

Registration Frameworks For IP Mobility Agents Hierarchies

Ayman Abdel-Hamid

Hussein Abdel-Wahab

Department of Computer Science, Old Dominion University

Norfolk, VA 23529, USA

{hamid, wahab}@cs.odu.edu

Abstract

Mobile IP is a network layer solution to the wide-area mobility problem in the Internet. Mobility agents' hierarchies and regional registrations have been introduced in the foreign domain as a local-area mobility support solution. In this paper, we critique the current regional registration proposal identifying some drawbacks with its registration mechanisms signaling design. In addition, we introduce novel registration frameworks for regional and home registrations associated with intra-hierarchy handoffs. We attempt to emphasize the local handoff aspect and benefit from the presence of a mobility support overlay network in the form of an agents' hierarchy. Performance evaluation results through network simulation demonstrate the effectiveness of the proposed techniques in reducing UDP packet loss, and maintaining better TCP throughput versus a base Mobile IP implementation, in the case of a distant home agent.

1 Introduction

An IP address reflects a host's point of attachment to the network. A mobile host continuously changing its network point of attachment creates a routing problem: *With no special handling to deal with host mobility, packets addressed to a mobile host will be routed to the mobile host's home network, not to its current location.* This problem occurs because an IP address serves a dual purpose: a *routing directive* in the network layer and an *end point identifier* in the transport layer [4].

Mobile IP presents a network layer solution to the host mobility problem in the Internet for both wired and wireless networks. Mobile IPv4 [11] uses a two level addressing architecture, and deploys *Mobility Agents* in the home network, and the visited network. The MH is associated with two IP addresses: its permanent home IP address which serves as an end point identifier, and a transient care-of IP address which reflects its current point of attachment, and serves as a routing directive at the network layer. The care-of address can be the address of a *Foreign Agent* (FA) in the visited network, or can be a co-located care-of address, which the mobile host

acquires on the visited network. The FA is a router in the foreign network that acts as a mobility agent. Whenever a mobile host is away from home, it registers its current care-of address with its *Home Agent* (HA). The HA is a router that acts as a mobility agent in the home network, and intercepts any datagrams destined to the mobile host's home address, and tunnels them to the registered care-of address. A host in the Internet communicating with the MH is termed a *Correspondent Host* (CH).

Mobile IP can handle wide-area mobility, and local-area mobility. Although, it is more suited to handle the former since a mobile host is required to register with its, possibly distant, HA whenever it changes its point of attachment. This results in large registration signaling overhead, and large handoff latencies in the local-area mobility case. One solution to handle local-area mobility in Mobile IPv4 deploys FA hierarchies within the foreign domain, and introduces the concept of regional registrations. Regional registrations are localized registrations processed by regional foreign agents in the hierarchy, hence shielding the HA from processing such local-movement and reducing packet loss during local handoffs [9].

In this paper, we present novel registration processing frameworks for intra-hierarchy handoffs in Mobile IPv4 foreign agent hierarchies in the foreign domain. We critique the current regional registration proposal [9] (MIP_RR) identifying several drawbacks and race conditions within its registration techniques. Consequently, we introduce an enhanced regional registration framework that avoids the identified drawbacks. In addition, we suggest two novel registration processing frameworks for home registrations involving local handoffs in which we identify the dual nature of such registrations, and attempt to emphasize the local handoff aspect. One technique, maintains tunneling of data packets to the MH through an old path until a home registration reply is received to set up the new path. In contrast, the other technique adopts a more proactive approach in switching immediately to the new path resulting in a reduction of the handoff latency. The proposed mechanisms are evaluated qualitatively and analytically, and their performance is investigated through network simulations using our extension of Columbia University's IP Micro-mobility software (CIMS) [6], an *ns-2* source code extension [10]. Our registration frameworks achieve a sizable reduction in UDP packet loss, and maintain better TCP throughput in the case of a distant home agent, versus a base Mobile IP implementation.

The rest of this paper is organized as follows. Section 2 presents an overview of related work. Section 3 critiques the current regional registration proposal identifying drawbacks and deficiencies within its registration processing techniques. Section 4 presents an enhanced regional registration framework. Section 5 introduces two novel techniques for processing home registrations involving local handoffs. Section 6 presents performance evaluation results of the proposed mechanisms through network simulations. Finally the paper is concluded in section 7, along with future work.

2 Related Work

A number of proposals, within Mobile IP framework, exist to handle local-area mobility, without incurring any large handoff latencies [7], [9]. Other researchers have optimized their local-area mobility solutions towards the wireless network environment, e.g., HAWAII [15], and Cellular IP [17]. In this section, we focus on registration techniques within the Mobile IP framework.

The *Regional Registration* approach introduced FA Hierarchies in the foreign domain [9]. An FA hierarchy (Figure 1) is rooted by a *Gateway Foreign Agent* (GFA), which has a publicly routable IP address. The MH registers, with its HA, the GFA address as its care-of address. This care-of address will not change when the MH changes FA under the same GFA. After registering the GFA IP address with the HA as care-of address, the MH is allowed to perform regional registration within the FA hierarchy as long as its registration with the HA did not expire. Hence, part of the HA functionality is delegated to the GFA, and any of the FAs beneath the GFA in the hierarchy become *Regional FAs* (RFA), i.e. the target of a regional registration from the MH. Figure 1 illustrates an example of the MH's home registration, and regional registration.

The crossover FA, the first common FA between the new and old path, has the task of issuing a registration reply to the MH's regional registration. The ability of a FA to correctly identify itself as the crossover FA for a regional registration request is crucial for correct registration processing. Such ability might be hindered because a MH is not required to inform a FA that it is no longer registered with it, relying on an eventual expiration of registration lifetime. Such approach reduces protocol messages overhead, but creates a *hierarchy tunneling consistency problem* for FA hierarchies: *a RFA not informed that the MH is no longer a current visitor might erroneously decide that it is the crossover FA and generate a regional registration reply in response to a regional registration request, although such request should be forwarded to upper level RFAs*. Hence, a mechanism, which we term a *tunneling consistency mechanism*, is required by which old regional foreign agents are informed that a MH is no longer a current visitor. The regional registration framework requires a smooth handoff mechanism [13] to be performed by the MH and a new FA in order to inform an old FA that the MH is no longer a current visitor. The old FA relays the smooth handoff binding update (BU) message, received from the new FA, upwards in the hierarchy (to its father FA) specifying itself as the care-of address of the MH. The father FA performs the following steps in response to receiving the BU message: delete its MH's visitor entry, create a binding cache entry for the MH with care-of address the child FA that sent the BU message, relay the BU message upwards in the hierarchy, and send back a binding acknowledge message to its child FA. Such process at each intermediate RFA repeats until the BU message reaches the crossover FA, which at this point generates a binding acknowledge message to the MH and sends it down the old path to the MH. The crossover FA deduces it is the crossover FA, and hence generates the binding acknowledge message to the MH, because it has received an earlier registration request from the MH's new path beneath it in the hierarchy. The same process is used for regional registrations and home registrations associated with local handoffs.

A number of FA hierarchies might be deployed in the same foreign domain, although existing prototype implementations, and simulations have considered deploying only one FA Hierarchy in the foreign domain [8], [14]. In [2], we introduced a cooperation-based registration processing framework for inter-hierarchy handoffs in order to further reduce packet loss for such handoffs.

