

IDMP-based Fast Handoffs and Paging in IP-based 4G Mobile Networks

Archan Misra

IBM T J Watson Research Center,
19 Skyline Drive, Hawthorne, NY 10533

Subir Das, Ashutosh Dutta, Anthony McAuley
Telcordia Technologies,
445 South Street, Morristown, NJ 07960

Sajal K Das

Center for Research in Wireless Mobility and Networking (CRWMan)
Department of Computer Science and Engineering
The University of Texas at Arlington, Arlington, TX 76019-0015

Abstract—In this paper, we consider the use of our recently proposed Intra-Domain Mobility Management Protocol (IDMP) in 4th generation (4G) mobile networks. On evaluating the heterogeneous access technologies, cellular layouts and application characteristics of 4G environments, we realize a need to reduce both the handoff latency and the frequency of mobility-related signaling. We first present IDMP's fast intra-domain handoff mechanism that uses a duration-limited, proactive packet 'multicasting' solution. We quantify the expected buffering requirements of our proposed multicasting scheme for typical 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

There has recently been an almost universal recognition that Mobile IP [1], the current standard for IP-based mobility management, needs to be enhanced to meet the need for future 4th-generation (4G) cellular environments. In particular, the absence of a location management hierarchy leads to concerns about the signaling scalability and the handoff latency, especially for a future infrastructure that must provide global mobility support to potentially billions of mobile nodes and accommodate the stringent performance bounds associated with real-time multimedia traffic.

In tandem with alternative proposals [2], [3] to develop a hierarchical IP-based mobility management framework in 3rd-generation (3G) cellular networks, we have recently proposed the *Intra-Domain Mobility Management Protocol* (IDMP) [4] for managing node mobility within a specific *domain*. Like most proposals, IDMP envisions that multiple IP-subnets are aggregated into a single domain; as long as the mobile node (MN) moves within a single domain, all its mobility-related signaling remains localized to specialized nodes within that domain. Since the MN changes domains fairly infrequently, this local-

ization drastically reduces both the global signaling load and the update latency. Conceptually, IDMP is a two-level generalization of the Mobile IP architecture, with a special node called the *Mobility Agent* (MA) providing an MN a domain-wide stable point of packet redirection. Based on our observation that the global communications infrastructure is likely to employ a multiplicity of protocols for ensuring ubiquitous global connectivity, we have intentionally designed IDMP to be independent of any specific solution for *global* (inter-domain) mobility management.

A careful consideration of the features of, and operating scenarios envisioned in, 4G networks however shows the need for further improvements and enhancements to IDMP. In this article, we emphasize the following characteristics of the 4G cellular environment and the consequent need for certain additional mobility-related features:

a) Heterogeneous Access Technologies:

The aim of 4G cellular networks is to develop a framework for truly ubiquitous IP-based access by mobile users, with special emphasis on the ability to use a wide variety of wireless and wired access technologies to access the common information infrastructure. While the 3G initiative is almost exclusively directed at defining *wide-area* packet-based cellular technologies, the 4G vision embraces additional *local-area* access technologies, such as IEEE 802.11-based wireless local area networks (WLANs) and Bluetooth-based wireless personal area networks (WPANs). The development of mobile terminals with multiple physical or software-defined interfaces is expected to allow users to seamlessly switch between different access technologies, often with overlapping areas of coverage and with dramatically different cell sizes.

