below we show that our scheme has strong security and privacy properties. First, we list our security assumptions. We assume that border gateway routers and AAA servers are secure. They share the group key (which we assume to be secure) used for generating the packet states. To increase the resistance to security analysis, we require that the temporary group key be changed periodically. This duration depends on the strength of the key. We do not assume that the base stations are secure (*i.e.*, they are more likely to be compromised).

*Defense against connection hijacking:* If the token is secure, then no one can hijack an ongoing connection between a correspondent host and mobile host because the traffic between

a correspondent host and mobile host cannot be redirected. This is because the forged token will not pass verification at border gateway router and will be dropped. Replaying the token by a malicious attacker will induce traffic from the correspondent host to the mobile host. The traffic will not reach the attacker. We note that our scheme is in contrast to the mobile IP solution, in which case if an attacker is on the path between correspondent host and the home agent of the mobile host, the attacker can hijack the connection.

*Resilience to DDoS attack:* For a packet destined to mobile networks, if it does not carry the packet state, it will be demoted to a low priority queue. Only attackers on the path between a legitimate correspondent host and the mobile network can spoof the packet state. Therefore, the attackers' ability to DDoS the mobile network is limited. Note that this means that the first packet from a correspondent host will be sent in the low priority queue. Note also that, if a correspondent host is inside an AS, its packets with a packet state will be not be verified until it either reaches the DBR or a BGP speaking router. Forging a packet state for an attacker inside the AS doesn't help. This is because it will be dropped as soon as it reaches a router with the group key before reaching the mobile network. Our DDoS solution is specific to packets to mobile networks. We remark that it can be combined with other schemes such as TVA so that a complete solution can defend against DDoS attacks to any host.

*Preservation of location privacy:* Since the correspondent host only communicates with the home address of a mobile host when it roams across an AS (the correspondent host does not understand the secure token), the true location of the mobile host is hidden from the correspondent host. Thus, location privacy is preserved in our solution.

## VII. IMPLEMENTATION ISSUES

We discuss issues related to implementing our proposed changes in router hardware and software as well as required protocol modifications. To support interdomain mobility, ASes involved in providing the mobility service need to recognize the mobility attributes encoded using the BGP community attributes. Furthermore, routers need to recognize mobile prefixes to prevent them from being propagated globally. Such information can be simply configured similar to bogon prefix list, as mobile prefixes will be well-known. Thus, to support interdomain mobility, our scheme requires that only router configurations or software be slightly modified without any changes on router hardware or the protocols. We believe that this approach is incrementally deployable and imposes minimal overhead. The major requirement for the intradomain optimization is that routers need to efficiently verify the correctness of the token or the packet mobility state. Commercial hardware solutions for 10 Gigabit encryption and decryption needed for IPsec already exist, e.g. Ciphertics [24]. In addition, the requirement of providing different priority queues preventing denial of service attacks is uniformly supported by common commercial routers today.

## VIII. EVALUATIONS

In this section, We evaluate the performance improvement and overhead using realistic AS and ISP topologies in both inter-domain and intra-domain settings.

#AS	#edges	#P/C	#Peering	#Sibling
23408	56002	44482	11085	435

TABLE I  
STATISTICS OF THE AS RELATIONSHIP DATASET.

### A. Evaluation Methodology

We simulate the mobility and routing changes using real Internet topology data. For simplicity, we treat each AS as one node in our AS topology graph, with each AS originating a single prefix. We assume that each AS selects and exports routes using the standard policy [25] based on AS business relationships (i.e., customer-provider relationship, peer-to-peer relationship and sibling relationship). In particular, the AS route selection policy is the conventional one implementable using local preference values: customer routes have the highest priority while provider routes have the lowest priority. We apply the algorithm in [26] to infer AS relationships from BGP tables obtained from RouteViews. The statistics of the AS topology is summarized in Table I.