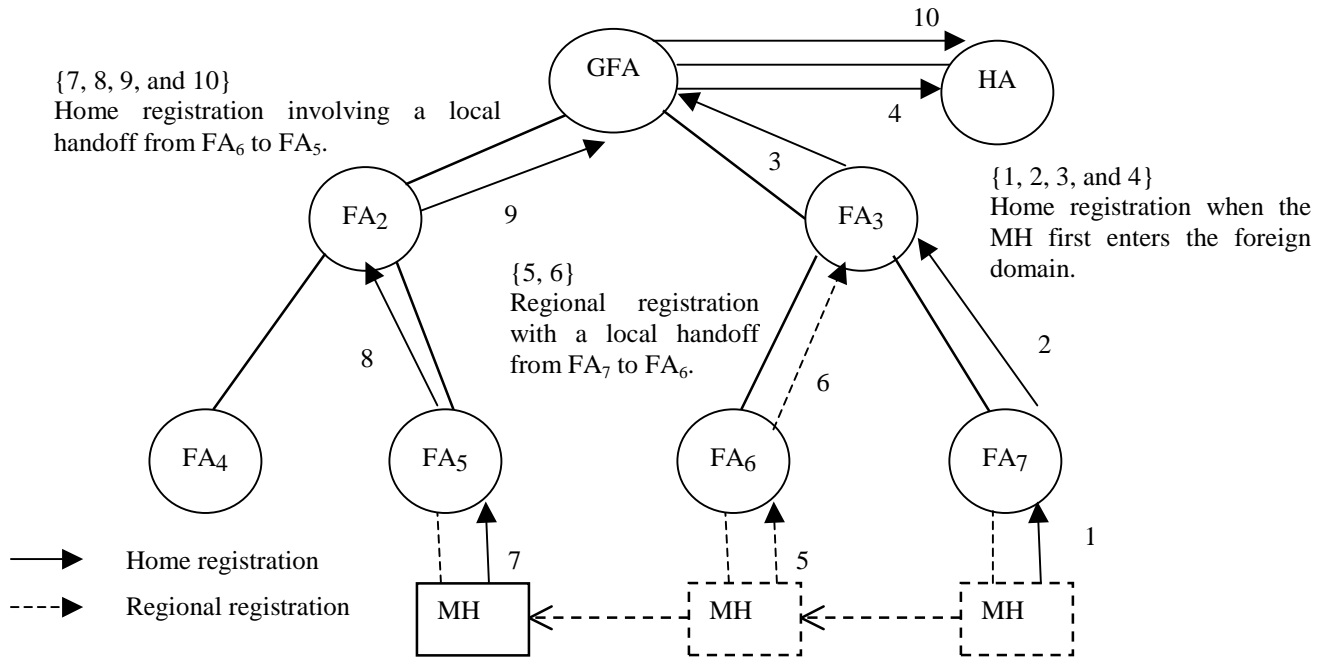


Figure 1: FA Hierarchy within a visited foreign domain.

The *Anchor FA* approach [7] introduces two registration methods to reduce handoff latencies within the visited domain: *local registration*, and *global indirect registration*. Either method requires the MH to perform a *global registration* (Mobile IP home registration) with its HA upon entering a visited *Zone*. In the *local registration* method, the current FA and the MH establish a shared security association, and the current FA acts as an *Anchor FA* for this MH, authenticating the MH while it moves within the same zone. When the MH changes FA within the same zone, the new FA performs local registration with the Anchor FA. The *global indirect registration* method is used when no security association could be established between the current FA and the MH, requiring the HA to always authenticate the MH registration. This approach has the disadvantage of requiring two shared security associations, one in each direction, between any two FAs within a zone. However, since any FA can become an Anchor FA, management of mobile hosts routing entries is distributed.

3 Critique of The Regional Registration Framework

Evaluating MIP_RR’s tunneling consistency mechanism, we point out the following advantages: it requires the smooth handoff mechanism [13], hence reduces potential packet loss until handoff completion; it introduces the binding acknowledgment of binding update (BU) messages between RFAs to insure message delivery; and the tunneling consistency mechanism is symmetrical for home and regional registrations. However, the following disadvantages can be observed: a MH that is not smooth handoff enabled would be denied service from within such FA hierarchy; and a potential race condition exists within the proposed message signaling, which might lead to inconsistent tunneling state within the hierarchy. If the BU message propagated through the old path reaches the crossover FA before the registration request propagated through the new path, a crossover FA can not deduce it needs to generate the MH’s binding acknowledge message and hence forwards the BU message to its father FA towards the GFA. Such scenario has the following consequences.

- Upper RFAs, higher than the crossover FA, possibly up to the GFA will replace the MH’s visitor entries with corresponding binding caches. The MH’s will be considered as not currently a visitor for these RFAs.
- The MH’s remaining registration lifetime, which the initial lifetime amount was initially granted by the HA, will be replaced with the specified lifetime in the BU message. Tunneling lifetime inconsistencies will exist in the hierarchy.

In the case of a home registration associated with a local handoff (HR-LH), such BU lifetime might expire before a registration reply is received from the HA. In the case of a regional registration, such BU lifetime is not consistent with the lifetime granted by the crossover FA in the generated regional registration reply, thus creating remaining lifetime inconsistencies within the RFAs. MIP_RR’s authors suggest that upper level RFAs should ignore the BU forwarded by the crossover FA since it does not supply any new care-of address. We argue that this condition is true for every binding update propagated through the old path and cannot be used as a special condition to ignore the BU for RFAs above the crossover FA.

The signaling design flaw highlighted here stems from allowing 2 messages, the BU through the old path, and the registration request through the new path, to simultaneously flow in 2 separate paths towards the crossover FA that identifies its “crossover” status when receiving a registration request through the existence of a visitor entry for this MH. Furthermore, such separate paths may not be symmetric in terms of available bandwidth, link delay, or current congestion status. In order to better quantify the conditions favoring the occurrence of the identified race condition, we introduce a set of relevant delay measures summarized in TABLE 1. The identified race condition with possibility of subsequent hierarchy tunneling inconsistencies will occur if the inequality in (1) holds.

$$D_{new-old} + D_{old-crossover} \leq D_{new-crossover} \tag{1}$$

TABLE 1
Summary Of Delay Measures

Delay Measure	Description
$D_{new-old}$	The delay required for a BU message to be generated and transmitted by the new FA to reach the old FA, when the new FA receives the MH's request.
$D_{old-crossover}$	The delay for a BU message to be generated and transmitted by the old FA and relayed by each intermediate RFA in the old path to reach the crossover FA, when the old FA receives the BU from the new FA.
$D_{new-crossover}$	The delay for a MH's registration request to be relayed by the new FA and propagated by each intermediate RFA in the new path to reach the crossover FA, when the new FA receives the MH's request.

4 A Regional Registration Processing Framework

In this section, we present a regional registration processing framework and associated tunneling consistency mechanisms. We opt for a signaling design methodology that does not require the smooth handoff mechanism to maintain tunneling consistency, leaving such functionality to be optionally used by the MH to further reduce potential packet loss. The crossover FA triggers the tunneling consistency mechanism upon receiving a regional registration request. Thus, race conditions, stemming from the crossover FA's inability to properly identify its status as a "crossover FA" for a certain registration request, are prevented.

4.1 Operational Overview

When the MH sends a regional registration request, it is propagated upwards in the hierarchy until it reaches the crossover FA. The crossover FA generates a regional registration reply switching the tunneling path for the MH from the old path to the new path. The regional registration reply is propagated down the new path until it reaches the MH. Any future data packets received at the crossover FA for the MH are tunneled through the new path, alleviating the need for the old path.