Figure 1 shows one example of this multi-technology vi-

sion at work in a corporate campus located in an urban environment. While conventional wide-area cellular coverage is available in all outdoor locations, the corporation also offers 802.11-based access in public indoor locations such as the cafeteria and parking lots, as well as Bluetooth-based access to the Internet in every individual office. As mobile users drive in to work, their ongoing Voice-over-IP (VoIP) calls are seamlessly switched, first from the wide-area cellular to the wireless LAN infrastructure, and subsequently from the 802.11 Access Point (AP) to the Bluetooth AP located in their individual cubicles or offices. Since a *domain* can comprise multiple access technologies, mobility management protocols should be capable of handling *vertical handoffs*, i.e., handoffs between heterogeneous technologies. Accordingly, we believe that it is preferable to define any additional mobility-related feature at the IP-layer, and not rely upon or assume the existence of specific features from the underlying link layers.

b) *Optional Fast Handoff Support:*

While current proposals for intra-domain mobility management do localize the scope of signaling messages generated during intra-domain movement to nodes within the domain, the resultant latency may still prove to be unacceptably high for certain 4G applications. Indeed, the support of various real-time or high-bandwidth multimedia services, such as VoIP and video conferencing, which have very strict bounds on the handoff delay, is central to the 4G vision.

c) *Paging Support:*

The base Mobile IP specification does not provide any form of paging support. Hence to maintain connectivity with the backbone infrastructure, an MN must generate location updates every time it changes its point of attachment, even if it is currently in dormant or standby mode. Such mobility-related signaling leads to significant wastage of an MN's battery power, especially in pico-cellular environments (such as 802.11-based cellular topologies) where the MN may change its subnet of attachment very frequently (possibly once every 10 seconds). It is thus absolutely essential to define some form of flexible paging support in the intra-domain mobility management scheme.

In this paper, we present the extensions to the base IDMP specifications [4] for fast handoff and paging support, and analyze their applicability to 4G cellular architectures. We assume an all-IP access network where IP-layer functionality extends all the way to the nodes at the wireless edge of the access domain. We refer to this IP-capable point of attachment to the wired infrastructure as the Base Station (BS); in this 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 specifically extend to the BS, are easy to

make.

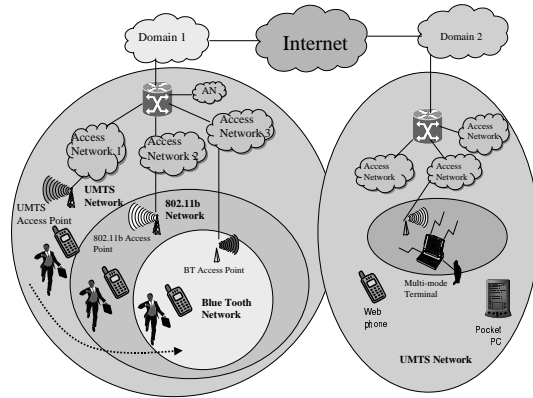


Figure 1: Example of Seamless Vertical Handoff in a 4G Environment

Under basic IDMP operation, an MN encounters a handoff delay equal to the time taken to update its designated MA about its new point of attachment; during this transient, packets may be incorrectly forwarded to the MN's old point of attachment and hence be lost. IDMP's *fast handoff* mechanism provides an IP-layer solution to reduce the duration of this service interruption and minimize the loss of in-flight packets during an inter-BS handoff. The mechanism uses 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 (see Figure 2). 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-layer paging mechanism allows an idle MN to be located even though it does not perform IP-layer registrations 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.

The rest of the paper is organized as follows. In section II, we present a summary of the current proposals for hierarchical IP-based mobility management and give an overview of the basic operation of IDMP. While section III describes IDMP's fast handoff mechanism and its prototype implementation, section IV details the enhancements needed to provide flexible IP-layer paging support. Finally, section V concludes the paper with a brief discussion of our plans for future work.

II. PREVIOUS AND RELATED WORK

Mobile IP (or MIP) [1], the current standard for IP-based mobility management, was designed primarily to provide transparent packet redirection to non-real time TCP applications running in conventional network hosts. Mobile IP works by assigning each mobile node a temporary care-of address (CoA),

in CDMA) with the new BS/SA.

- *IP Subnet Registration* (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 Acknowledgment message. If the co-located mode is used, the MN must exchange DHCP configuration messages with the DHCP Server before obtaining a valid LCoA.
- *Intra-domain Update Delay* (Δ_3): The MN must finally inform the MA of this new LCoA via an Intra-domain location update message. Upon reception of this update, the MA will redirect packets to the MN’s new LCoA.

