

IDMP-based FAST HANDOFFS AND PAGING IN IP-BASED CELLULAR NETWORKS

Archan Misra, Subir Das, Ashutosh Dutta, Anthony Mcauley
{archan,subir,adutta,mcauley}@research.telcordia.com
Telcordia Technologies,
445 South Street, Morristown, NJ 07960, USA.

Sajal K Das
Center for Research in Wireless Mobility and Networking (CreWMaN)
Department of Computer Science and Engineering
The University of Texas at Arlington, P.O. Box 19015, Arlington, TX 76019-0015, USA.

Abstract—In this paper, we consider the use of the recently proposed Intra-Domain Mobility Management Protocol (IDMP) in 3rd and 4th generation (3/4G) wireless cellular networks to reduce the latency of intra-domain location updates and the mobility signaling traffic. We first present enhancements to basic IDMP that provide fast intra-domain handoffs by using a duration-limited, proactive packet ‘multicasting’ scheme. We quantify the expected buffering requirements of our proposed multicasting scheme for typical 3/4G network characteristics and compare it with alternative IP-based fast handoff solutions. We also present a paging scheme under IDMP that replicates the current cellular paging structure. Our paging mechanism supports generic paging strategies and can significantly reduce the mobility-related IP signaling load.

I. INTRODUCTION

The Intra-Domain Mobility Management Protocol (IDMP) [1] uses a two-level hierarchy to manage node mobility in future IP-based cellular networks. By aggregating multiple subnets into a mobility domain, IDMP localizes the scope of most location update messages and drastically reduces both the global signaling load and the update latency. IDMP is conceptually a two-level generalization of the Mobile IP architecture, with a special node called the Mobility Agent (MA) providing a mobile node (MN) a domain-wide stable point of packet redirection. Depending on the mobility-related requirements of a specific application, IDMP can be combined with multiple global binding protocols. For example, the TeleMIP architecture [2] combines IDMP with Mobile IP [3] (MIP) to provide seamless packet redirection at the network layer for TCP-based applications. Alternatively, we have shown [4] how SIP-based application-layer global mobility management can be combined with IDMP to define a scalable and flexible mobility solution for Voice-over-IP (VoIP) traffic.

In this paper, we consider two extensions to the base IDMP specifications, and analyze their applicability to future IP-based cellular network architectures (see figure 1), where the base stations (BS) are IP-enabled and link layer-specific functions are confined only to the wireless interface between the MN and the Base Station (BS). In such a scenario, we assume that the Subnet Agent (SA), a specialized IDMP node that provides subnet-

specific support to the MN, is co-located with the BS. Generalizations to alternative architectures, where IP functionality does not extend into the RAN, are easy to make.

We first present IDMP’s optional *fast handoff* mechanism, which provides an IP-layer solution to reduce the service interruption during an inter-BS handoff. The mechanism allows the use of either MN-initiated or BS-initiated handoff triggers and is applicable to multiple future link-layer technologies. Such triggers notify the MA of an impending handoff, whereupon the MA proactively multicasts in-bound packets (destined to the MN) to the set of neighboring SAs. By caching these packets for a specified duration, the SA can minimize the loss of in-flight packets and forward such packets to the MN immediately after the MN refreshes its IP configuration parameters at the new subnet. IDMP’s IP-based fast handoff technique provides a secure fast handoff solution that does not assume the existence of specific layer-2 authentication functions or require the adjacent BSs to be aware of each other’s identity.

We then describe IDMP’s IP-layer paging mechanism, which allows an idle MN to be located even though it does not perform IP-layer registration/configuration at every change in subnet (BS). By performing the essential paging functions at the IP layer, we can make the mechanism relatively independent of the radio technology.

A. Previous and Related Work

Mobile IP (MIP) [3], the standard approach to IP-based mobility management, was designed primarily to provide transparent packet redirection to non-real time TCP applications running in conventional network hosts. Accordingly, for cellular environments with a large number of MNs and real-time VoIP traffic, MIP 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 non-hierarchical MIP solutions, such as MIP-RO [5] and MIPv6 [6]. SIP-based mobility mechanisms [7], [8] 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 send-

ing 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, as it allows a CN to send traffic directly to the MN's co-located CoA (without tunneling), and as it permits the application to control the characteristics of an ongoing session at the MN changes subnets.