We suggest using a *Deregistration mechanism*, as a tunneling consistency mechanism, triggered by the crossover FA, by which a *binding update message with lifetime equal to 0* is propagated through the MH's old path originating from the crossover FA. A similar approach, albeit relying on home registrations only, was previously proposed and implemented in the Dynamics Hut Hierarchical mobile IP implementation [8]. Furthermore, we require binding update delivery acknowledgement by the receiving FA in response to the deregistration message. Consequently, each RFA, beneath the crossover FA in the old path, receiving the binding update message from its parent FA performs the following steps: note the current tunnel endpoint for the MH, delete the MH's visitor entry, generate a deregistration message to the noted tunnel endpoint (one of its children FA), and generate a binding acknowledgment message back to the sender FA (its parent FA). This process repeats at each intermediate RFA until the BU message reaches the leaf FA that was previously serving the MH (old FA). If a RFA does not receive a binding acknowledgment message from the tunnel endpoint (one of its children FA) after a specific time interval, it is responsible for resending the BU message, until an

acknowledgment is received. The proposed tunneling consistency mechanism ensures that the old path entries for the MH are cleared in a timely fashion. Figure 2 illustrates the proposed signaling message flow for regional registrations along with the associated tunneling consistency mechanism invoked by the crossover FA. If the MH is requesting simultaneous binding in its regional request, such deregistration mechanism is not initiated.

Any binding update or acknowledgement message exchange between foreign agents is authenticated by a route optimization authentication extension [13] based on the pre-established security associations between each parent and child FA. Replay protection is provided for binding updates by usage of the identification field within the binding update and acknowledgment messages according to the replay protection style between each 2 pairs of foreign agents. The MH authenticates its regional registration request using its registration key shared with the foreign agent hierarchy by appending a MH-GFA authentication extension (a subtype of the generalized authentication extension [12]). Replay protection for regional registration requests is provided through the usage of the identification field within the request and reply messages [11]. In [2], we suggested a mechanism to disseminate the new identification value generated by the crossover FA, in the case of timestamp or nonce replay protection, to upper levels of the hierarchy to ensure future successful registration processing by such upper levels.

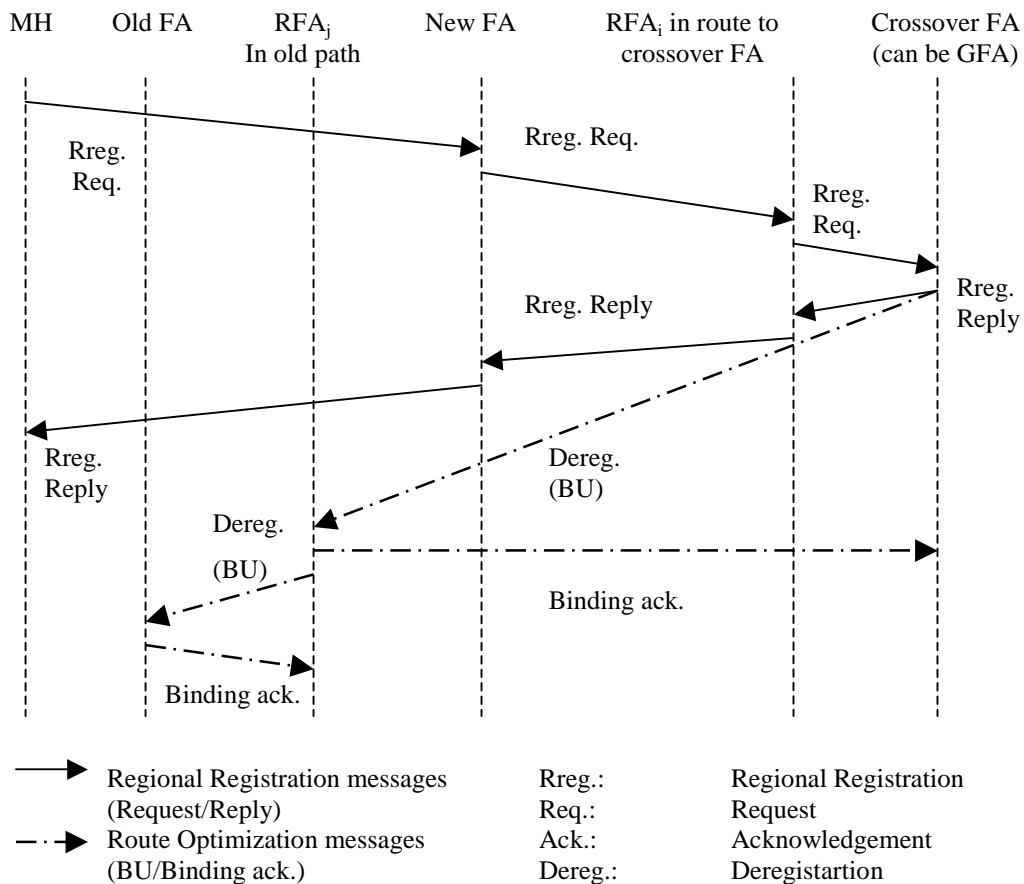


Figure 2: Proposed signaling message flow for regional registration.

4.2 Analysis and Comparison

We analyze the proposed regional registration framework and associated tunneling consistency mechanism using the delay measures introduced earlier in TABLE 1. In addition, we introduce a new delay measure $D_{crossover-new}$ to be the delay to generate and transmit the regional registration reply by the crossover FA, to be relayed by all intermediate RFAs to reach the new FA, starting when the crossover FA receives the MH's registration request. We define $D_{crossover-old}$ to be the delay to generate and transmit a BU down the old path by the crossover FA, to be relayed by all intermediate RFAs to reach the old FA, starting when the crossover FA receives the MH's registration request. In TABLE 2, we compare the proposed approach versus MIP_RR, in terms of the following aspects.

1. The delay before the tunneling consistency mechanism is initiated ($D_{InitConsistent}$) starting when the new FA receives the MH's request.
2. The delay before the tunneling consistency mechanism runs to completion ($D_{Consistent}$).
3. The delay before the old FA can forward any packets (buffered or received afterwards) to the new FA ($D_{Forward}$), if any, starting when the new FA receives the MH's request.

For our proposed approach, we are assuming the availability of a mechanism that propagates the new FA IP address information to the old FA (see section 5.1), while the MH is not using the smooth handoff mechanism. If the MH is using the smooth handoff mechanism, only $D_{Forward}$ is affected, according to whether the BU from the new FA, or the tunneling consistency BU from the parent FA reaches the old FA first.

TABLE 2
Quantitative Comparison Of The Proposed Approach Versus MIP_RR

Measure	Proposed	Proposed with MH using smooth handoff	MIP_RR
$D_{InitConsistent}$	$D_{new-crossover}$	$D_{new-crossover}$	$D_{new-old}$
$D_{Consistent}$	$D_{InitConsistent} + D_{crossover-old}$	$D_{InitConsistent} + D_{crossover-old}$	$D_{InitConsistent} + D_{old-crossover}$
$D_{Forward}$	$D_{new-crossover} + D_{crossover-old}$	Min $\{D_{new-old},$ $D_{new-crossover} + D_{crossover-old}\}$	$D_{new-old}$

Analyzing the formulas presented in TABLE 2, we conclude that the initiation of the tunneling consistency mechanism in our proposed framework is dependent on the number of RFAs between the new FA and the crossover FA (number of intermediate levels), and the corresponding link and queuing delays on the new path. On the other hand, for MIP_RR such measure is dependent on link, queuing and routing delays between the new and old FA, respectively. The mathematical relationship (\leq, \geq) between $D_{new-old}$ and $D_{new-crossover}$ holds the key to which approach is faster in ensuring tunneling consistency along the old path. Utilizing the smooth handoff mechanism as the basis for the tunneling consistency mechanism, allows initiating such mechanism from the earliest possible point on the RFA tree, which is the new FA. However, the usage of such approach coupled with the dependence on the time of receiving the corresponding regional registration request by the crossover FA

open the door for potential race conditions. Even in the absence of such potential race condition, we argue that our framework presents a viable approach to ensure hierarchy tunneling consistency and inform the old FA about the new FA when the MH is not enabled to use the smooth handoff mechanism, or the foreign agents are not advertising any means for personal identification such as their IP addresses, or their network access identifiers [3].

