


IDMP: AN INTRADOMAIN MOBILITY MANAGEMENT PROTOCOL FOR NEXT-GENERATION WIRELESS NETWORKS

SUBIR DAS, ANTHONY MCAULEY AND ASHUTOSH DUTTA, TELCORDIA TECHNOLOGIES INC.

ARCHAN MISRA, IBM T. J. WATSON RESEARCH CENTER

KAUSHIK CHAKRABORTY, UNIVERSITY OF MARYLAND

SAJAL K. DAS, THE UNIVERSITY OF TEXAS AT ARLINGTON



There has been much interest in efficient IP-based micro-mobility management schemes for next-generation wireless networks. Such schemes are essential to achieve seamless integration of cellular networks with existing IP-based data networks.

ABSTRACT

This article describes a lightweight *intradomain mobility management protocol* (IDMP) for managing mobility within a domain, commonly known as *micromobility management*, for next-generation wireless networks. IDMP is modular and simple because it leverages existing protocols, such as Mobile IP or SIP as global mobility management, for locating roaming nodes. Unlike other proposed intradomain mobility management schemes, IDMP uses two dynamically auto-configured care-of addresses for routing the packets destined to mobile nodes. The global care-of address is relatively stable and identifies the mobile node's attachment to the current domain, while the local care-of address changes every time the mobile changes subnets and identifies the mobile's attachment to the subnet level granularity. After describing the lightweight base protocol, we discuss possible enhancements to reduce the latency of intradomain updates during handoffs, which are critical for real-time applications both for wide area cellular networks and enterprise wireless LANs. We also discuss mechanisms to incorporate paging support in IDMP and hence reduce the mobility-related signaling load on a mobile node. Detailed implementation and performance results from experiments on our testbed are also presented.

INTRODUCTION

In recent years, there has been much interest in developing efficient IP-based *micromobility* management schemes to handle node mobility within a domain in next-generation wireless networks. Such schemes are essential to achieve seamless integration of cellular networks with existing IP-based data networks, popularly known as the Internet. Many cellular network providers and operators have already realized the need for an IP-based mobility management solution to support real- and non-real-time applications in next-

generation networks. However, Internet protocols are currently unable to support the additional performance guarantees these applications require at the user level. Based on a survey of requirements, we can identify, at a high level, the following features desired of any intradomain mobility management solution.

Support for fast handoffs: The mobility management architecture and protocol should be able to seamlessly redirect packets to the mobile's new point of attachment with minimum latency. To support real-time IP applications, including voice over IP (VoIP), the latency typically associated with the registration process must be decreased and bounded.

Reduction in packet loss during movement: With new emerging applications, especially for cellular networks that use unreliable transport protocols [1] for packet transport, the packet loss during handoffs should be minimum.

Support for paging: Paging is important in power-conscious environments since it enables a mobile node to significantly reduce its mobility-related signaling traffic. Next-generation cellular networks are likely to see a proliferation of power-conscious miniature devices and appliances. Any mobility management protocol for such networks should have the option to provide paging support.

Support for multipath distribution techniques: The intradomain mobility management protocol should be able to support multiple traffic paths, typically used for providing redundancy and greater transmission reliability. This support must be optional and configurable only when the link and physical layer technologies permit. At this point, it is not very clear whether such support is necessary for currently emerging wireless access technologies.

The current standard for IP-based mobility management, namely Mobile IP [2], was designed primarily for environments where the mobile node (MN) was assumed to have a well-defined home network and a topologically correct care-of address (CoA) in the foreign network. In such pre-

dominantly static environments, the frequency and volume of global registration messages generated by mobile nodes are not a major concern. Mobile IP ensures transparency to TCP connections by preserving the fixed home address of the MN and performing packet redirection (using tunneling) at the network layer. In the absence of a set of viable real-time or delay-sensitive applications, the latency involved in updating the remote home agent (HA) or correspondent node (CN) on every subnet change was also not a topic of practical concern. Moreover, the base Mobile IP assumes that the rate of subnet change by an MN is not too rapid; the specifications state that Mobile IP is intended for situations where the MN does not change subnets more than once every second [2].

On the other hand, the signaling overhead in next-generation wireless networks, where every active node is likely to exhibit significant mobility, can become very large. In practical wide-area cellular networks, topology considerations, frequency, and address space limitations (e.g., in IPv4) may also cause an IP subnet to span a fairly limited geographical area. Thus, a mobile may change subnets fairly frequently, especially if the trend toward picocellular networks in urban areas continues. Therefore, a separate protocol for supporting intradomain mobility becomes necessary. The *intradomain mobility management protocol* (IDMP) proposed in this article fosters a more modular network architecture and allows static Internet hosts to communicate with mobile nodes without any changes. This fits nicely with the requirements for a variety of applications in next-generation cellular networks. Additionally, unlike the conventional Internet, where backward compatibility is not a major concern, cellular networks have no “IP legacy” issues.

The rest of the article is organized as follows. We briefly discuss the related work. We give an overview of the base protocol IDMP. The next two sections describe the enhancements to IDMP to support fast handoffs and paging, which are critical for next-generation networks. Prototype implementation and testbed layout of our protocol (IDMP) are presented. Finally, we compare the signaling load of IDMP with that of basic Mobile IP. We also tabulate preliminary initial performance results of IDMP. The final section concludes the article.

RELATED WORK

Recently various enhancements have been proposed to overcome the shortcomings of base Mobile IP (or MIP), for example, MIP-RO [3, 4], MIPv6 [5], HMIPv6 [6], HAWAII [7], Cellular IP [8], and fast handoffs [9–11].

As pointed out earlier, for cellular environments with a large number of MNs and real-time VoIP traffic, Mobile IP suffers from several shortcomings, including high update latency, large global signaling load, and lack of paging support. These problems are also present in various other nonhierarchical MIP solutions, such as MIP-RO [3] and MIPv6 [5].