IDMP is one of several proposed hierarchical mobility management solutions. All such schemes localize the signaling on intra-domain movement to nodes within the domain. One approach to intra-domain mobility management is the route-modification approach, characterized by Cellular IP (CIP) [9] and HAWAII [10]: 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 approach 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) [12] 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 intra-domain movement. Similarly, Hierarchical MIPv6 (HMIPv6) [11] introduces an agent called the MAP to localize the management of intra-domain mobility. Comparisons with alternative fast handoff and paging proposals will be discussed in the relevant later sections.

A.1 IDMP Overview

IDMP is also a multi-CoA intra-domain mobility solution. However, unlike HAWAII, MIP-RR or HMIPv6, IDMP is designed as a stand-alone solution for intra-domain mobility 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) is similar to a MIP FA and provides subnet-specific mobility services. Under IDMP, an MN obtains two concurrent CoAs:

- *Local Care-of Address (LCoA)*: This is similar to MIP's CoA in that it identifies the MN's present subnet of attachment. 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 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 figure 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 [1], [2].

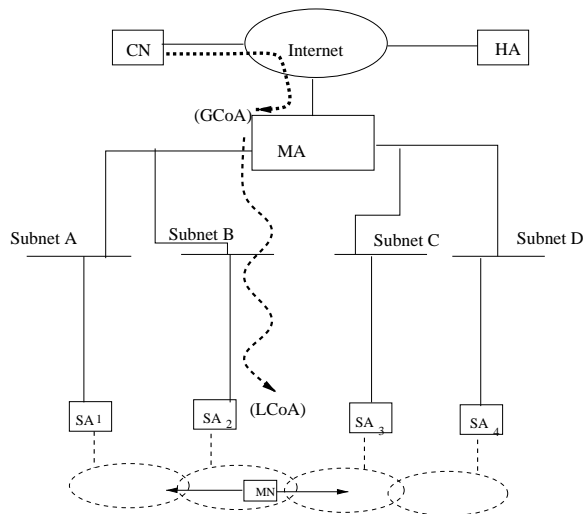


Figure 1: IDMP Logical Elements & Architecture

II. FAST HANDOFF SCHEME IN IDMP

Under basic IDMP, the handoff delay equals the time taken for the MA to become aware of the MN's new point of attachment (LCoA). In an IPBS (IP-based Base Station) architecture, this delay 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 operations such as slot-specification in TDMA or code synchronization in CDMA.
- *IP Subnet Configuration (Δ_2)*: An MN must use IP-layer configuration protocols to obtain the new LCoA. If IDMP's SA mode is used, then the MN must obtain an Agent Advertisement beacon and then request a new LCoA. The SA will then respond with an Acknowledgement message. If the co-located 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 Intra-domain location update message. The MA will redirect packets to the MN's new LCoA only after receiving this message.

Δ_1 , while a link-layer specific parameter, can be expected to be quite low. For example, in CDMA-based soft handoffs, Δ_1 is effectively 0, since, in a well-designed network, communication with the old BS is not discontinued until the connection with the new BS is firmly established. 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.

A. The Fast Handoff Procedure

IDMP's fast handoff procedure is based on the assumption that a layer-2 trigger will be available (either to the MN or to the old BS) indicating an imminent change in connectivity. We explain the fast handoff mechanism using figure 2, which shows an MN moving from SA_2 to SA_3 . To minimize the 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 per-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) with SA_3 , SA_3 can immediately forward all such buffered packets over the wireless interface, without waiting for the MA to receive the corresponding Intra-domain Location Update.