5 A Home Registration Processing Framework

The MH performs a home registration when it first enters the foreign domain, and periodically to maintain its home mobility binding. If no handoff is involved, no tunneling consistency mechanism is required since there is no crossover FA, and no old path with visitor entries to be cleared. On the other hand, if a local handoff is involved, the registration request is propagated in an un-established path until it reaches a crossover FA that forwards the request over an already established path until it reaches the GFA. In the latter case, simply applying a tunneling consistency mechanism that deletes the MH's visitor entries in the old path maintains the tunneling consistency, but degrades handoff performance since the MH can not be reached until the registration reply from the HA establishes the new path.

In this section, we present 2 novel approaches for processing of home registrations involving local handoffs within the same foreign agent hierarchy. Both approaches attempt to exploit the hierarchy structure, in order to optimize the MH handoff while waiting for the home registration reply to be received from the HA. The KOPA (Keep Old Path Alive) approach (section 5.1) follows the same line of thought as MIP_RR [9] in attempting to keep the old path "alive" and tunnel packets to the MH's new FA until the handoff completes by receiving a home registration reply. Nevertheless, such task is performed without relying on the MH's usage of the smooth handoff mechanism as a required component, and designed to prevent previously identified race conditions (section 3). The SINP (Switch Immediately to New Path) approach (section 5.2) emphasizes the local handoff aspect and switches the MH's tunneling path within the hierarchy immediately to follow the new path without waiting for the home registration reply.

5.1 KOPA: Keep Old Path Alive Approach

The KOPA approach relies on the crossover FA to initiate a mechanism by which the MH's old path is kept alive until the home registration is received, and consequently creating visitor entries in the new path. "Keeping the old path alive" implies performing the following two steps.

1. Replace visitor entries in the old path with binding cache entries, with a specified lifetime, that point to the same visitor entry's tunnel endpoint for the MH (one of the children FAs),
2. Inform the old serving FA about the MH's new FA to tunnel to it any already buffered or future data packets that arrive at the old FA.

Such procedure ensures future tunneling consistency since a visitor entry is replaced with a binding cache avoiding future erroneous decisions by RFAs when receiving a future MH's regional registration request. Propagating the new FA information to the crossover FA and then down the old path is necessary since we do not rely on the MH's usage of the smooth handoff mechanism.

The new FA, upon receiving the MH's home registration request, propagates the registration request upwards in the new path, appending his own IP address information using a *local care-of address* extension [1]. This registration request eventually reaches the crossover FA. The crossover FA generates a binding update message with an estimated lifetime down the old path to its visitor entry tunnel endpoint appending the new FA information it extracts from the received registration request. The binding update lifetime is computed as a function of the perceived home registration latency at the crossover FA and the remaining registration lifetime [1] as follows.

$$BU \text{ lifetime} = \text{Max} \{ \text{home reg. latency}, \alpha * \text{remaining reg. lifetime} \} \quad (2)$$

Where $0 < \alpha \leq 1$.

The fraction α represents the percentage of the remaining registration lifetime to be used in comparison to the home registration latency. We suggest using a value of 0.5 for α , since the initial value for the registration lifetime is set by the MH and later approved or modified by the HA, and the MH most probably will issue a home registration request long enough before the registration expiration. Hence, the remaining registration lifetime might be much larger than the home registration latency, depending on when the MH initiates a home registration request.

The crossover FA propagates the home registration request upwards towards the GFA which sends it to the HA. The crossover FA only propagates the new FA information down the old path, but does not use such information for tunneling data packets to the MH even after a home registration reply is received. This is due to the fact that the next MH's regional registration might terminate at a crossover FA that is at a lower hierarchy level than the current crossover FA, hence the current crossover FA will not be informed to update its MH's tunnel endpoint. Thus, the crossover FA always establishes the MH's tunnel endpoint as its child FA that originally forwarded the registration request.

Each RFA, beneath the crossover FA in the old path, upon receiving the binding update message from its parent FA performs the following steps: note the current tunnel endpoint for the MH, delete the MH's visitor entry, create a binding cache entry with the specified lifetime (pointing to the noted tunnel endpoint), generate a binding update message to the noted tunnel endpoint (one of its children FA), and generate a binding acknowledgment message back to the sender FA (its parent FA). This process repeats until the BU message along with the new FA information reaches the FA that was previously serving the MH. The old FA, armed with

the new FA information, can send any already buffered, or future data packets to the new FA that delivers them to the MH. The binding caches in the old path will eventually expire. Meanwhile, the home registration reply should be received from the HA switching the MH's tunneling path to the new path alleviating the need for the old path. Figure 3 illustrates the signaling message flow for the KOPA approach.

In the KOPA approach, the crossover FA acts upon the MH's home registration and generates binding updates down the MH's old path, hence altering the MH's tunneling state in affected RFAs. If the crossover FA cannot authenticate the received home registration, then such approach would not be feasible. Hence, we exploit the presence of a hierarchy structure and the existence of an old and new path to the MH by extending such home registration, a HR-LH request, as a special case of a *combined home and regional registration*. Consequently, we require that the MH authenticate any HR-LH requests by using a MH-GFA authentication extension. Moreover, the home registration request does not contain regional identification information, creating a regional replay protection problem. Thus, we introduce a *local replay protection extension* [1] to be supplied by the MH with a HR-LH request to ensure the feasibility of the KOPA approach. The MH uses the local replay protection extension to supply its current regional identification value; to enable the crossover FA's processing mechanism of the home registration request.

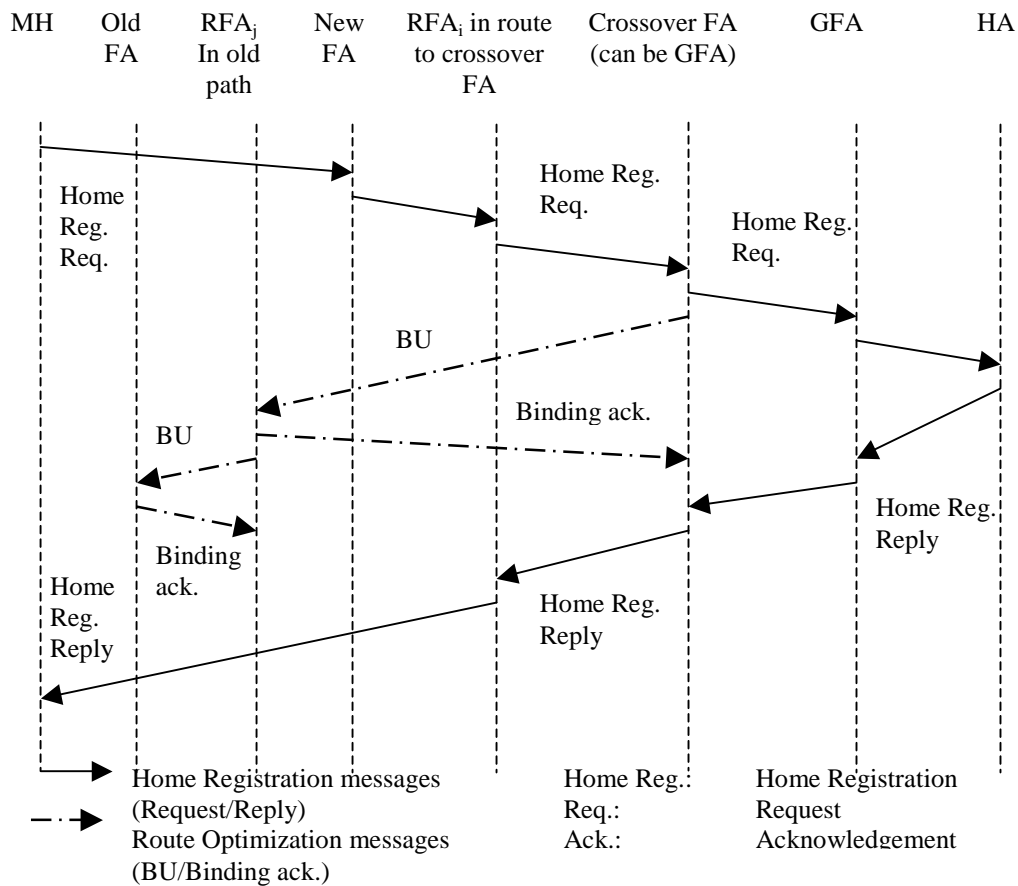


Figure 3: Signaling message flow in the KOPA approach.