We evaluate our intradomain approach using the POP-level topologies of five large ISPs. These topologies are constructed from the Rocketfuel data [27]. We annotate each link in the topologies with an approximate delay value.

We use the SSFNet [28] as our main simulation tool. When computing the path inflation ratio where the optimal path is calculated based on algorithm in [29], we write a stand-alone program for better efficiency.

### B. Effectiveness of Inter-domain Mobility Support

The main goal of our inter-domain solution is to obtain low-stretch or efficient routes to mobile network when compared with native BGP routes, while bounding the BGP overhead due to updates. We demonstrate the effectiveness of our solution by comparing it with two major existing solutions for network mobility: NEMO and Boeing's Connexion. NEMO is a direct extension of mobile IP to support network mobility, which causes no routing updates but suffers from triangular routing and significant performance degradation. On the other hand, Boeing's Connexion completely removes triangular routing among ASes by propagating the updates for the mobile network via BGP. However, it introduces updates globally for each mobile prefix and hence could lead to serious BGP churns and instabilities.

First we compare the three solutions by their average path inflation ratio and average path length (in terms of number of AS hops). The path inflation ratio of a given path is defined as the path length given by an algorithm over the optimal path length. We conduct 500 rounds of simulations. In each round, we randomly pick one AS as the attachment point AS of a mobile network. For each mobility solution, we compute the average path length between the attachment AS and all other ASes (which is equivalent to all other ASes serving as corresponding networks). For NEMO, since the home agent is also important for calculating AS path length, for each (correspondent AS, attachment AS) pair, we randomly select 500 home agent locations and average the path length. Since Connexion propagates routing information across the entire Internet, it achieves the best route possible under BGP policy constraints. Hence we normalize the path inflation of the other two solutions against Connexion.

The results on path inflation are summarized in Table II. From Table II, the normalized average path length ratio of

	Connexion	NEMO	WINMO
Avg Path length	3.951	7.475	4.403
Avg Path Length Ratio	1	1.89	1.11

TABLE II  
AVERAGE AS PATH LENGTH.

WINMO is only 11% higher than Connexion, while that of NEMO is 89% higher. An explanation of the result is that in most situations, traffic from a stub AS has to use its provider's route to reach the destination. In WINMO, all traffic will be delivered through tier-1 ASes and the general hop count between an AS and the top tiers is only around 3 or 4; hence the increase in path length in WINMO is very small compared with the BGP optimal Connexion.




Fig. 4. CDF of path inflation ratio: WINMO incurs minimal path inflation.

We also plot the cumulative distribution of the normalized path inflation in Fig. 4. A particularly interesting observation from the figure is that sometimes WINMO selects even a shorter path in terms of AS hop count compared with Connexion, whose path is selected by BGP. This is because the route selection of BGP takes into account various policies and preference (e.g., customer route is preferred over provider route). This sometimes results in suboptimal paths that traverse through an AS's customers. However, in WINMO, the provider route is selected with a shorter AS hop count than the customer route chosen in the Connexion case. We see that, for WINMO more than 80% of the paths has an inflation less than 30% when compared with Connexion. However, for NEMO, over 80% of the paths has an inflation of 90%.

There is a trade-off between routing optimality measured by path length and the amount of updates introduced. Thus, one concern is that the efficient paths chosen by our solution could have been achieved at the cost of injecting a large number of BGP updates to the inter-domain routing system. Our next results show that our solution obtains near-optimal routing while incurring little BGP update overhead.

We obtain BGP update count by simulating BGP dynamics using the SSFNet BGP simulator. To evaluate the trends in the number of BGP updates and the BGP convergence time when Internet grows, we extract sub-topologies from the complete AS topology. We vary the number of ASes in a sub-topology from 100 to 1500. For extracted sub-topology, we preserve the business relationships among the ASes.