To gain an idea of the handoff latency incurred by current IP-based intra-domain mobility protocols, we performed experiments with our implementation of IDMP, and the publicly available implementation of Cellular IP, on our 802.11-based testbed. Figure 3 shows the number of consecutive packets lost during a handoff as the packet inter-arrival time is varied. We can see that the handoff delay with IDMP in this scenario (where the MA was located one hop away from the SAs) is approximately around 12-15 msec.

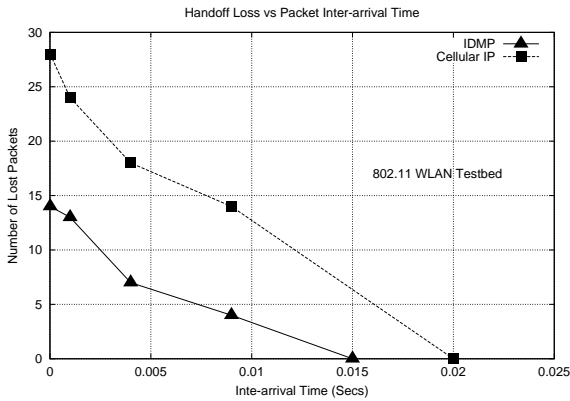


Figure 3: Packet Loss in IDMP and CIP Handoff

IDMP’s fast handoff mechanism is designed to eliminate the Δ_3 component in service interruption. 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. *IDMP’s IP-based fast handoff technique provides a solution that does not require IP-layer information to be carried in layer-2 signaling messages or require the adjacent BSs to be aware of each other’s identity.* Since we do not re-establish packet forwarding until the MN has performed a subnet-level registration at the new BS, IDMP’s fast handoff process does not eliminate Δ_2 , the delay incurred in the subnet-level registration process.

A. The Fast Handoff Procedure

We assume that a layer-2 trigger will be available (either to the MN or to the old BS) indicating an imminent change in connec-

tivity. We explain the fast handoff mechanism using Figure 4, 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 registration (using IDMP) with SA_3 , this subnet agent 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.

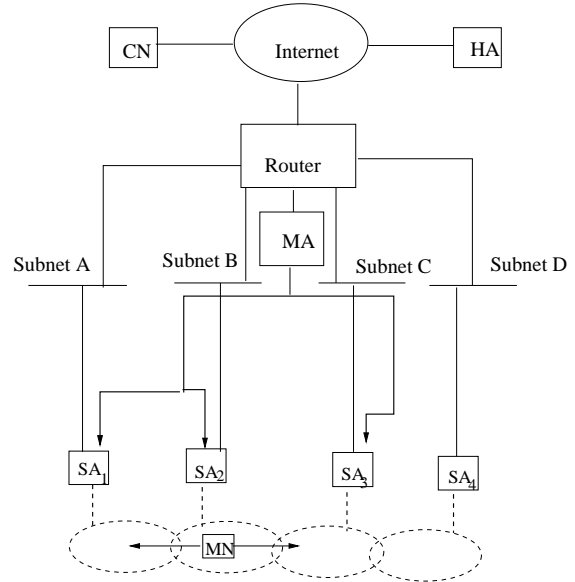


Figure 4: IDMP Fast Handoff

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

- Unlike other existing 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 registration 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. It is the MA which decides the set of target BSs to which in-flight packets are multi-

cast. This is especially useful in scenarios where the MN may be in contact with multiple BSs and is unable to specify the future point of attachment exactly. While current cellular networks use a network-controlled handoff technique (where the base station controller (BSC) determines the candidate BS based on link-layer measurements supplied by the MN or BS), 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.