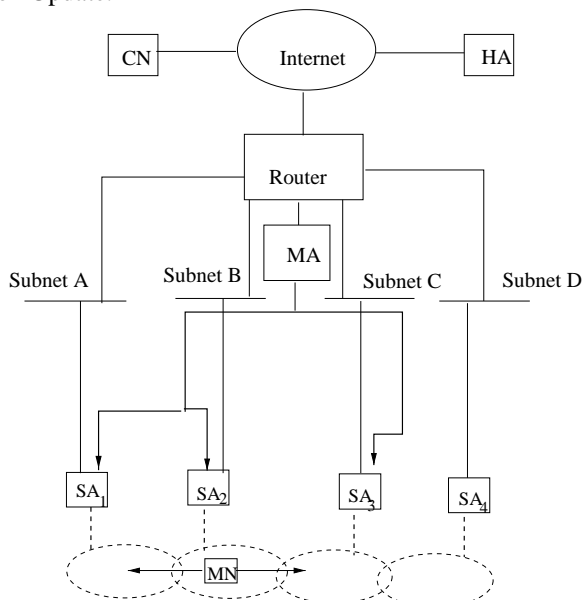


Figure 2: IDMP Fast Handoff

Several features of this proposal make it attractive for future IP-based networks.

- Unlike other fast handoff proposals, IDMP's *MovementImminent* message does not specify the IP address of the target (new) BS in this message. This keeps the message size short; in fact, the MN can piggyback such a message simply by setting a bit in frames used for existing link-layer signaling. Moreover, by allowing the old BS to generate this message for the MA, we also accommodate some possible (hard) handoff scenarios, where the MN loses connectivity with the old SA before establishing radio-connectivity with the new BS. In such a scenario, the *MovementImminent* message is forwarded in parallel to the establishment of the new link. Accordingly, multicast forwarding is likely to be invoked concurrently with the subnet-level IP configuration phase, thereby reducing (if not completely eliminating) the Δ_3 component of the delay.

- IDMP utilizes a *network-controlled* (network or mobile-initiated) handoff technique. The MA (effectively the network element) decides the set of target BSs to which in-flight packets are multicast. This is especially useful in scenarios where the MN may be in contact with multiple BSs and is unable to exactly specify the identity of its exact point of attachment. While current cellular networks use a network-controlled handoff technique (where the BSC determines the candidate BS based on link-layer measurements supplied by the MN), the IP mobility model is typically MN-driven, with the MN selecting an FA from a list announced via agent advertisements. IDMP preserves the network-controlled handoff model for future IP-based cellular networks, without compromising the MN's ability to select such fast handoff support.
- Unlike proposals such as [13], the BS in our solution does not transmit all arriving multicast packets over the wireless interface. Such packets are only temporarily buffered in per-MN buffers. Any such proactively multicast packet is forwarded to the MN over the wireless interface by the new BS alone, thus preventing unnecessary wastage of wireless bandwidth.
- IDMPs' fast handoff scheme does not eliminate Δ_2 from the handoff delay; it merely delays the transmission of packets arriving during this interval. By buffering packets during this transient, we are however able to avoid the loss of in-flight packets. VoIP receivers are typically able to tolerate variation in the per-packet delay, as long as the packets are not actually lost. It is thus acceptable to suffer longer delay (due to Δ_2), as long as we can minimize the handoff-related packet loss.

B. Implementing Fast Handoff

For a prototype implementation, we use IP multicast to proactively distribute such packets to possible points of attachment. IDMP requires only one multicast group per neighbor set; all the BSs that are neighbors of a specific BS are members of this multicast group. Since a single BS can be a neighbor of multiple BSs, each BS can indeed be a member of multiple multicast groups. Unlike other multicast-based forwarding schemes [13], [14], the IDMP approach does not require the establishment of separate multicast groups for individual MNs. Also, the group membership is not dynamic; given a fixed network topology, the set of neighboring BSs stays constant. Each BS is thus permanently subscribed to one or more multicast groups, each of which always has a well-defined distribution tree. Accordingly, the fast handoff scheme does not require a BS to dynamically join or leave a group, and hence, does not suffer from any transient tree-establishment latencies. Standard protocols, such as PIM [15] or DVMRP can be used to establish the multicast tree.

On receiving a *MovementImminent* message, the MA encapsulates an in-flight packet and then tunnels it to the appropriate multicast address. (For such multicast forwarding, the MA does not perform the conventional tunneling towards the current LCoA). On receiving such a tunneled multicast packet, each SA will first decapsulate the outer-most header. It then buffers the decapsulated packet in a per-user buffer, using the destination address in the inner-header (which is unique to a specific MN)

as an index. When an MN subsequently obtains its new subnet-specific configuration parameters from the new SA (say SA_3 in figure 2), that SA can then forward any cached packets to the MN before the intra-domain location update process is complete.