If the crossover FA is able to successfully authenticate the MH, it initiates the binding updates down the old path as previously explained, and forwards the MH's request unchanged upwards towards the GFA. Such processing allows RFAs in higher levels than the crossover FA to know the current regional identification used by the MH. Eventually, this request reaches the GFA, which removes any regional information extensions and forwards the request to the HA. If the MH is using nonce replay protection, then a new nonce value needs to be generated and sent back to the MH. In this case, we suggest that the GFA perform this function by placing the new nonce value in a local replay extension to be appended at the end of the home registration reply received from the HA, and authenticated using a MH-GFA authentication extension. The registration reply flows down the hierarchy to reach the MH, informing all RFAs of the new nonce value generated for this MH. If timestamp replay protection is used, then such processing is not needed.

5.2 SINP: Switch Immediately to New Path Approach

The SINP approach reinforces the local handoff aspect while the MH is sending a HR-LH request. We note that the HA does not know which local FA is currently serving the MH, since the home registered care-of address for the MH is the GFA. We argue here that the establishment of the new path within the hierarchy should not be dependent on a home registration reply that only indicates that the GFA has been established as the MH's care-of address. Hence, we capitalize on such issue and view a home registration with an involved local handoff as truly a combined home and regional registration. The regional aspect of the registration is handled by the crossover FA in order to switch the MH's tunneling path from the old path to the new path immediately without waiting for a home registration reply, while the home registration aspect is handled by the HA to renew the MH's home mobility binding. Therefore, we suggest that *the crossover FA switches immediately the MH's tunneling path from the old path to the new path by issuing a regional registration reply in response to the MH's home registration reply*. To enable such functionality, similar to the KOPA approach, the MH formulates the home registration request by supplying any current regional protection information authenticating the request using a MH-GFA authentication extension (see section 5.1). In addition to generating a regional registration reply, the crossover FA forwards the home registration request upwards in the hierarchy towards the GFA for normal home registration processing by the HA. Furthermore, to ensure tunneling consistency, the old path to the MH is cleared by using a tunneling consistency mechanism triggered by the crossover FA similar to the deregistration mechanism introduced in section 4. Figure 4 depicts signaling message flow in the SINP approach.

Similar to the KOPA approach, the crossover FA is required to authenticate the MH request by inspecting the authenticator value in the MH-GFA authentication extension and checking the validity of the regional identification value provided in the local replay protection extension. If the MH fails either test; a regional registration reply is generated with appropriate error code. Nevertheless, the home registration portion of the

request is propagated upwards towards the GFA for normal home registration processing. For requests successfully processed by the crossover FA, the MH request is propagated upwards towards the GFA that removes any regional extensions and forwards the request to the HA. For timestamp replay protection, the crossover FA forwards the MH request unchanged leaving the local replay protection extension, hence allowing for upper-level RFAs to know the current identification value used by the MH. For nonce replay protection, the crossover FA puts a new nonce value in its regional registration to the MH. Two alternatives are available to inform upper-level RFAs about such new nonce value: propagate a separate replay protection update message upwards in the hierarchy towards the GFA [1], or append a local replay protection extension containing the new nonce value to the end of the MH request. In the latter case, the resulting message is authenticated using a FA-FA authentication extension.

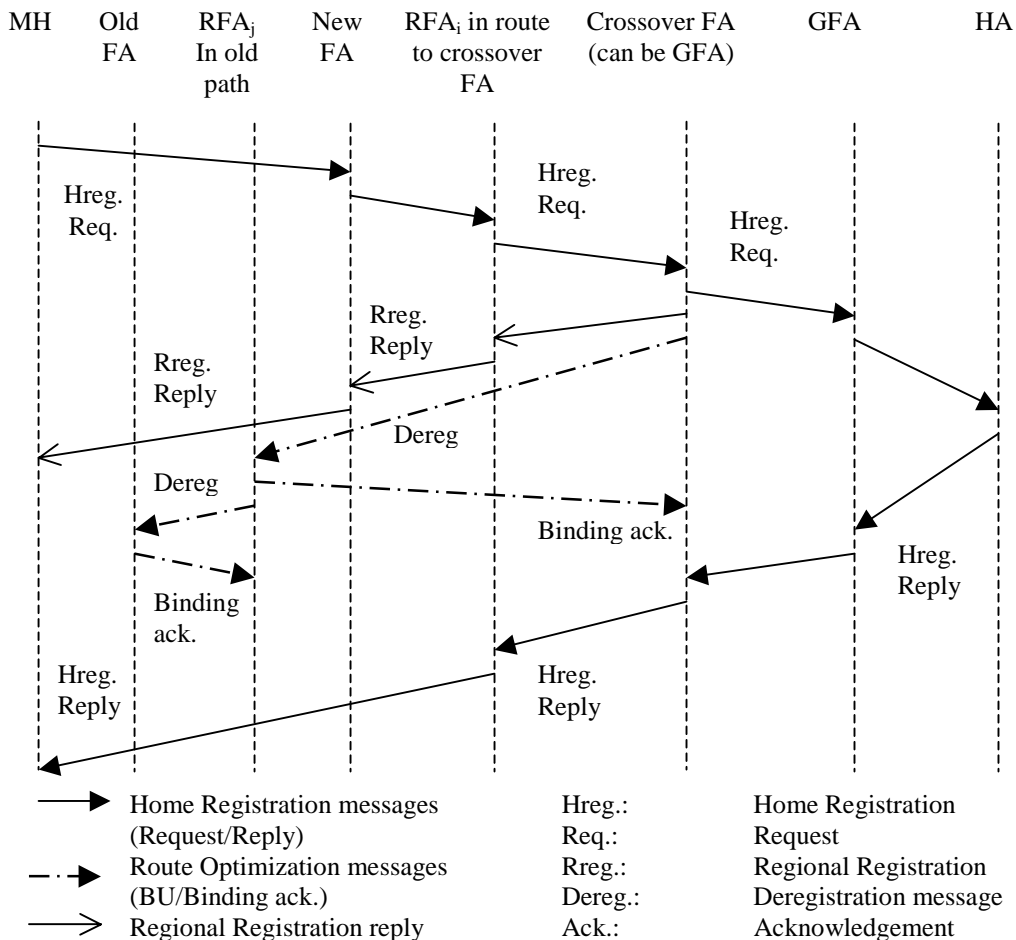


Figure 4: Signaling message flow in the SINP approach.

6 Performance Evaluation

We evaluate the performance of the proposed approaches through network simulation. We implemented the proposed registration processing frameworks by extending the Columbia IP micro-mobility software (CIMS) [5], [6], which is an *ns-2* network simulator [10] source code extension implementing a 1-level foreign agent hierarchy below the GFA. CIMS was extended to simulate *n*-level foreign agent hierarchies, model a true foreign domain, and encompass a local-area mobility network simulation framework [1]. We investigate the performance of both TCP and UDP traffic for KOPA, SINP, and Base Mobile IP (MIP) approaches, and attempt to enforce an adequate number of home registrations involving local handoffs. For comparison, we include a “naïve” HR-LH processing approach where the old path is cleared by the crossover FA upon receiving the home registration propagated through the new path, hence treating HR-LH similar to regional registrations. We term this “naïve” approach the *Delete Old Path* (DOP) approach. Observing the performance of the DOP approach gives an insight of the benefits achieved by the KOPA and SINP approaches.