Fig. 5 and 6 depict disruption time and the number of updates for both WINMO and Connexion. The disruption time is defined as the time duration when a router doesn't have a route to reach the mobile prefix. The maximum disruption of Connexion is not shown in the figure because of its two orders of magnitude higher value. Even the worst case of WINMO is still much better than the average cases for Connexion. Of particular interest are the results on the number of BGP updates. We observe that in Connexion, the number of BGP




Fig. 5. Disruption time comparing Fig. 6. Number of BGP updates for WINMO with Connexion.

updates grows exponentially as the number of ASes increases. As a contrast, WINMO successfully controls the propagation of BGP updates, and thus generates much fewer updates. Note that we show results only for sub-topologies with size up to 1500 ASes. Connexion generates from 32 to as much as 124 times more BGP updates than WINMO.

### C. Effectiveness of Intra-domain Mobility Support

In the preceding section, our evaluation focusing on AS-level Internet topology demonstrates WINMO's effectiveness when mobile networks move across ASes. Evaluation using AS-level topology hides the details of intra-domain mobility. In this section, we evaluate the effectiveness of WINMO when mobile networks move within an AS. To achieve this, we use network topologies with finer granularity. We construct five POP-level topologies (each belongs to one of five large ISPs) with long distance link delay from Rocketfuel project [27]. We assume that the intra-domain protocol is OSPF and the shortest path is used to route packets.

We assume a mobile host in a mobile network roams within an AS. The correspondent host can be anywhere in the Internet. Assuming the entry point (EP) to the AS where the mobile host currently visits is the same for NEMO and WINMO, while the rest of the Internet topology is irrelevant. For WINMO, packets will be delivered directly to the current attached base station (BS). For NEMO, packets first go to the home agent of the mobile host, then get tunneled to the current BS. We focus on this inefficiency using performance metrics of path and delay inflation ratio. For each correspondent host and mobile host pair, it is the path length (delay) of NEMO over that of WINMO.

ISP	Telstra	Tiscali	Sprintlink	Ebone	Exodus
#POPs	61	50	43	25	23
Avg path infl	2.410	2.217	2.433	2.509	2.350
Avg delay infl	3.453	3.871	6.148	4.042	6.581

TABLE III  
AVERAGE PATH/DELAY INFLATION OF NEMO NORMALIZED BY WINMO.

We first show the overall inflation ratios in Table III using average ratios: For each pair of correspondent host (CH) and mobile host (MH), we calculate the average inflation ratio by averaging over all possible entry points; we then compute the average over all possible (CH, MH) pairs. We observe that across all 5 topologies, the path inflation of NEMO is consistently at least 2.2 times of our solution. The delay inflation of NEMO compared with our solution is even higher: at least 3.4 in all 5 cases.

We also compute the cumulative distributions of the path and delay inflation ratios. Fig. 7 and 8 show the results.






Fig. 7. CDF of path inflation of Fig. 8. CDF of delay inflation of NEMO normalized by WINMO. NEMO normalized by WINMO.

In particular, for delay inflation, we observe that there exist cases where the delay of NEMO is 10 times that of our solution. These cases can lead to substantially reduced user performance. The poor performance of NEMO with respect to WINMO in terms of path length and delay demonstrates the effectiveness of our solution in handling intra-domain mobility.

### IX. CONCLUSIONS AND FUTURE WORK

In this paper, we have proposed WINMO, a simple, systematic, novel solution for wide-area IP network mobility. Through scoped BGP updates, route aggregation and tunneling as well as mobility packet state, we have achieved low stretch global Internet routing for mobile networks roaming across wide areas with minimal inter-domain routing overhead. Our extensive evaluation shows that, the average path length of WINMO is only 11% more when compared with Connexion which uses BGP as it is; the BGP update overhead of WINMO is orders of magnitude smaller than Connexion.

Mobility support is one of the major challenges facing the Internet, and there are many avenues for further study. In particular, as a typical engineering design, our design has made a number of tradeoffs. Specific deployments may need to make different tradeoffs according to user and network requirements. We believe that our design is flexible and adaptable to many settings, and we will evaluate our design in more settings.