Simple calculations with typical 3G data rates indicate that even a small user buffer is effective in reducing the loss of in-flight packets. For example, if the intra-domain update latency (L) is 200 msec, and the incoming traffic rate (R) is 144 Kbps, then a buffer size of ($L \cdot R$) 3.6 KBytes is able to protect against buffer overflow due to multicast packets transmitted during the handoff transient.

C. Alternative Fast Handoff Suggestions

Two alternative schemes for providing fast handoff, within the MIP context, have also been recently proposed. Under the pro-active handoff proposal [16], layer-2 triggering causes the old FA to effectively establish a transient tunnel to the new FA. This mechanism assumes that layer-2 mechanisms always provide the old FA with the IP address of the new FA. This proposal also assumes that layer-2 authentication mechanisms are adequate to (temporarily) authenticate the MN at the new FA; consequently, all arriving packets are forwarded to the MN by the new FA as soon as the link-layer connectivity (Δ_1) is established. While this approach makes the handoff transient smaller than IDMP (since it eliminates Δ_2 , the delay for subnet-specific IP configuration at the new BS), the mechanism makes assumptions about the authentication and signaling capabilities at the radio layer. In contrast, IDMP does not forward any packets until the IP-level configuration at the new subnet is complete, and hence, does not assume the presence of link-layer security mechanisms.

The alternative proposal [17] implements fast handoff by having the MN initiate a new MIP registration with the new FA through the old FA. The old FA thus tunnels such a MIP Registration request to the new FA. This proposal eliminates the sequential delay due to $\Delta_1 + \Delta_2$, since the new MIP registration (via the old BS) can occur concurrently with the establishment of radio connectivity at the new FA (Δ_1). However, in contrast to IDMP, this approach does not provide the network any control over the handoff process, since the MN unilaterally decides the identity of the new BS. Moreover, both [16], [17] require the MN to inform the old BS of the identity of the new BS; while current BSC-controlled handoff schemes do provide such information, IDMP does not require transmission of such information.

III. PAGING SUPPORT IN IDMP

While IDMP's use of multicasting for fast handoffs minimizes the loss of in-flight packets during an intra-domain handoff, it does not reduce the frequency of intra-domain location updates. In the absence of paging support, an MN must obtain a local care-of address and re-register 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.

A. Paging Operation for Idle Hosts

To motivate IDMP's paging solution, note that the 'multicasting' scheme described for fast handoff support in section II inherently sends multiple copies of the same data to multiple FAs/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 BSs) are grouped into Paging Areas (PA) identified by 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., FA Agent Advertisements). In fact, such IP-layer advertisements may optionally be combined with link-layer beacons.

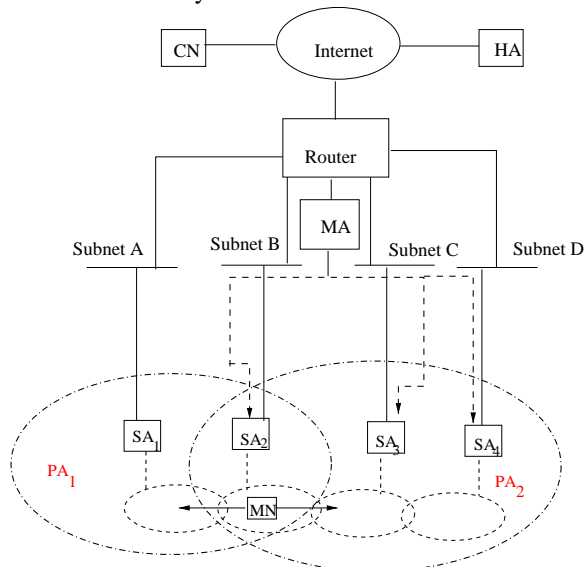


Figure 3: IDMP Paging Mechanism

IDMP's paging scheme is visually illustrated in figure 3. In this model of operation, Subnets B, C and D belong to the same PA, while subnet A is part of a different PA. 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, it does not even 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 PA.