Figure 5 depicts the simulated network topology. The FA hierarchy is a perfect 4-level binary tree. Leaf foreign agents provide wireless access to the MH, whereas other foreign agents in the hierarchy do not possess such capability. Neighboring base stations’ coverage areas have an overlap region of 30 meters. Each FA is only connected to its children foreign agents through individual 100 Mbps duplex links and link delay LD_{FA-FA} in milliseconds (default link delay is 0.5 ms). The GFA is connected to the MH’s HA through a 1.5 Mbps duplex link with delay LD_{GFA-HA} ms (default link delay is 20 ms). We simulate a single MH within the hierarchy communicating with a fixed *Correspondent Host*.

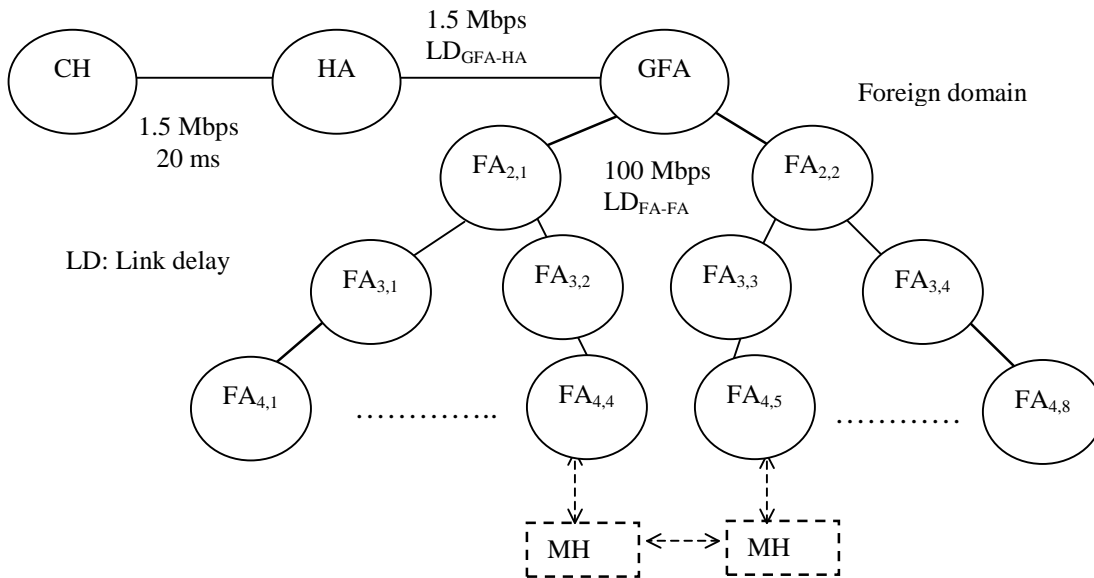


Figure 5: Simulated network topology.

6.1 UDP Traffic

We simulate the behavior of a 64 kbps audio application by applying constant bit rate UDP traffic from the CH to the MH. The MH is moving periodically between $FA_{4,4}$ and $FA_{4,5}$ (3-hop handoffs) at a speed of 10 m/sec until an adequate number of handoffs is attained (more than 100 handoffs). Performance measures include the average number of lost packets per handoff, average number of lost packets per HR-LH, and in the KOPA case, the average number of encapsulated packets from the old FA to the new FA per HR-LH.

6.1.1 Effect of LD_{GFA-HA}

We investigate the effect of the delay to the HA to show the effect of having a distant HA. Figure 6 illustrates the average lost packets per handoff while varying LD_{GFA-HA} from 5 to 50 ms and fixing LD_{FA-FA} at 0.5 ms. The KOPA and SINP approaches outperform both base MIP and DOP, with DOP outperforming base MIP. In general, the number of lost packets per handoff increases linearly with the link delay increase in base MIP and DOP, while it is not affected in KOPA and SINP. For instance, with LD_{GFA-HA} equal to 50 ms (a distant HA) the average lost packets in KOPA and SINP is almost the same, representing a reduction in packet loss of almost 96% compared to base MIP. We note that KOPA and SINP perform similarly in terms of packets lost since the link delays on the old path are the same as on the new path, and KOPA is able to keep up the same performance as SINP, since the number of encapsulated packets from the old FA to the new FA increases linearly with the increase of LD_{GFA-HA} as illustrated in Figure 7. Such linear increase is attributed to the increase of LD_{GFA-HA} while LD_{FA-FA} is fixed, allowing for longer use of the old path by the crossover FA as the MH's tunneling path.

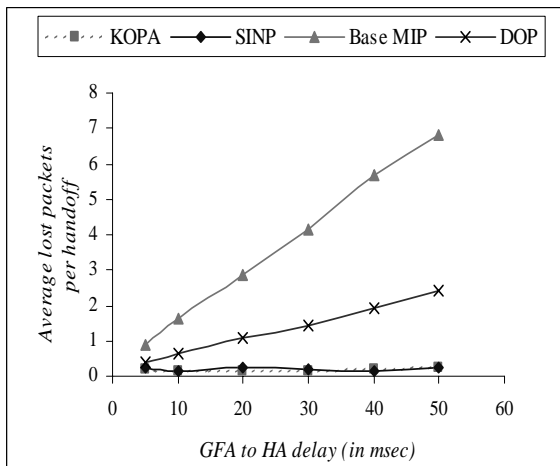


Figure 6: Average lost packets per handoff versus LD_{GFA-HA} .

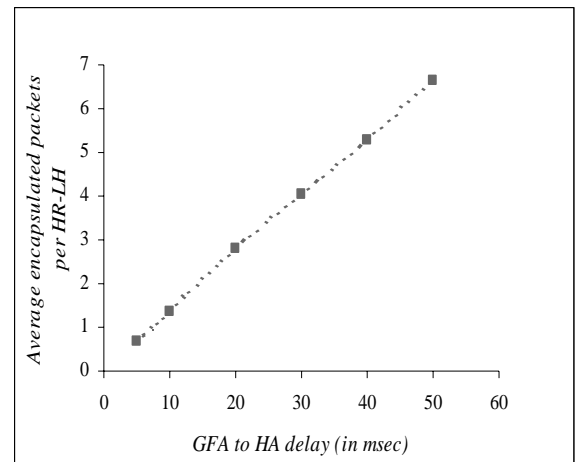


Figure 7: Average encapsulated packets per HR-LH versus LD_{GFA-HA} in the KOPA approach.

We demonstrate the reduction in packet loss achieved by introducing special processing for home registrations involving local handoffs. Figure 8 depicts the average number of lost packets during HR-LH. In DOP, the number of lost packets per handoff increases linearly with the link delay increase, while it is not

affected for KOPA and SINP. Comparing Figure 6 and Figure 8 in the case of DOP, we conclude that the total number of lost packets during HR-LH represents a very high percentage compared to the total number of lost packets over all handoffs. This ratio is bounded between [75-93]% for LD_{GFA-HA} between [5-50] ms; hence the necessity of the KOPA and SINP approaches to handle home registration processing during HR-LH.




Figure 8: Average lost packets per HR-LH versus LD_{GFA-HA} .