One approach to intradomain mobility management is the route modification approach, characterized by Cellular IP (CIP) [8] and HAWAII [7]: the MN is assigned a CoA that is valid

throughout the domain, and host-specific routes are used to track the MN’s precise location in the domain. The other one is the multi-CoA approach: an MN is assigned multiple CoAs, each resolving the MN’s location at an intermediate level in the hierarchy. Among these schemes, Mobile IP Regional Registration (MIP-RR) [4] uses a gateway foreign agent (GFA) to provide an MN a stable global CoA; the GFA acts as a proxy for the HA during any subsequent intradomain movement. Similarly, Hierarchical MIPv6 (HMIPv6) [6] introduces an agent called the MAP to localize the intradomain mobility management. Two alternative schemes [10, 11] for providing fast handoff, within the MIP context, have also been recently proposed. The advantages and disadvantages of these schemes are available in [12].

Mobility mechanisms based on Session Initiation Protocol (SIP) [13, 14] also provide an alternative application-layer mobility management technique, especially for real-time multimedia applications. In general, the SIP-based solution is analogous to MIPv6, with the MN sending each active correspondent node (CN) a Re-INVITE (asking it to rejoin at the new CoA) and the appropriate SIP server a new REGISTER (updating the binding between the SIP UserID and the current CoA). VoIP traffic benefits from such a mechanism, since it allows a CN to send traffic directly to the MN’s collocated CoA (without tunneling), and permits the application to control the characteristics of an ongoing session as the MN changes subnets.

The recently proposed Telecommunication Enhanced Mobile IP (TeleMIP) [15] is a scalable and hierarchical IP-based architecture that provides lower handoff latency and signaling overhead compared to Mobile IP. However, like HAWAII and Cellular IP, TeleMIP uses Mobile IP as the global mobility management protocol.

IDMP OVERVIEW

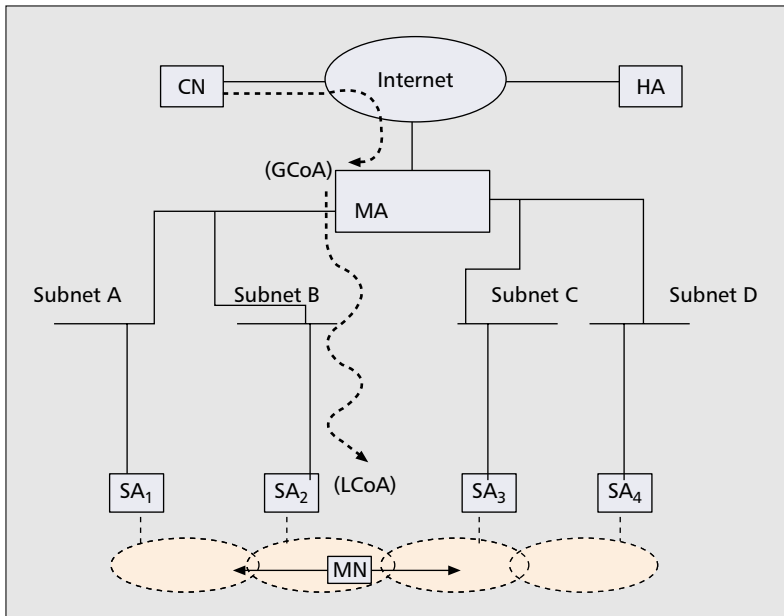
The IDMP proposed in this article is an extension to the base intradomain protocol used in TeleMIP. An Internet draft on IDMP [16] was recently proposed for supporting several additional mobility features, such as minimally interrupted handoff and paging, within the mobility domain for highly mobile users. This separation of intradomain mobility from interdomain mobility is intended to allow a common base protocol to coexist with multiple alternatives for global mobility management, including Mobile IP and SIP. An architecture called Dynamic Mobility Agent (DMA) [17] was also recently proposed and uses IDMP as the base mobility management protocol to provide a scalable and robust mobility management framework.

This section describes the base protocol. Its enhancement to support fast handoffs and paging will be presented in subsequent sections. Packet formats and other detailed specifications are presented in [16].

BASE PROTOCOL

IDMP offers intradomain mobility by using multi-CoAs. However, unlike HAWAII, MIP-RR, or HMIPv6, our protocol IDMP is designed as a standalone solution for intradomain mobili-

The recently proposed TeleMIP is a scalable and hierarchical IP-based architecture that provides lower handoff latency and signaling overhead compared to Mobile IP. However, like HAWAII and Cellular IP, TeleMIP uses Mobile IP as the global mobility management protocol.

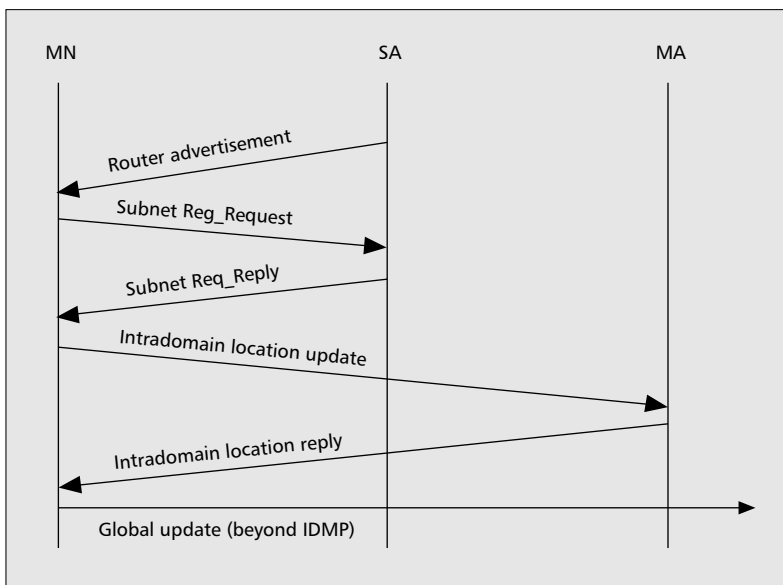


■ Figure 1. IDMP logical elements and architecture.

ty and does not assume the use of MIP for global mobility management. Figure 1 depicts the functional layout of IDMP. The mobility agent (MA) is similar to a MIP-RR GFA and acts as a domain-wide point for packet redirection. A subnet agent (SA) (similar to a MIP FA in CoA mode and DHCP/DRCP [19, 20] server in collocated CoA mode) provides subnet-specific mobility services. Under IDMP, an MN obtains two concurrent CoAs:

Local care-of address (LCoA): This identifies the MN's attachment to the subnet. Unlike MIP's CoA, the LCoA in IDMP only has local (domain-wide) scope. By updating its MA of any changes in the LCoA, the MN ensures that packets are correctly forwarded within the domain.