When the MA receives packets for an MN which is currently registered, but which does not have a valid LCoA assigned, it 'multicasts' a *PageSolicitation* packet to all the subnets associated with the MN's current PA (to SA_2 , SA_3 and SA_4) and buffers the incoming packet. When the MN re-registers with the MA, buffered packets are forwarded to the MN. We assume that temporary buffering is acceptable as the intra-domain 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, call setup delays are typically around 2.5 sec [18]; accordingly the paging latency is expected to fall within the targeted bounds.

B. Paging Implementation

Each PA is then identified by a unique domain-specific multicast address; an SA belonging to a specific PA must permanently subscribe to the corresponding multicast group. Note also that, similar to the overlapping RA concept in current cellular networks, an SA can subscribe to multiple multicast groups and hence, be associated with multiple PAs (figure 3).

The base IDMP specification needs minor modifications for supporting paging. An MN must now actively inform its MA when it switches from the active to the idle state, thereby activating the paging functionality at the MA. In the absence of active “idle state notification”, the MN would move to neighboring subnets without performing the subnet-level reconfiguration, while the MA would continue to (mistakenly) unicast arriving packets to the MN’s last registered local care-of address. Moreover, when an MN changes its LA in idle state, it first performs a local re-configuration to obtain a new local care-of address and then informs the MA of its new LAI.

C. Comparison with Alternative IP Paging Schemes

IDMP’s paging mechanism differs from alternative paging proposals presented in CIP and HAWAII in the following manner:

- IDMP does not need intermediate nodes to cooperate for paging support; hence, upgrades of intermediate routers are not needed. Also, only specialized nodes (MA and SAs) take part in the paging process; *this localizes the nodes where future upgrades would need to be installed*. Moreover, paging operations are distributed among different MAs (each serving different MNs); there is thus no single point of failure for the entire domain.
- The scope of a PA in CIP is determined by the location of the Paging Cache closest to the wireless edge. If this cache is higher in the hierarchy than the LCA (crossover router or least common ancestor), the ‘paging packet’ (actually a regular data packet) is broadcast over a larger area (wasting resources); if the nearest paging cache is located at a lower level than the LCA, the ‘paging packet’ is then broadcast only over a subset of the PA. Such a placement-based PA definition hinders the co-existence of overlapping PAs of arbitrary size; changing PAs also requires explicit manipulation of paging cache locations inside the cellular domain. In contrast, IDMP follows the existing Registration Area-based cellular paging architecture and permits the co-existence of arbitrarily sized PAs. The size of a PA can be changed by simply altering the subscription of FAs/subnets to well-known multicast groups.

IV. CONCLUSIONS AND FUTURE WORK

In this paper, we presented two enhancements to the IDMP solution for IP-based hierarchical mobility management. We considered an IP-based BS architecture, where radio-specific management functions terminate at the wireless interface and motivate the need for both fast handoff and paging solutions.

To minimize packet loss during intra-domain handoffs, we presented a *time-bound* localized ‘multicasting’ approach. By proactively informing its associated Mobility Agent (MA) of an impending change, an MN enables the MA to multicast packets for a limited duration to a set of *neighboring* subnets. Specific nodes on those subnets (SAs/ designated routers) buffer such multicast packets for a short while; if the MN enters its subnet, such a node is able to immediately forward these packets to the mobile, significantly eliminating packet loss and delays. Our approach is consistent with a mobile-initiated, network-controlled handoff scheme and reduces the handoff delay to the latency incurred in performing a new IP-layer registration at the new BS. While this latency is higher than other schemes that assume some layer-2 coordination, our temporary buffering mechanism appears to provide acceptable latency variation for most applications of interest.

We also extended this localized ‘multicasting’ idea to provide paging support under IDMP. In our approach, each subnet would be associated with one or more Paging Areas (PA). A non-active MN would perform intra-domain location updates only when it changes its PA; to determine the MN’s exact location within its current PA, the MA would ‘multicast’ a paging packet to all subnets to this PA. Unlike other suggested IP-based paging schemes, our mechanism does not assume a tree-like topology and allows easy configuration of variable-size PAs.

We currently have a Linux-based implementation of IDMP deployed in our testbed. The Linux Mobile IP code of Stanford University MosquitoNet project [19] is used as a basis for IDMP implementation. The mobility agent daemon of IDMP is a modified version of the home agent daemon, while the mobile host daemon has been upgraded to support IDMP. Additional details of the implementation, as well as preliminary performance results, are available in [20].