- c) Our multicasting scheme prevents unnecessary wastage of wireless bandwidth, since a BS does not unilaterally transmit all arriving multicast packets over the wireless interface. Such proactively multicast packets are temporarily buffered by an BS in per-user buffers and forwarded to the MN over the wireless interface only if it happens to register at that BS. In case the MN does not register at a particular SA, the buffered packets are discarded after a specified maximum time interval.
- d) IDMP's fast handoff scheme does not eliminate Δ_2 from the service interruption time; it merely delays the transmission of packets arriving during this instant. By requiring the MN to first perform an IDMP subnet-registration before receiving any such packets, we avoid the need to define new extensions to layer-2 protocols for transporting IP-layer information (such as the CoA). By buffering packets during this transient, we are however able to avoid the loss of in-flight packets. Most multimedia applications, such as VoIP, are typically able to tolerate variations 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. Our approach does not require the establishment of dynamic multicast groups for individual MNs. The membership of the neighborhood set is also 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.

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 performs a subnet-specific registration with an SA (say SA_3 in Figure 4), the SA can then forward any cached packets to the MN before the intra-domain location update process is complete.

Simple calculations 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

The IETF is currently considering [11] several alternative approaches for supporting fast handoffs within the Mobile IPv4 context. In the pre-registration approach, the MN is assumed to initiate a new Mobile IP registration with the new FA via its current FA; the current FA tunnels such a registration request to the new FA. This mechanism requires the current FA to be aware of the identity of the neighboring (candidate) FAs and to advertise their IP-layer identities using *proxy* router advertisements. This proposal eliminates the sequential delay due to $\Delta_1 + \Delta_2$, since the new MIP registration (via the old FA) can occur concurrently with the establishment of radio connectivity at the new FA (Δ_1). However, unlike IDMP's fast handoff mechanism, this approach requires FAs to be aware of the identities of their neighbors and to pre-establish secure tunnels with them.

In the post-registration approach, layer-2 triggers are used to establish a transient tunnel between the current FA and the new FA. This mechanism requires the current FA to proactively tunnel all inbound packets to the new FA via this layer-2 triggered tunnel, thus allowing an MN to maintain its IP connectivity via the old FA even while it is performing a new Mobile IP registration at the new FA. As before, this approach requires an FA to be either aware of the identity of its neighboring FAs or to instantaneously form secure tunnels with a designated FA. Moreover, this proposal assumes that layer-2 authentication is adequate to (temporarily) authenticate the MN at the new FA; packets are forwarded to the MN by the new FA before the IP subnet registration process (Δ_2) is completed.

Similar fast handoff mechanisms have also been proposed [10] for Mobile IPv6. These mechanisms essentially rely on the establishment of a transient tunnel between the MN's old and new point of attachments, and the ability to use link-layer signaling to indicate IP-layer connectivity parameters.

IV. 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:

- i) A single device may have multiple wireless interfaces and hence, need to maintain multiple simultaneous bindings.
- ii) Pico-cellular layouts, especially for local area wireless access technologies, can lead to very frequent changes in the point of attachment.

A. Paging Operation for Idle Hosts

IDMP's IP-layer paging solution provides a flexible and radio-technology independent solution to this important problem and helps to minimize the power wasted by an MN in unnecessary mobility-related signaling. To motivate IDMP's paging solution, note that the 'multicasting' scheme described for fast handoff support in section III 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, called Paging Area Identifiers (PAI). 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., SA Agent Advertisements). In fact, such IP-layer advertisements may optionally be combined with link-layer beacons.

IDMP's paging scheme is visually illustrated in Figure 5. 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 SA_1 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 update latency between the MN and its MA). For example, call setup delays in VoIP are typically around 2.5 sec; accordingly the paging latency is expected to fall within the targeted bounds.