### REFERENCES

- [1] C. Perkins, "Mobile IP," *IEEE Communications Magazine*, vol. 35, pp. 84–99, May 1997.
- [2] A. Dul, *Global IP Network Mobility using Border Gateway Protocol*, Mar. 2006. [Online]. Available: <http://www.mobilenetworks.org/nemo/ietf62/nemo-ietf62-global-mobility-b%gp.pdf>
- [3] T. Underwood, "Tracking plane flight on Internet," Available from [http://www.renesys.com/blog/2006/04/tracking\\_plane\\_flight\\_on\\_inter.shtm%1](http://www.renesys.com/blog/2006/04/tracking_plane_flight_on_inter.shtm%1).
- [4] A. Campbell, J. Gomez, S. Kim, and C. Wan, "Comparison of IP micro-mobility protocols," *IEEE Wireless Comm*, pp. 72–82, Feb. 2002.
- [5] R. Moskowitz and P. Nikander, *Host Identity Protocol Architecture, RFC 4423*, Aug. 2005.
- [6] F. Teraoka, K. Uehara, H. Sunahara, and J. Murai, "VIP: A protocol providing host mobility," *Communications of the ACM*, vol. 37, no. 8, Aug. 1994.
- [7] F. Teraoka, Y. Yokote, and M. Tokoro, "A network architecture providing host migration transparency," in *Proceedings of ACM SIGCOMM '91*, Zurich, Switzerland, Sep. 1991.
- [8] A. C. Snoeren and H. Balakrishnan, "An end-to-end approach to host mobility," in *Proceedings of the Sixth International Conference on Mobile Computing and Networking (Mobicom)*, Boston, MA, Aug. 2000.
- [9] B. Y. Zhao, A. D. Joseph, and J. D. Kubiatiowicz, "Supporting rapid mobility via locality in an overlay network," University of California at Berkeley, Berkeley, CA, USA, Tech. Rep., 2002.
- [10] S. Zhuang, K. Lai, I. Stoica, R. Katz, and S. Shenker, "Host mobility using an internet indirection infrastructure," *Wireless Networks*, vol. 11, no. 6, pp. 741–756, 2005.
- [11] A. Dutta, T. Zhang, S. Madhani, K. Taniuchi, K. Fujimoto, Y. Katsube, Y. Ohba, and H. Schulzrinne, "Secure universal mobility for wireless internet," *SIGMOBILE Mob. Comput. Commun. Rev.*, vol. 9, no. 3, 2005.