6.1.2 Effect of LD_{FA-FA}

We investigate the effect of hierarchy link delay LD_{FA-FA} to show the effect of the delay between the new FA and the crossover FA. Figure 9 illustrates the average lost packets per handoff when varying LD_{FA-FA} between 0.5 and 5 ms while fixing LD_{GFA-HA} at 20 ms. All approaches exhibit average lost packets per handoff linearly increasing with the increase of LD_{FA-FA} . However, a sizable packet loss reduction can be achieved with KOPA and SINP. For example, for an LD_{FA-FA} value of 5 ms, the average lost packets per handoff in KOPA and SINP represent a reduction in packet loss of 57% versus base MIP. The increase in propagation delay between the new FA and the crossover FA is the cause for the linear increase in packet loss in KOPA and SINP, since it takes longer for the MH's request to get to the crossover FA, allowing for increased packet loss over the old path in KOPA while the old FA is not informed about the new FA, and in SINP before the switch to the new path is performed by the crossover FA. The SINP approach switches the MH's tunneling path immediately, while KOPA uses the old path for a period of time until the home registration reply is received at the crossover FA.

Figure 10 shows the average number of lost packets per HR-LH. The previous discussion applies here that the longer it takes for the registration request to reach the crossover FA, the higher the resulting packet loss. Comparing Figure 9 and Figure 10, when LD_{FA-FA} is 5 ms, the percentage of packets lost during HR-LH compared to the total number of packets lost is 36% and 35% in the KOPA and SINP approaches, respectively. We conclude that the number of hops (number of intermediate RFAs) between the new FA and the crossover FA affect the performance of KOPA and SINP. In general, the usage of the smooth handoff mechanism by the MH

and the serving FAs would reduce the observed packet loss, since the new FA (the earliest point of contact in the FA hierarchy) sends a BU message to the old FA on behalf of the MH. We note that our proposed approaches optimize registration processing within the FA hierarchy resulting in packet loss reduction compared to base MIP, and are not meant as a replacement for the smooth handoff mechanism.

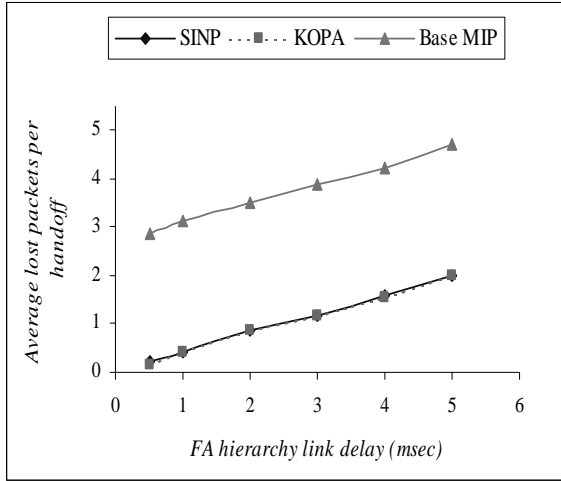


Figure 9: Average lost packets per handoff versus LD_{FA-FA} .

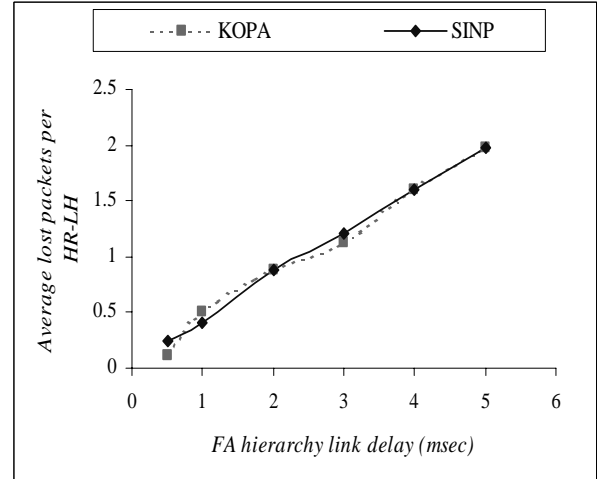


Figure 10: Average lost packets per HR-LH versus LD_{FA-FA} .

6.1.3 Effect of tunneling in KOPA

The efficiency of the tunneling approach in KOPA depends on the following factors.

- *How fast the old FA is informed about the new FA.* This affects the observed number of lost packets if the old FA is not providing packet-buffering services to the MH, or the MH takes too long to find a new FA.
- *The propagation delay between the old FA and new FA.* Tunneled packets experience such delay increasing packet latency, which affects multimedia applications attempting to deal with network jitter.

In addition, the existence of an old path with a tunnel and a new path for data packets can create out of order packets. We consider the case where LD_{FA-FA} is 5 ms and observe the received packet order. On the average 1 packet was received out of order during HR-LH in KOPA, while no such behavior was observed with SINP. Moreover, a constant playout delay is maintained and varied while we measure the average number of application-dropped packets per handoff in KOPA and SINP. A packet playout time is calculated to be the packet send time + playout delay. If a packet is received after its playout time it is dropped. Figure 11 illustrates the average number of application-dropped packets per handoff versus the playout delay in ms. We note that zero dropped packets can be achieved in SINP by maintaining a smaller playout delay (around 60 ms) than

KOPA (around 90 ms). In our network topology, the old path is symmetric to the new path, and hence the added latency in KOPA is due to the tunneling process from the old FA to the new FA during HR-LH.

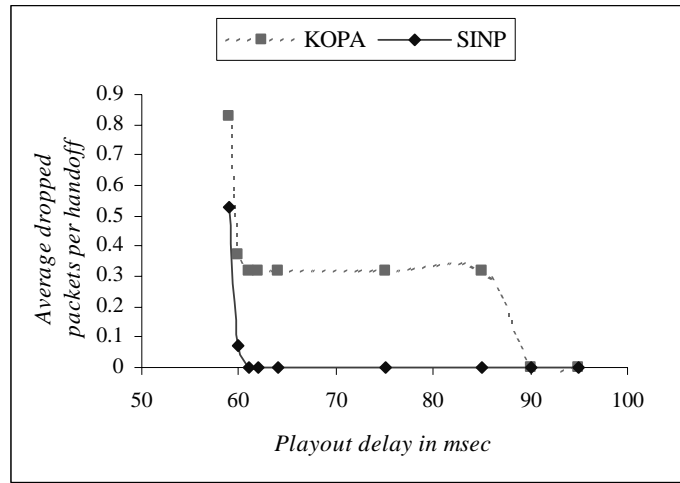


Figure 11: Average number of dropped packets per handoff versus playout delay.

6.2 TCP Traffic

We simulate the behavior of a long-term FTP session between the MH and the CH where the MH is downloading a very large file from the CH. The MH is periodically moving between $FA_{4,4}$ and $FA_{4,5}$ at a speed of 20 m/s until an adequate number of handoffs is attained (> 100 handoffs). We vary LD_{GFA-HA} between 5 and 50 ms and measure the observed application-level TCP throughput. TCP Tahoe [16] was used for the purpose of this simulation.

Figure 12 illustrates TCP throughput in Mbps versus LD_{GFA-HA} . As expected, the application-perceived TCP throughput degrades with the link delay increase due to increased round trip times. Base MIP exhibits the worst throughput degradation since the throughput at an LD_{GFA-HA} value of 50 ms represents a drop of 30% compared to an LD_{GFA-HA} value of 5ms. KOPA and SINP demonstrate similar behavior with the throughput degradation at link delay of 50 ms representing a drop of 9% and %10 respectively compared to a 5 ms link delay. At smaller link delays, e.g., 5 ms, all investigated approaches achieve a comparable throughput value. With the link delay increase, KOPA and SINP attain higher throughput than base MIP and DOP, e.g., at a link delay of 50 ms KOPA represents a throughput increase of 34% and 8% over base MIP and DOP, respectively. Such higher sustained throughput is exhibited due to the reduction in packet loss, consequently requiring fewer number of TCP retransmissions, which translates into higher TCP throughput. We investigate TCP's retransmission behavior during the simulation by measuring the retransmission ratio, calculated as the number of retransmitted packets to the total number of transmitted packets (Figure 13). Base MIP is able to keep up with the increased number of lost packets by increasing the number of retransmitted packets, hence leading to an overall lower throughput value.