Global care-of address (GCoA): This address resolves the MN's current location only up to a



■ Figure 2. IDMP message flow during the initial intradomain location update.

domain-level granularity and hence remains unchanged as long as the MN stays within a single domain. By issuing global binding updates that contain this GCoA, the MN ensures that packets are routed correctly to its present domain.

Under IDMP, packets from a remote CN are forwarded (with or without tunneling) to the GCoA and are intercepted by the MA. As shown in Fig. 1, the MA then tunnels these packets to the MN's current LCoA. Since global binding updates are generated only when the MN changes domains and obtains a new GCoA, this approach drastically reduces the global signaling load. Further details of IDMP, and its use with MIP, are available in [15, 16].

Basic Packet Redirection and Mobility Support — When the MN first moves into a domain, it obtains an LCoA (this LCoA is SA₂'s address in Fig. 1) by performing a subnet-specific registration using IDMP. As requested by IDMP, the serving SA (SA₂ in this case) dynamically assigns the MN an MA during this subnet-specific registration process. The MN then performs an intradomain location update by communicating its current LCoA to the designated MA. The MA includes either its address or a separate GCoA in the intradomain location update reply. Subsequently, the MN is responsible for generating a global location update (registration) to the necessary remote nodes (e.g., HA if Mobile IP is used for global mobility management or Registrar [LR] if SIP is used); this is, however, independent of the IDMP specifications. The IDMP call flow when the MN first moves into a new domain is illustrated in Fig. 2.

After the initial intradomain registration process, IDMP now allows the MN to retain its GCoA as long as it stays within the same domain. Whenever MN changes subnets within this domain, it performs a new subnet-specific registration with the new SA. Since the MN indicates that it has an existing valid registration, the SA does not allocate it a new MA address in this case. The MN then performs a new intradomain location update and informs its MA of its new LCoA. No global messages are generated in this case, since the GCoA remains unchanged. As with other hierarchical mobility management schemes, the localization of intradomain mobility significantly reduces the latency of handoffs across subnets within the same domain and also dramatically decreases the frequency of global signaling traffic. Figure 3 describes the IDMP call flow during subsequent intradomain movement.

FAST HANDOFF SCHEME IN IDMP

In the basic mode of IDMP, the handoff delay (or the service interruption time) equals the time taken from a disconnection until the MA becomes aware of the MN's new point of attachment (LCoA) and begins redirecting packets correctly again. In a cellular network architecture where an IP-based base station (IPBS) is used, this delay essentially consists of three components:

Radio channel establishment delay (Δ_1): The MN must establish a new radio channel at the new BS. This is a link-layer-specific function, and could involve even operations such as slot specification in time-division multiple access

(TDMA) or code synchronization in code-division multiple access (CDMA).

IP subnet configuration delay (Δ_2): An MN must use IP-layer configuration protocols to obtain the new LCoA. If IDMP's SA mode is used, the MN must obtain an *agent advertisement* through router discovery or some other beacon and then request a new LCoA. The SA will then respond with an Acknowledgment message. If the collocated mode is used, the MN must exchange DHCP configuration messages with the DHCP server before obtaining a valid CoA.

Intra-domain Update Delay (Δ_3): The MN must finally inform the MA of this new LCoA via an intradomain location update message. The MA will redirect packets to the MN's new LCoA only after receiving this message.

The parameter Δ_1 , although link-layer-specific, can be expected to be quite low. For example, in CDMA-based soft handoffs, Δ_1 is effectively 0, since in such a network communication with the old BS is not discontinued until the connection with the new BS is firmly established (commonly known as *soft handoff*). Even under the hard handoff scenario, no disruption to the radio-level connectivity should occur in a well-designed system: the various elements should coordinate to ensure a synchronized switch to the new point of attachment. IDMP's fast handoff mechanism is designed to eliminate the Δ_3 component in the handoff delay. To make IDMP's operation independent of current or future link-layer techniques, we do not provide IP-level connectivity until the MN has performed a subnet-level configuration at the new BS. IDMP's fast handoff process thus does not eliminate Δ_2 , the delay incurred in the subnet-level configuration process.

THE FAST HANDOFF PROCEDURE

IDMP's fast handoff procedure is based on the assumption that a layer 2 trigger will be available (to either the MN or the old BS) indicating an imminent change in connectivity. We explain the fast handoff mechanism with the help of Fig. 4, which shows an MN moving from SA_2 to SA_3 . To minimize service interruption during the handoff process, IDMP requires either the MN or the old SA (SA_2) to generate a *MovementImminent* message to the MA serving the MN. Upon reception of this message, the MA multicasts all inbound packets to the entire set of neighboring SAs (SA_3 and SA_1 in this case). Each of these candidate SAs buffers such arriving packets in individual MN buffers, thus minimizing the loss of in-flight packets during the handoff transient. When the MN subsequently performs a subnet-level configuration (using IDMP messages) with SA_3 , the latter can immediately forward all such buffered packets over the wireless interface, without waiting for the MA to receive the corresponding intradomain location update. Several features of this proposal make it attractive for future IP-based networks, such as:

- Unlike other fast handoff proposals, IDMP's *MovementImminent* message does not specify the IP address of the target (new) BS in this message.
- IDMP utilizes a *network-controlled* (network or mobile-initiated) handoff technique.
- IDMP's fast handoff scheme does not eliminate