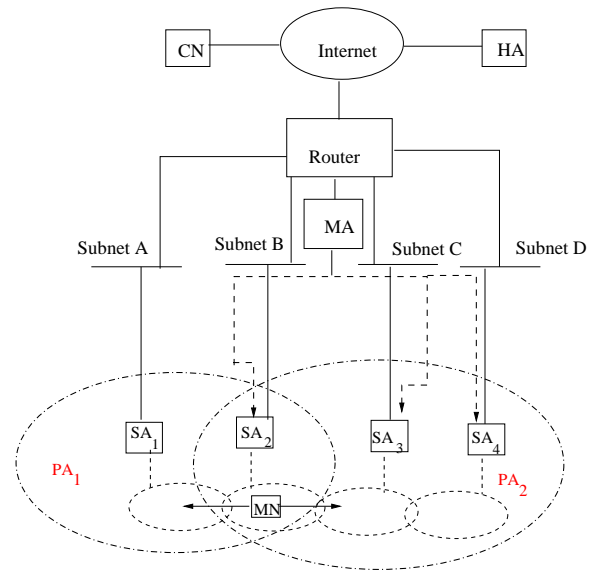


Figure 5: IDMP Paging Mechanism

B. Paging Implementation

Each PA is identified by a unique domain-specific multicast address; a BS belonging to a specific PA must permanently subscribe to the corresponding multicast group. Note also that, similar to the overlapping *registration area* (RA) concept in current cellular networks, a BS can subscribe to multiple multicast groups and hence, be associated with multiple PAs (Figure 5).

The base IDMP specifications need 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 local registration, while the MA would continue to (mistakenly) unicast arriving packets to the MN's last registered local care-of address. Moreover, when an MN is in idle state, it performs an intra-domain location update by indicating its new PAI only when it changes its current PA; otherwise, the mobile is content to passively monitor the periodic link-layer multicast beacons.

C. Comparison with Alternative IP Paging Schemes

We now compare IDMP's paging mechanism with some of the recently proposed approaches to IP-layer paging in wireless cellular networks:

- The P-MIP proposal [12] for paging in Mobile IP is very similar to our IDMP paging specifications. Like IDMP, P-MIP associates a Mobile IP FA with one or more paging area identifiers; an MN in dormant (idle) mode performs Mobile IP registrations upon crossing into a new paging area. The P-MIP specifications, however, assume the use of Mobile IP as the underlying mobility management protocol and requires all FAs to be aware of the identity of their neighboring FAs. In contrast, IDMP's paging func-

tionality does not assume the use of Mobile IP and hence, does not rely on Mobile IP security mechanisms; moreover, IDMP requires only the MA to be aware of the identity of all Subnet Agents associated with a particular paging area.

- Proposals for paging in Cellular IP and HAWAII essentially distribute the paging operation over multiple nodes in the access domain. While these protocols also use the paging area concept to reduce the number of updates generated by dormant MNs, the extent of a specific PA is only *implicitly* defined by the location of *paging caches*. To resolve the location of a currently unregistered MN, a paging request is broadcast over all domain nodes that lie downstream of this paging cache; it is thus difficult to dynamically alter the size or shape of a specific PA without actually changing the location of the paging cache. Moreover, such an implicit definition of a paging area makes it difficult to define overlapping PAs; it is well known that paging areas need to be overlapped to avoid the problems of ping-ponging. Perhaps, most importantly, it is difficult to apply such paging mechanisms in 4G environments, where an access domain may comprise multiple access technologies with overlapping cellular layouts, and thus does not exhibit a strictly hierarchical topology. IDMP, in contrast, defines a paging area in a topology-independent manner, based on an explicit set of SAs that subscribe to the corresponding multicast group.

V. CONCLUSIONS AND FUTURE WORK

In this paper, we presented two enhancements to the IDMP solution for IP-based hierarchical mobility management. We considered a 4th generation cellular environment, and motivated the need for developing IP-layer fast handoff and paging solutions that would work across heterogeneous access technologies.

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 does not assume specific link-layer capabilities (such as the ability to transport specific IP parameters) and does not require an individual SA to establish a transient tunnel with another explicitly defined neighboring SA.