Figure 12: TCP throughput versus LD_{GFA-HA} .




Figure 13: TCP Retransmission ratio versus LD_{GFA-HA} .

We demonstrate the effect of increasing the handoff rate on TCP throughput. We consider 3 speeds for the MH: 5, 10, and 20 m/s corresponding to average handoff rates of 5, 10, and 15 handoffs per minute. In addition, we select 2 values for LD_{GFA-HA} : 5 and 50 ms and compare base MIP, KOPA, and SINP. Figure 14 depicts the TCP throughput versus the MH speed (handoff rate). Solid lines represent a link delay of 5 ms (nearby HA), while the dashed lines represent a link delay of 50 ms (distant HA). For a nearby HA, the throughput drop with the handoff rate increase is not that significant for all approaches. For a distant HA, base MIP exhibits the worst throughput drop with the handoff rate increase attributed to an increase in the number of lost packets (at 20 m/s a throughput drop of 20% is reported compared to 5 m/s), whereas KOPA and SINP report a modest throughput drop with the handoff rate increase (at 20 m/s, throughput drops are 3% and 2% for KOPA and SINP, respectively, compared to 5 m/s).

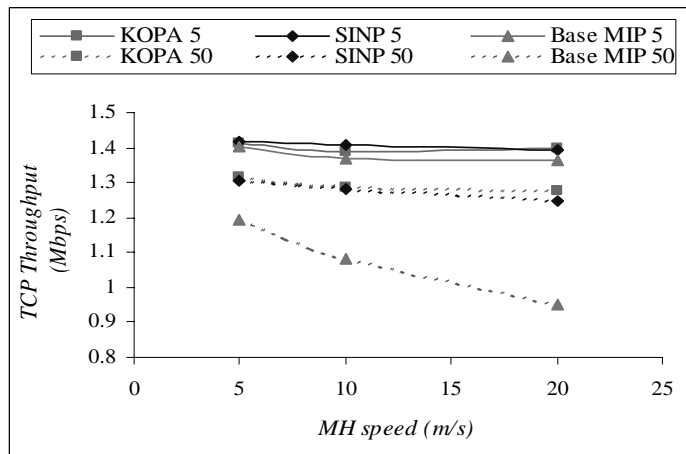


Figure 14: Effect of handoff rate on TCP throughput.

7 Conclusion and Future Work

In this paper, we presented registration frameworks for IPv4 mobility agents' hierarchies for intra-hierarchy handoffs. We introduced an enhanced regional registration framework that prevents some of the identified deficiencies with the current regional registration proposal. In addition, we introduced two novel home registration frameworks when associated with local handoffs. The proposed techniques achieve the same level of security as base Mobile IP by providing authentication and replay protection for all messages. Network simulation experiments have demonstrated the effectiveness of the proposed techniques in achieving up to 96% reduction in UDP packet loss, and maintaining up to 34% increase in TCP throughput, versus a base Mobile IP implementation, for a distant HA. Our techniques constitute a building block within a local-area mobility support based on deploying multiple cooperating foreign agent hierarchies in the foreign domain [1], [2].

Future work includes investigating an approach for the automatic set up and maintenance of FA hierarchies in the foreign domain, extending and adapting the registration frameworks when the mobility support overlay network is actually a forest and not a tree, implying the possibility of existence of multiple parent FAs for a specific FA, and investigating techniques to improve the fault tolerance of the proposed registration techniques. Furthermore, we intend to compare through network simulations between a suite of local-area mobility support solutions including the FA hierarchy approach, Cellular IP, and HAWAII.

References

- [1] A. Abdel-Hamid, "A Framework Of Cooperating Agents Hierarchies For Local-area Mobility Support", Ph.D. Thesis, Department of Computer Science, Old Dominion University, Norfolk, VA 23505, Nov. 2002.
- [2] A. Abdel-Hamid and H. Abdel-Wahab, "Local-area Mobility Support through Cooperating Hierarchies of Mobile IP Foreign Agents," in *Proc. of the 6th IEEE International Symposium on Computers and Communications (ISCC'2001)*, Hammamet, Tunisia, July 2001, pp. 479-484.
- [3] B. Aboba and M. Beadles, "*The Network Access Identifier*", Request For Comments (Proposed Standard) 2486, Internet Engineering Task Force, Jan. 1999.
- [4] P. Bhagwat, C. Perkins, and S. Tripathi, "Network Layer Mobility: an Architecture and Survey", *IEEE Personal Communications*, Vol. 3, No. 3, pp. 54-64, June 1996.
- [5] A. T. Campbell, J. Gomez, S. Kim, Z. Turanyi, C-Y. Wan, and A. Valko, "Comparison of IP Micro-Mobility Protocols," *IEEE Wireless Communications*, vol. 9, no. 1, pp. 72-82, Feb. 2002.
- [6] Columbia IP Mirco-Mobility Software (CIMS). [Online]. Available: <http://www.comet.columbia.edu/micromobility/>
- [7] G. Dommety and T. Ye, "Local and Indirect Registration for Anchoring Handoffs", Internet Draft (Work in Progress), draft-dommety-mobileip-anchor-handoff-03.txt, July 2001.
- [8] D. Forsberg, J. T. malinen, J. K. Malinen, T. Weckstrom, and M. Tiusanen, "Disributing Mobility Agents Hierarchically under Frequent Location Updates", in *Proceedings of 6th IEEE International Workshop on Mobile Multimedia Communications (MoMuC)*, San Diego, CA, USA, Nov. 15-17, 1999.
- [9] E. Gustafsson, A. Jonsson, and C. Perkins, "Mobile IPv4 Regional Registration", Internet Draft (Work in Progress), draft-ietf-mobileip-reg-tunnel-06.txt, Mar. 2002.
- [10] S. McCanne and S. Floyd, *ns Network Simulator*. [Online]. Available: <http://www.isi.edu/nsnam/ns/>.
- [11] C. Perkins (Editor), "*IP Mobility Support for IPv4*", Request for Comments (Proposed Standard) 3344, Internet Engineering Task Force, Aug. 2002.
- [12] C. Perkins and P. Calhoun, "Mobile IPv4 Challenge/Response Extensions", Request For Comments (Proposed Standard) 3012, Internet Engineering Task Force, Nov. 2000.

- [13] C. Perkins and D. Johnson, "Route Optimization in Mobile IP", Internet Draft (Work in Progress), draft-ietf-mobileip-optim-011.txt, Sept. 2001.
- [14] C. Perkins and K.-Y. Wang, "Optimized Smooth Handoffs in Mobile IP", in *Proceedings of 4th IEEE Symposium on Computers and Communications (ISCC)*, pp. 340-346, Red Sea, Egypt, July 6-8, 1999.
- [15] R. Ramjee, T. La Porta, S. Thuel, K. Varadhan, and S. Wang, "HAWAII: A Domain-based Approach for Supporting Mobility in Wide-area Wireless Networks", in *Proceedings of 7th IEEE International Conference on Network Protocols (ICNP)*, Toronto, Canada, 1999.
- [16] W. Stevens, *TCP/IP Illustrated, Volume 1: The Protocols*, Reading, MA: Addison-Wesley, 1994.
- [17] A. Valko, "Cellular IP: A New Approach to Internet Host Mobility", *ACM Computer Communication Review*, Vol. 29, No. 1, pp. 50-65, Jan. 1999.