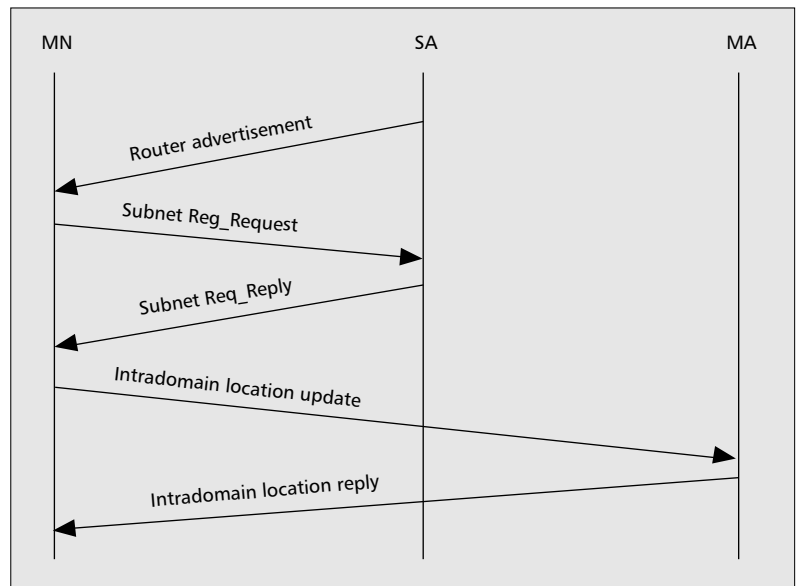


Figure 3. IDMP call flow during subsequent intradomain movement.

Δ_2 from the service interruption time; it merely delays the transmission of packets arriving during this instant. Details of this proposal are described in [17].

PAGING SUPPORT IN IDMP

While IDMP's use of multicasting for fast handoffs minimizes the loss of in-flight packets during an intradomain handoff, it does not reduce the frequency of intradomain location updates. In the absence of paging support, an MN must obtain a LCoA and reregister with its MA every time it changes its current subnet. This can lead to significant power wastage, especially in future 4G networks where a single device may maintain multiple simultaneous bindings with multiple radio technologies. IDMP's IP-layer paging solution provides a flexible and radio-technology-independent solution to this important problem.

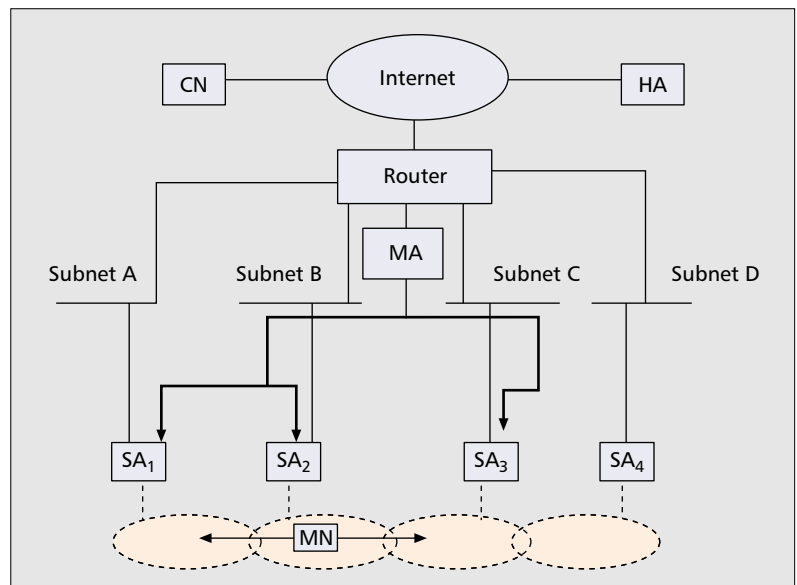
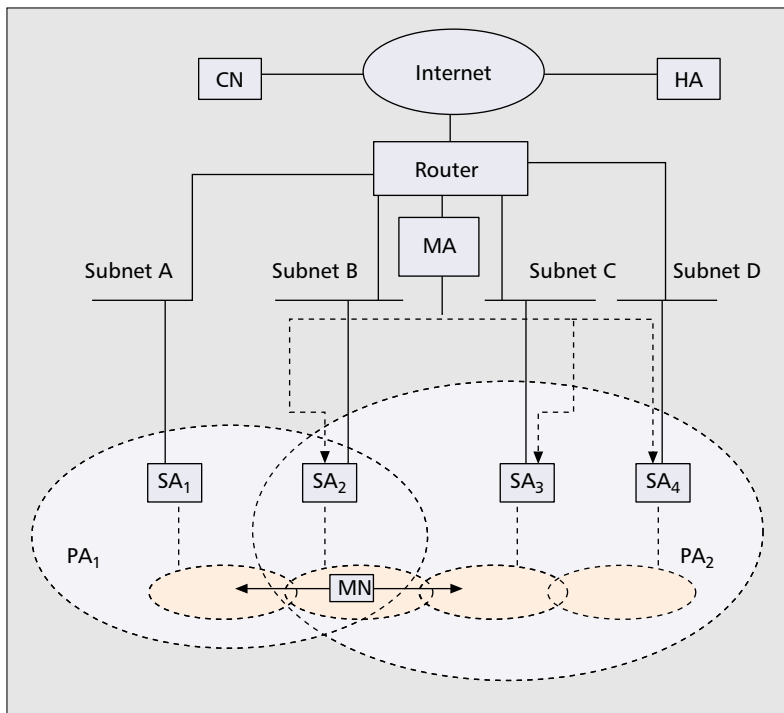
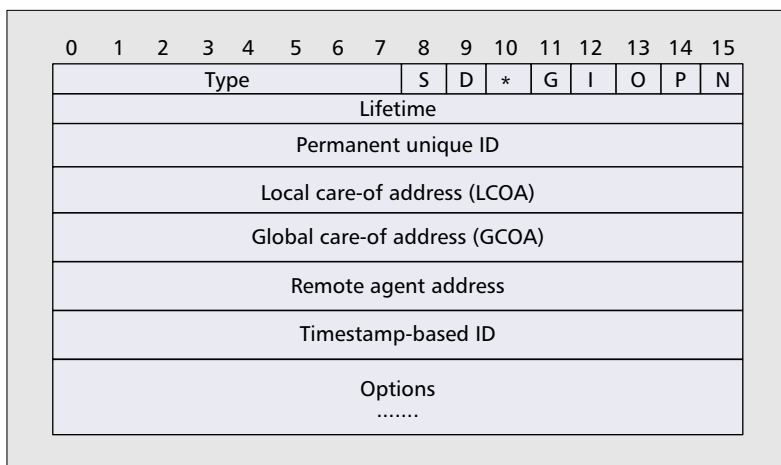


Figure 4. IDMP fast handoff.