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 some other proposals for IP-based paging, our mechanism does not assume a tree-like topology and

allows easy configuration of variable-size and overlapping PAs. Table I summarizes the commonalities and differences between the mechanisms used for fast handoffs and paging in IDMP.

TABLE I
OVERVIEW OF IDMP FAST HANDOFF AND PAGING

Function	Primary Mechanism	Handling of In-flight Packets
Fast Handoff	Multicasting from MA to candidate SAs	Buffered by candidate SAs
Paging	Multicasting by MA to SAs within Paging Area	Buffered by MA

We currently have a Linux and FreeBSD based implementation of IDMP deployed in our testbed. Additional details of the implementation, as well as preliminary performance results, are available in [6]. We are currently working to measure the performance improvements offered by our fast handoff and paging procedures.

VI. ACKNOWLEDGMENTS

This article is based on our previously published material at 3G Wireless 2001 conference by Delson Group (delson.org). Ashutosh Dutta would like to acknowledge Henning Schulzrinne for many helpful discussions related to multicast. The work of S. K. Das at UTA is partially supported by NSF grants EIA-0086260, EIA-0115885, and ITR IIS-0121297.

REFERENCES

- C. Perkins, “IP Mobility Support for IPv4, revised”, draft-ietf-mobileip-rfc2002-bis-08.txt, IETF, September 2001.
- A. Campbell, J. Gomez, S. Kim, Turanyi Z., C-Y. Wan and A. Valko, “Design, Implementation and Evaluation of Cellular IP,” IEEE Personal Communications Systems Magazine, August 2000, pp. 42-49.
- R. Ramjee, T. La Porta, L. Salgarelli, S. Thuel, K. Varadhan, L. Salgarelli and L. Li, “IP-based Access Network Infrastructure for Next-Generation Wireless Data Networks,” IEEE Personal Communications Systems Magazine, August 2000, pp. 34-41.
- 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, IETF, July 2000, Work in Progress.
- S. Das, A. Misra, P. Agrawal and S. K. Das, “TeleMIP: Telecommunication Enhanced Mobile IP Architecture for Fast Intra-Domain Mobility”, IEEE Personal Communications Systems Magazine, August 2000i, pp. 50-58.
- S. K. Das, A. Misra and S. Das, “Qos Support in the DMA Architecture: Implementation and Performance Analysis”, to appear in IEC Annual Review of Communications, Volume 55.
- 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.
- H. Soliman, C. Castelluccia, K. El Malki and L. Bellier, “Hierarchical MIPv6 Mobility Management (HMIPv6)”, draft-ietf-mobileip-hmipv6-05.txt, IETF, July 2001, Work in Progress.
- E. Gustafsson, A. Jonsson and C. Perkins, “Mobile IP Regional Registration”, draft-ietf-mobileip-reg-tunnel-05.txt, September 2001, IETF, Work in Progress.
- G. Dommety, editor, “Fast Handovers for Mobile IPv6”, draft-ietf-mobileip-fast-mipv6-03.txt, IETF, July 2001, Work in Progress.
- K. ElMalki, editor, “Low Latency Handoffs in Mobile IPv4”, draft-ietf-mobileip-lowlatency-handoffs-v4-03.txt, IETF, November 2001, Work in Progress.
- X. Zhang, J. Gomez and A. Campbell, “P-MIP: Paging in Mobile IP”, Proceedings of International Workshop on Wireless Mobile Multimedia (WoW-MoM), July 2001.

Archan Misra (archan@us.ibm.com) is currently a Research Staff Member with the Pervasive Security and Networking Department of IBM Research, Hawthorne, NY, where he investigates protocols and architectures for future Wireless Internet services. As part of his previous work at Telcordia Technologies, he has 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 has published more than forty research papers in the areas of wireless and pervasive networking, Internet congestion control and buffer management, and filed for four US patents in wireless mobile networks. He received the Best Paper Award in IEEE MILCOM 2001. Dr. Misra received his Ph.D. in Electrical and Computer Engineering from the University of Maryland at College Park in May, 2000.