- [12] E. Sakhaee and A. Jamalipour, "The global in-flight Internet," *IEEE Journal on Selected Areas in Communications*, vol. 24, no. 9, Sep. 2006.
- [13] E. Perera, V. Sivaraman, and A. Seneviratne, "Survey on network mobility support," *SIGMOBILE Mob. Comput. Commun. Rev.*, vol. 8, no. 2, pp. 7–19, 2004.
- [14] IETF, "IETF NEMO working group," <http://www.ietf.org/html.charters/nemo-charter.html>.
- [15] V. Devarapalli, R. Wakikawa, A. Petrescu, and P. Thubert, *Network mobility (NEMO) basic support protocol, RFC 3963*, Jan. 2005.
- [16] H. Petander, E. Perera, and K. L. A. Seneviratne, "Measuring and improving the performance of network mobility management in IPv6 networks," *IEEE Journal on Selected Areas in Communications*, vol. 24, no. 9, Sep. 2006.
- [17] C. Huang, C. Lee, and J. Zheng, "A novel SIP-based route optimization for network mobility," *IEEE Journal on Selected Areas in Communications*, vol. 24, no. 9, Sep. 2006.
- [18] M. Calderon, C. Bernardos, M. Bagnulo, I. Soto, and A. De La Oliva, "Design and experimental evaluation of a route optimization solution for NEMO," *IEEE Journal on Selected Areas in Communications*, vol. 24, no. 9, Sep. 2006.
- [19] H. Cho, T. Kwon, and Y. Choi, "Route optimization using tree information for nested mobile networks," *IEEE Journal on Selected Areas in Communications*, vol. 24, no. 9, Sep. 2006.
- [20] A. Baig, L. Libman, and M. Hassan, "Performance enhancement of on-board communication networks using outage prediction," *IEEE Journal on Selected Areas in Communications*, vol. 24, no. 9, Sep. 2006.
- [21] Y. Han, J. Choi, and S. Hwang, "Reactive handover optimization in IPv6-based mobile networks," *IEEE Journal on Selected Areas in Communications*, vol. 24, no. 9, Sep. 2006.
- [22] X. Zhang, P. Francis, J. Wang, and K. Yoshida, "Scaling ip routing with the core router-integrated overlay," in *ICNP '06*, 2006.
- [23] IRTF, "The secure multicast research group (SMuG)," <http://www.ipmulticast.com/community/smuG/>.
- [24] "The First 10Gig IPsec Encryption Solution," <http://www.cipheroptics.com/products/sg10g.html>.
- [25] L. Gao and J. Rexford, "Stable Internet routing without global coordination," *IEEE/ACM Trans. on Networking*, vol. 9, no. 6, Dec. 2001.
- [26] L. Gao, "On inferring autonomous system relationships in the Internet," *IEEE/ACM Trans. on Networking*, vol. 9, no. 6, Dec. 2001.
- [27] N. Spring, R. Mahajan, and D. Wetherall, "Rocketfuel: An ISP topology mapping engine," Available from <http://www.cs.washington.edu/research/networking/rocketfuel/>.
- [28] "SSFNNet BGP: SSF implementation of BGP-4," <http://www.ssfnet.org/hgp/doc/>.
- [29] X. Dimitropoulos, D. Krioukov, M. Fomenkov, B. Huffaker, Y. Hyun, K. Claffy, and G. Riley, "AS Relationships: Inference and Validation," *ACM Computer Communication Review*, vol. 37, no. 1, 2007.

### APPENDIX

#### Proof of Proposition 1

*Proof:* Let's denote the border router in AS  $i$  that received the UPDATE message  $s$  with  $PE$ . Let's denote the customer border router that sent the message  $s$  with  $CE$ .  $PE$  (whether AR or non-AR) will propagate  $s$  to its closest AR. As all ARs are connected, eventually, all ARs will receive  $s$ . There are two cases. If a router  $R$  is an AR. Since  $s$  preserves the tunnel endpoint information of  $CE$ , the router will be able to tunnel packets destined to  $p$  to  $CE$  which in turn detunnels, and knows how to reach  $p$ . If a router  $R$  is a non-AR, since  $R$  has a default route to its closest AR, packets to  $p$  will be tunneled to that AR. This becomes case 1. Note that there can be multiple UPDATE messages from a number of routers like  $CE$ . There can also be stale entries. However, stale entries will eventually be withdrawn. Afterwards, all tunnel endpoints should be valid. Each AR can choose the best tunnel endpoints. ■

#### Proof of Theorem 1

*Proof:* We assume routing converged in our proof. There are two cases. One is that an AS has a route to  $q$  where  $p \subset q$ . The other is that an AS has only default router to providers. For the first case, since an AS announces  $q$ , it must know how to reach each sub-prefix of  $q$ . For the second case, packets to  $p$  will be sent by default to providers, if packets go through a non-tier-1 AS with a routing entry  $q$  such that  $p \subset q$ , this goes back to case 1; if it reaches all the way to tier-1 AS. The router should either have a routing entry to the root prefix  $r$  of  $p$  or a specific prefix  $q$  such that  $p \subset q \subset r$  (by Proposition 1 and properties of scoped BGP updates). For the former, packets will be tunneled to the closest AR which will have a tunnel endpoint for  $p$ . For the latter,  $q$ 's tunnel endpoint (aggregated entry) will reach  $p$ . ■