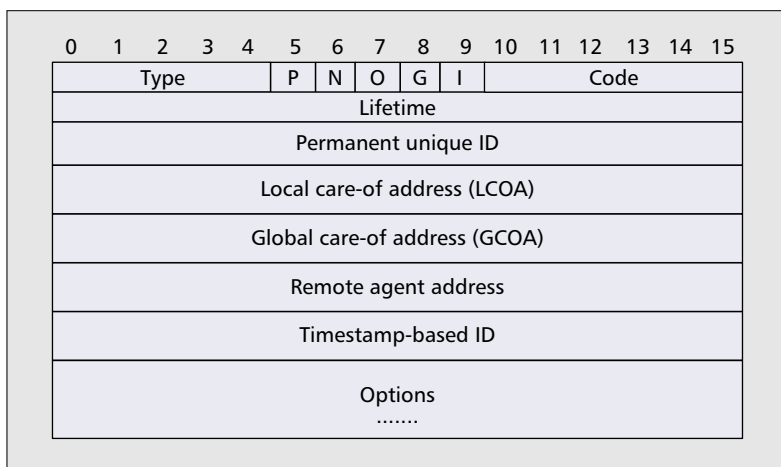
PAGING OPERATION FOR IDLE HOSTS



■ Figure 5. IDMP paging mechanism.



■ Figure 6. IDMP intradomain location update packet format.



■ Figure 7. IDMP intradomain registration reply packet format.

To motivate IDMP's paging solution, we refer to the *multicasting* scheme described for fast hand-off support in the last section. In fast handoff, the MA essentially sends multiple copies of the same data to multiple SAs/subnet routers that are judged to be in the vicinity of the MN's current point of attachment. Since limited broadcast of solicitations is really the central feature of paging, the idea of multicast groups can be extended to provide paging support as well. IDMP's paging operation assumes that SAs (subnets or IPBSs) are grouped into *paging areas* (PAs) identified by some unique identifiers. An MN in passive/idle mode is then able to detect changes in its current PA by listening to these unique identifiers in the subnet-level advertisements (e.g., subnet agent advertisements). In fact, such IP-layer advertisements may optionally be combined with link-layer beacons.

IDMP's paging scheme is illustrated in Fig. 5. In this model of operation, subnets B, C, and D belong to the same PA₂, while subnet A is part of a different PA, PA₁. We assume that the MN switches to idle state in subnet B. Then, as long as it moves to C or D, it detects changes in its subnet of attachment but no change in its current PA. Consequently, not only does the MN not update its MA about its current LCoA, but also does not bother to obtain a new LCoA. However, when it moves to subnet A and realizes that it has changed to a new PA, the MN obtains a new LCoA at SA₁ and sends a location update to the MA, indicating the new paging area.

When the MA receives packets for an MN that is currently registered, but does not have a valid LCoA assigned, it multicasts a *PageSolicitation* packet to all the subnets associated with the MN's current PA (i.e., to SA₂, SA₃ and SA₄) and buffers the incoming packets. When the MN reregisters with the MA, the buffered packets are forwarded to the MN. We assume that temporary buffering is acceptable since the intradomain location update process is assumed to have reasonably low latency ($\sim 2 \times \Delta$, where Δ is the delay between the MN and its MA). For VoIP, the call setup delays are typically around 2.5 s [20]; accordingly the paging latency is expected to fall within the targeted bounds. A comparison with alternative IP paging schemes is discussed in [16].

IMPLEMENTATION OF IDMP

This section describes our basic prototype implementation of IDMP. Before describing the IDMP packet formats we briefly discuss here some functional enhancements to our base implementation.

The MA handles local registration requests from MNs currently in its domain, and provides temporary bindings to the MNs as long as they remain in the domain. As far as the handling of such registration (or location update) requests is concerned, there is little functional difference between HA and MA. Unlike the HA, which has a permanent list of mobility bindings for each MN associated with its home network, the MA maintains a dynamic list of mobility bindings for currently registered MNs. The major functional

difference between HA and MA is in terms of packet forwarding to the MN. When the MN is away from the home network, the HA is responsible for collecting all the packets directed to the MN's permanent IP address and tunneling the packets to the GCoA (which is also the IP address of the MA interface). The task of the MA is simpler; it receives the packets automatically, and after decapsulating the packets redirects the inner IP packet to the MN's LCoA.

In fact, the HA is potentially unaware of the use of IDMP and the presence of the MA. As in conventional MIP, it simply has to intercept all packets intended for the MN from the home network, encapsulate them, and forward them to the CoA specified in the MN-HA registration message. The registration request and reply message formats for global registrations are, in fact, identical to MIP with a single exception: the reserved bit in the flags field in [21] is now used to indicate whether the MN is operating in a DMA-based network.

For simplicity, our implementation of IDMP was based on enhancements to the Stanford University MosquitoNet [21] Linux Mobile IP code. Implementation details are available in [22].

IDMP PACKET FORMATS

MNs in our architecture use IDMP messages to register their local care-of address with the designated MA. While IDMP packet formats and location update messages are based on MIP, they have been modified to support additional intradomain mobility features. Figures 6 and 7 show the IDMP packet formats for intra-domain registration request and reply messages respectively. Our current implementation supports only the collocated mode for local addressing. An MN thus uses DHCP [18] or DRCP [19] to obtain an LCoA; subnet-level registrations (between the MN and an SA) are consequently not described in this article. For additional details on the individual message fields, refer to [16]. Since support for paging and fast handoff is not supported in our current implementation, the corresponding flags (P and O bits) are set to 0.