Subir Das (subir@research.telcordia.com) received the 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, India. Since 1999 he has been at Telcordia Technologies Inc. and currently a Research Scientist in the Wireless IP Research Lab. During 1997-99, Dr. Das was a faculty member in Electronics and Electrical Engineering Department, Indian Institute of Technology, Kharagpur, India. Subir has designed several protocols and architecture for next generation wireless networks particularly in the area of autoconfiguration, registration and mobility management. He has also developed several prototypes for Mobile IP and SIP-based seamless connectivity in wireless LAN-based networks. His current research interests include mobility management in next generation wireless networks, wireless multimedia, user authentication, auto-configuration of adhoc, mobile networks. He has published nearly 25 papers in these areas. He is a member of IEEE and currently co-Chair in IETF pana Working Group.

Ashutosh Dutta (adutta@research.telcordia.com) is currently a Research Scientist in Telcordia Technologies's Advanced Networks Systems Research Laboratory. He has a Bachelors degree in Electrical Engineering from R.E.C Rourkela, India, Masters Degree in Computer Science from NJIT, and Professional Degree in Electrical Engineering from Columbia University, USA. Ashutosh has been working in Computer Systems and Networking area for the last 15 years. He has dealt with variety of high-speed networks, and has been responsible for implementing many systems and networking related projects. Prior to joining Telcordia Ashutosh was the Director of the Central Research Facilities in Columbia University's Computer Science Department for eight years. His research interests include Session Control Protocols, Internet architecture, streaming multimedia, 3G-IP, multicast. Ashutosh is a member of IEEE and ACM. He is also currently pursuing his part-time Ph.D. degree in Electrical Engineering at Columbia University, New York.

Anthony McAuley (mcauley@research.telcordia.com) received a PhD from Hull University, England in 1985. He was a research fellow in Caltech from 1985-1987. Since 1987 he has been at Telcorida 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. His current research includes work on automatic configuration and reconfiguration, security, mobility and service differentiation. Anthony 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++). Anthony has also worked on IP multicasting in ad-hoc and large-scale networks (including the Comprehensive Test Ban Treaty network), efficient error correction and detection codes, and design of VLSI chips (for everything from microprocessors to asynchronous packet switches).

Sajal K. Das(das@cse.uta.edu) received B.Tech. degree in 1983 from Calcutta University, M.S. degree in 1984 from Indian Institute of Science, Bangalore, and PhD 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 (CRWMan) 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). Dr. Das is a recipient of UNT's Honor Professor Award in 1991 and 1997 for best teaching and outstanding research; and UTA's Outstanding Senior Faculty Research Award in Computer Science in 2001. He has visited numerous universities, research organizations and industry research labs for collaborative research and invited seminar talks. He is frequently invited as a keynote speaker at international conferences and symposia. His current research interests include resource and mobility management in wireless networks, mobile and pervasive computing, wireless multimedia and QoS provisioning, mobile Internet, network architectures and protocols, parallel/distributed processing, grid computing, performance modeling and simulation. He has published over 200 research papers in these areas, directed numerous projects funded by industry and government, and filed four US patents in wireless mobile networks. He received the Best Paper Awards in the ACM MobiCom'99, ACM MSWiM 2000, and ACM/IEEE PADS'97. Dr. Das serves on the Editorial Boards of Computer Networks, Journal of Parallel and Distributed Computing, Parallel Processing Letters, Journal of Parallel Algorithms and Applications. He was General Vice Chair of ACM MobiCom 2000 and IEEE HiPC 2001, General Chair of IEEE MASCOTS 2002, ACM WoWMoM 2001, and as TPC member of numerous IEEE and ACM conferences. He is a member of the IEEE TCPP Executive Committee and Advisory Boards of several cutting-edge companies.