REFERENCES

- [1] A. Misra, S. Das, A. McAuley, A. Dutta and S. K. Das, “IDMP: An Intra-Domain Mobility Management Protocol using Mobility Agents”, `draft-mobileip-misra-idmp-00.txt`, July 2000, Work in Progress.
- [2] S. Das, A. Misra, P. Agrawal and S. K. Das, “TeleMIP: Telecommunication Enhanced Mobile IP Architecture for Fast Intra-Domain Mobility”, *IEEE PCS Magazine*, August 2000.
- [3] C. Perkins, “IP Mobility Support for IPv4, revised”, `draft-ietf-mobileip-rfc2002-bis-02.txt`, IETF, July 2000.
- [4] A. Misra, S. Das and A. McAuley, “Hierarchical Mobility Management for VoIP Traffic”, under submission.
- [5] C. Perkins and D. Johnson, “Route Optimization in Mobile IP”, `draft-ietf-mobileip-optim-09.txt`, IETF, Work in Progress, February 2000, Work in Progress.
- [6] D. Johnson and C. Perkins, “Mobility support in IPv6”, `draft-ietf-mobileip-ipv6-13.txt`, IETF, September 2000, Work in Progress.
- [7] E. Schooler, M. Handley, H. Schulzrinne and J. Rosenberg, “SIP: Session initiation protocol,” RFC 2543, IETF, March 1999.
- [8] E. Wedlund and H. Schulzrinne, “Mobility support using SIP,” *Proceedings of Second ACM International Workshop on Wireless Mobile Multimedia, ACM/IEEE*, August 1999, pp. 76–82.
- [9] A. Campbell, J. Gomez, C-Y. Wan, S. Kim, Z. Turanyi and A. Valko, “Cellular IP,” `draft-ietf-mobileip-cellularip-00.txt`, IETF, January 2000, Work in Progress.
- [10] R. Ramjee, T. La Porta, S. Thuel, K. Varadhan and L. Salgarelli, “IP micro-mobility support using HAWAII,” `draft-ietf-mobileip-hawaii-01.txt`, July 2000, Work in Progress.
- [11] H. Soliman, C. Castelluccia, K. El Malki and L. Bellier, “Hierarchical MIPv6 Mobility Management”, `draft-soliman-mobileip-hmipv6-01.txt`, IETF, September 2000, Work in Progress.
- [12] E. Gustafsson, A. Jonsson and C. Perkins, “Mobile IP Regional Tunnel

- Management”, draft-ietf-mobileip-reg-tunnel-03.txt, July 2000, Work in Progress.
- [13] C. Tan, S. Pink and K. Lye, “A fast handoff scheme for wireless networks”, Proc. of Second ACM International Workshop on Wireless Mobile Multimedia (WOWMOM), ACM SIGMOBILE, August 1999, pp. 83-90.
 - [14] S. Seshan, H. Balakrishnan and R. H. Katz, “Handoffs in cellular wireless networks: The daedalus implementation and experience”, Kluwer International Journal on Wireless Personal Communications, January, 1997.
 - [15] D. Estrin, et al, “Protocol independent multicast-sparse mode (PIM-SM): protocol specification”, RFC 2117, IETF, June 1998.
 - [16] P. Calhoun, et al, “Foreign Agent Assisted Handoff”, draft-calhoun-mobileip-proactive-fa-02.txt, September 2000, Work in Progress.
 - [17] K. ElMalki and H. Soliman, “Hierarchical Mobile IP v4/v6 and Fast Handoffs”, draft-elamalki-soliman-hmipv4v6-00.txt, March 2000, Work in Progress.
 - [18] AT&T, “AT&T sets the standard for network reliability”, <http://www.att.com/network/standrd.html>, October 2000.
 - [19] MosquitoNet Project, “Mobile IP resources.”, <http://gunpowder.Stanford.EDU/mip/resource.html>, Technical report, 1998.
 - [20] K. Chakraborty, et al, “Implementation and Performance Evaluation of TeleMIP”, to appear in Proceedings of IEEE International Communications Conference (ICC), June 2001.