For ease of implementation, we use MIP as the global mobility management protocol. The permanent home IP address is assumed to be the unique identifier for the MN. The MN uses the IP address of its HA in the remote agent address field in its location update message. Like [18], we have provided timestamp-based replay protection in the location update process, with two distinct timestamps for the local (MN-MA) and global (MN-HA) registrations. Similarly, the security association between the HA and MN is distinct from the security association between the MN and MA; currently the only authentication method supported is keyed-MD5.

TESTBED SETUP

Figure 8 shows our experimental network testbed used for evaluating IDMP. We considered a single MN served by its HA (Durga = 192.4.20.44) in its home network (10.10.5.0), with home IP address 10.10.5.10. The home interface address of Durga is 10.10.5.1. Two MAs, MA_1 (Lakshmi = 192.4.20.43) and MA_2 (Saraswati = 192.4.20.45), are connected to routers serving

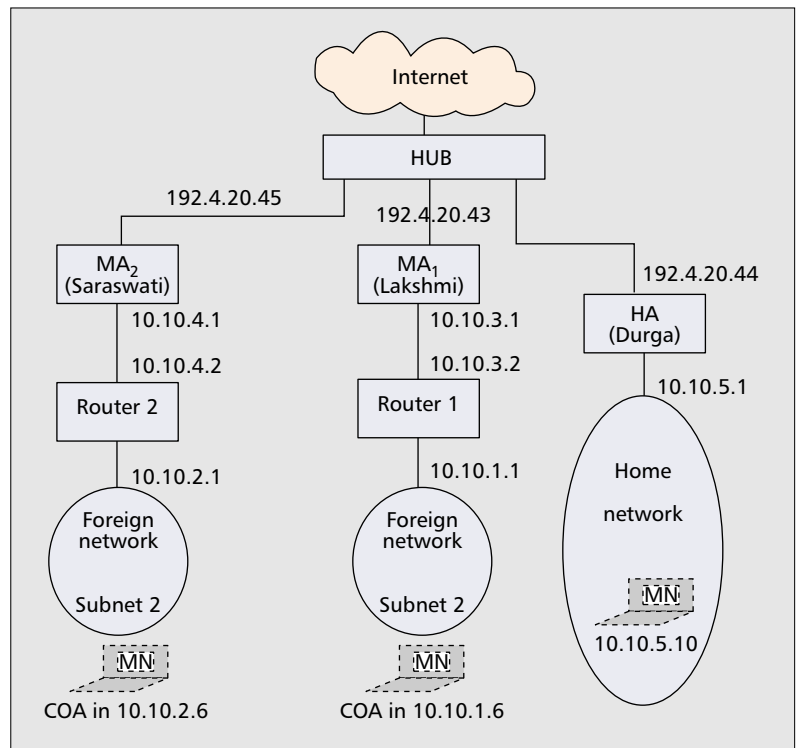


Figure 8. Test network configuration.

subnets 10.10.1.0 and 10.10.2.0, respectively. We assume that our mobility domain comprises both subnets 10.10.1.0 and 10.10.2.0. Accordingly, both Lakshmi and Saraswati can serve as MAs for our MN as long as it stays within this domain.

As the MN enters into the subnet 10.10.1.0, it receives a locally scoped collocated address 10.10.1.6 and the IP address of MA_1 (192.4.20.43) as its GCoA. The MN accordingly first informs MA_1 of its LCoA (10.10.1.6) and subsequently registers with the HA using 192.4.20.43 as its CoA. Afterward, the MN roams into the subnet 10.10.2.0 and gets a new LCoA, 10.10.2.6. Since MA_1 is still its MA, the MN simply performs an intradomain location update, informing MA_1 of its new LCoA.

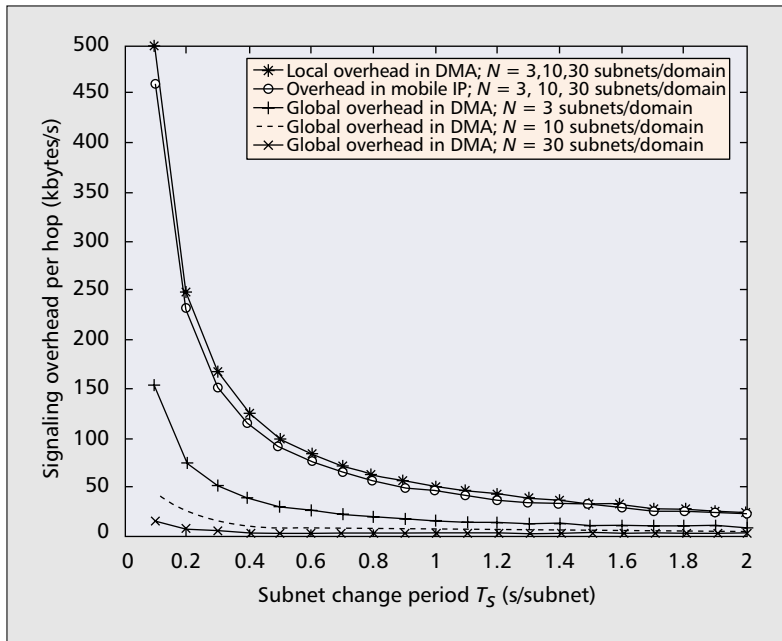
To test the case of interdomain (global) mobility, we subsequently configured the DRCP server to provide a new MA address, say MA_2 (Saraswati = 192.4.20.45), to the MN. In this case, the MN performs both the intradomain and interdomain registrations.

ANALYSIS OF SIGNALING OVERHEAD

In this section we compare the signaling overhead associated with DMA architecture with that of base MIP.

We use the following parameters to express the signaling overhead of both DMA and MIP:


- $L_g = 46$: Size of global registration packet (in bytes)
- $L_l = 50$: Size of local registration packet (in bytes) (Note that $L_g < L_l$, since the global registration request does not contain the LCoA field.)
- T_s : Average duration for which MN remains in a subnet (seconds per subnet)



■ Figure 9. Global and local signaling overhead in DMA.

- T_d : Average duration for which MN remains in a domain (seconds per domain)
- N : Average number of subnets in a domain
- $N_{MA} = 2$: Average number of hops from MN to MA when the MN is in a foreign network
- $N_{HA} = 5$: Average number of hops from MN to HA when the MN is in a foreign network (2 and 5 are arbitrary numbers)

Clearly, T_s and T_d depend on the network topology and mobility pattern of the MN. For the sake of simplicity, in our analysis we assume $T_d = N * T_s$. Table 1 displays the expressions for signaling overhead in basic MIP and DMA architecture in terms of the parameters listed above. In each expression, the factor of 2 is due to the fact that each registration attempt involves



■ Figure 10. Total network signaling overhead.

exchange of a registration request and a corresponding reply message.

The global and local signaling overhead per hop in DMA architecture against T_s for different values of N (3, 10, and 30) are plotted in Fig. 9. These numbers are chosen arbitrarily in order to reflect the increase in number of subnets per domain. As expected, global signaling overhead in DMA architecture is significantly less than the corresponding local overhead. Also the signaling overhead goes down as the MN stays longer in a subnet (and domain). As the number of subnets in a domain increases, the global signaling overhead reduces, whereas the local signaling overhead remains unchanged. In other words, global signaling overhead in basic MIP and local overhead in DMA does not depend on N .

Since global signaling messages travel over a larger number of hops (and hence consume a larger portion of network resources), we would also like to compare DMA and Mobile IP in terms of the total network capacity (aggregated over all hops) used, as shown in Fig. 10. From these plots it is clear that DMA results in a significant reduction in the network signaling overhead, especially when mobiles change subnets more frequently and when a larger number of subnets form a single domain. As N_{HA} increases, the reduction in signaling overhead in DMA becomes more significant. For example, if we use DMA in a 30 subnets/domain network instead of a 3 subnets/domain network, the percentage gain in terms of signaling overhead will be approximately 14 keeping the subnet mobility rate constant.

CONCLUSION

In this article we discuss the design of a lightweight and modular intradomain mobility management protocol and a prototype implementation. Along with the base protocol design, we present two important enhancements for fast handoffs and paging that are necessary for supporting delay-sensitive applications and power-constrained devices, respectively.

The implementation of the protocol elements are based on modifications to Stanford University's MosquitoNet Project Linux code. We demonstrated the basic operation of IDMP in our testbed and present preliminary experimental results. We also used standard packet formats to quantitatively compare DMA's signaling overhead with that of MIP. However, we realize the need for a more comprehensive comparison of IDMP with other existing protocols, especially from the standpoint of handoff latency for real-time traffic, such as voice and video. We expect to shortly incorporate fast handoffs and paging features in our implementation and study their performance in our test-bed in greater detail.

At this moment, IDMP borrows the ideas of replay protection and security associations from Mobile IP. Although it currently appears that IDMP has the same security considerations as Mobile IP, we need to investigate the security aspects further. For example, mobility agents may need additional authentication and authorization functions when a mobile node first registers in a domain. Moreover, the security architecture will require the distribution of transient session keys

(shared secret between the domain and the MN) by the MA to the relevant SAs. With these session keys, SAs can accept new registrations without needing to verify with the MA. Since such functions are not part of the base protocol, we may need to enhance the IDMP message formats with additional fields in the future.

ACKNOWLEDGMENTS

The work of S. K. Das at UTA is partially supported by NSF grants EIA-0086260, EIA-0115885, and ITR IIS-0121297. Ashutosh Dutta would also like to acknowledge Henning Schulzrinne of Columbia University for helpful discussion on some multicast issues.

REFERENCES

- [1] H. Schulzrinne *et al.*, "RTP: A Transport Protocol for Real-Time Applications," RFC 1889, Jan. 1996.
- [2] C. Perkins, Ed., "IP Mobility Support for IPv4," RFC 3320, Jan. 2002.
- [3] C. Perkins and D. Johnson, "Route Optimization in Mobile IP," draft-ietf-mobileip-optim-11.txt, IETF, Sept. 2001, work in progress.
- [4] E. Gustafsson, A. Jonsson, and C. Perkins, "Mobile IP Regional Registration," draft-ietf-mobileip-reg-tunnel-06.txt, Mar. 2002, work in progress.
- [5] D. Johnson, C. Perkins, and J. Arko, "Mobility Support in IPv6," draft-ietf-mobileip-ipv6-17.txt, May 2002, work in progress.
- [6] H. Soliman *et al.*, "Hierarchical MIPv6 Mobility Management," draft-soliman-mobileip-hmipv6-05.txt, IETF, July 2001, work in progress.
- [7] R. Ramjee *et al.*, "IP Micro-mobility Support Using HAWAII," draft-ietf-mobileip-hawaii-01.txt, July 2000, work in progress.
- [8] A. Campbell *et al.*, "Cellular IP," draft-ietf-mobileip-cellularip-00.txt, IETF, Jan. 2000, work in progress.
- [9] C. Tan, S. Pink, and K. Lye, "A Fast Handoff Scheme for Wireless Networks," *Proc. 2nd ACM Int'l. Wksp. Wireless Mobile Multimedia (WOWMOM), ACM SIGMOBILE*, Aug. 1999, pp. 83–90.
- [10] K. El Malki, "Low Latency Handoffs in Mobile IPv4," draft-ietf-mobileip-lowlatency-handoffs-v4-03.txt, Nov. 2001, work in progress.
- [11] K. El Malki and H. Soliman, "Hierarchical Mobile IP v4/v6 and fast handoffs," draft-elmalki-soliman-hmipv4v6-00.txt, Mar. 2000, work in progress.
- [12] A. Misra *et al.*, "IDMP-based Fast Handoffs and Paging in IP-based Cellular Networks," *Proc. Int'l. Conf. 3GWireless and Beyond*, May/June 2001, pp. 427–32.
- [13] E. Schooler *et al.*, "SIP: Session Initiation Protocol," RFC 2543, IETF, Mar. 1999.
- [14] E. Wedlund and H. Schulzrinne, "Mobility Support Using SIP," *Proc. 2nd ACM Int'l. Wksp. Wireless Mobile Multimedia, ACM/IEEE*, Aug. 1999, pp. 76–82.
- [15] S. Das *et al.*, "TeleMIP: Telecommunication Enhanced Mobile IP Architecture for Fast Intra-Domain Mobility," *IEEE Pers. Commun.*, vol. 7, no. 4, Aug. 2000, pp. 50–58.
- [16] A. Misra *et al.*, "IDMP: An Intra-Domain Mobility Management Protocol using Mobility Agents," draft-mobileip-misra-idmp-00.txt, July 2000, work in progress.
- [17] A. Misra *et al.*, "Autoconfiguration, Registration and Mobility Management for Pervasive Computing," *IEEE Pers. Commun.*, vol. 8, no. 4, Aug. 2001, pp. 24–31.
- [18] R. Droms, "Dynamic Host Configuration Protocol," RFC 2131, IETF, Mar. 1997.
- [19] A. McAuley, S. Das, S. Baba and Y. Shobatake, "Dynamic Registration and Configuration Protocol (DRCP)," draft-itsumo-drcp-01.txt, July 2000, work in progress.
- [20] AT&T, "AT&T sets the standard for network reliability," <http://www.att.com/network/standrd.html>, Oct. 2000.
- [21] MosquitoNet Project, "Mobile IP resource," <http://gunpowder.stanford.edu/mip/resource.html>, tech. rep., 1998.
- [22] K. Chakraborty, *et al.*, "Implementation and Performance Evaluation of TeleMIP," *Proc. IEEE ICC*, Helsinki, Finland, vol. 8, June 2001, pp. 2588–92.

BIOGRAPHIES

SUBIR DAS [M] (subir@research.telcordia.com) received his M.Tech. degree from the Institute of Radio Physics and Electronics, Calcutta University in 1990, and Ph.D. degree in 1996 from Indian Institute of Technology, Kharagpur. Since

Architecture	Signaling overhead (bytes/sec)		
	Local per hop	Global per hop	Total in network
Mobile IP	0	$2L_g/T_s$	$2N_{HAL}L_g/T_s$
MA	$2L_l/T_s$	$2L_g/T_d$	$2N_{HAL}L_g/T_d + 2N_{MAL}L_l/T_s$

■ Table 1. Expressions for signaling overhead.

1999 he has been at Telcordia Technologies Inc. and is currently a research scientist in the Wireless IP Research Lab. During 1997–1999 he was a faculty member in the Electronics and Electrical Engineering Department, Indian Institute of Technology, Kharagpur. He has designed several protocols and architecture for next-generation wireless networks particularly in the areas of autoconfiguration, registration, and mobility management. He has also developed several prototypes for MIP and SIP-based seamless connectivity in wireless LAN-based networks.

ANTHONY MCAULEY (mcauley@research.telcordia.com) received a Ph.D. from Hull University, England in 1985. He was a research fellow at Caltech, 1985–1987. Since 1987 he has been at Telcordia and is currently a director in the Wireless IP Research group. He has worked on all aspects of TCP/IP, particularly as they relate to wireless and dynamic environments. He wrote the proposals that won funding for these projects that are being applied to both commercial and military research projects. He has built several mobile internetworking prototypes on Linux used in experiments over 802.11, cellular, and satellite networks. These prototypes demonstrated software he has written for protocol boosters, multicast proxies, Mobile IP with DHCP, and a novel Transport Protocol (TP++). He has also worked on IP multicasting in ad-hoc and large-scale networks (including the Comprehensive Test Bed Treaty network), efficient error correction and detection codes, and design of VLSI chips (for everything from microprocessors to asynchronous packet switches).

ASHUTOSH DUTTA [M] (adutta@research.telcordia.com) is currently a research scientist in Telcordia Technologies' Advanced Networks Systems Research Laboratory. He has a Bachelor's degree in electrical engineering from R.E.C Rourkela, India, a Master's degree in computer science from New Jersey Institute of Technology, and a professional degree in electrical engineering from Columbia University, New York. He has been working in computer systems and networking for the last 15 years. He has dealt with a variety of high-speed networks, and has been responsible for implementing many systems and networking related projects. Prior to joining Telcordia he was director of the Central Research Facilities in Columbia University's Computer Science Department for eight years.

ARCHAN MISRA (archan@us.ibm.com) is currently a research staff member with the Pervasive Security and Networking Department of IBM Research, Hawthorne, New York, where he investigates protocols and architectures for future wireless Internet services. As part of his previous work at Telcordia Technologies he designed protocols, such as IDMP and DCDP, for IP-based mobility management and auto-configuration of dynamic networks. He is presently working on middleware for secure and auto-configuring location-aware services. He received his Ph.D. in electrical and computer engineering from the University of Maryland at College Park in May 2000.

KAUSHIK CHAKRABORTY (kauchaks@glue.umd.edu) received a B.Tech. in electrical and communications engineering from the Indian Institute of Technology, Kharagpur, India in 1998 and an M.S. in electrical and computer engineering from the University of Maryland at College Park in 2000, where he is currently working toward a Ph.D.

SAJAL K. DAS [M] (das@cse.uta.edu) received his B.Tech. degree in 1983 from Calcutta University, M.S. degree in 1984 from Indian Institute of Science, Bangalore, and Ph.D. degree in 1988 from the University of Central Florida, Orlando, all in computer science. Currently he is a full professor of computer science and engineering and also the founding director of the Center for Research in Wireless Mobility and Networking (CREWMan) at the University of Texas at Arlington (UTA). Prior to 1999 he was a professor of Computer Science at the University of North Texas (UNT).