Efficient Macro Mobility Management for GPRS IP Networks

Seong Gon Choi, Rami Mukhtar, Jun Kyun Choi, and Moshe Zukerman

Abstract: **Mobile IP has several inefficiencies, and was not originally designed for situations where both peers are highly mobile. We present a mobility management solution that retains compatibility with existing Internet protocols, whilst increasing the efficiency of communications between two GPRS mobile hosts. Our proposal eradicates triangle routing and minimizes handover latency. We show by numerical analysis that the routing optimization improves the performance of TCP controlled data flows, reducing buffering requirements and minimizing the recovery time after a handover occurs.**

Index Terms: **GPRS, handover, mobile IP, mobility management, TCP performance.**

I. INTRODUCTION

The increasing need for Internet Protocol (IP) portable computing devices has spawned a great deal of research that has stemmed from the original Mobile IP proposal [1]. The majority of these proposals have focused on providing wide area mobility support to network connections that span over a range of network fabrics. In order to remain compatible with existing IP network infrastructure and protocols, Mobile IP confined its operation to the IP layer and made minimal assumptions about the network environment. For these reasons, Mobile IP has enjoyed success as a macro-mobility protocol, providing mobility across heterogeneous networks. In contrast, micro-mobility, i.e., mobility within a single homogeneous network, is commonly provided by mechanisms that are specific to the underlying link layer fabric.

Mobile IP has several inefficiencies due to its confinement to the IP layer, the most predominant being triangular routing and location update latency. Triangular routing refers to bypassing packets through the Home Agent (HA), increasing the distance between two peers. This problem is further exacerbated when both peers are mobile, in this case the route becomes rectangular - the apexes defined by the peers and their respective HAs. Solutions to this problem have been proposed [2], [3] and [5]. However, they require wide spread modifications to Internet and/or network protocols, and are not optimized for situations in which both peers are mobile. Since Mobile IP depends on an IP network to transmit all signaling information, Mobile IP location updates depend on the performance of the IP layer. This can result in in-flight packet loss, or packet re-ordering. These effects not only reduce the perceived quality of service, but also adversely affect the performance of TCP controlled data flows. This paper is motivated by the recent activities in the 3rd Generation Partnership Project (3GPP) which is finalizing its standards for a high speed General Packet Radio Service (GPRS) [4] over the GSM air interface. There has been an explosion of data exchange between GSM terminated peers facilitated by the deployment of the Short Messaging Service (SMS) (part of the GSM standard), which provides a facility to exchange short text messages. Recently, the SMS service has evolved to Multimedia Messaging Service (MMS) [9] and [10] that provides a facility for the transfer of multimedia rich messages between GPRS peers. These developments are already fueling an increase in data exchange between mobile peers that is likely to continue into the future.

MMS has its limitations as a peer-to-peer data transfer mechanism; similar to SMS, it is based on a "cache and deliver" paradigm. Data objects are first sent to a proxy-rely server where they are cached before being delivered to the destination. This paradigm may be suitable for the exchange of multi-media rich messages, but is limited in application. We argue this claim using the following two distinct examples. There is usually a fixed charge for sending an MMS message fixed by the service provider. This is usually much higher than the retail cost of sending the equivalent amount of data using the network. The service provider will often justify this disparity by claiming that MMS is a service that requires significant additional infrastructure. This makes MMS economically unsuitable for general peer-to-peer exchange of data objects when both peers are known to be active. Another important application is the exchange of general time sensitive data, e.g., interactive peer-to-peer gaming, which may prove to be the next "killer application" of cellular data services. Delayed delivery of this type of content is of no value to either of the peers. Furthermore, these sessions are likely to have a very long lifespan, and hence many handovers between networks are likely to occur throughout the session duration. This motivates the development of a mech-

---

Manuscript received May 14, 2002.

This paper was originally submitted to the special issue on 3G/4G, December 2002.

S. G. Choi is with Information and Communications University(ICU) Yusong, Daejon, Korea, e-mail: sgchoi@icu.ac.kr.

R. Mukhtar is with ARC Special Research Centre for Ultra Broadband Information Networks (CUBIN), Dept. of Electrical and Electronic Engineering, Univ. of Melbourne, e-mail: r.mukhtar@ee.mu.oz.au.

J. K. Choi is with ICU Yusong, Daejon, Korea, e-mail: jkchoi@icu.ac.kr.

M. Zukerman is with CUBIN, Dept of Electrical and Electronic Engineering, University of Melbourne, and he is visiting the Electronic Engineering Department, City University of Hong Kong between November 2002 and July 2003, email: m.zukerman@ee.mu.oz.au.

This research was supported in part by KOSEF (Korea Science and Engineering Foundation) and OMIC (Ministry of Information and Communication) of Korean government through OIRC(Optical Internet Research Center) project. Part of this work was supported by the Australian Research Council.

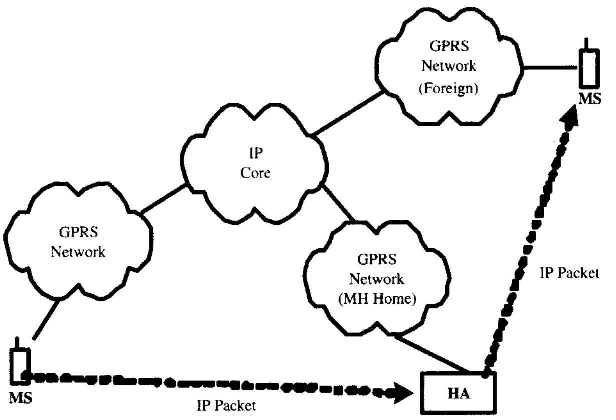

Fig. 1. Existing mobile IP solution.

anism that provides general peer-to-peer connectivity, where peers located on a GPRS network can form direct connections between each other, exchanging data objects at will. Furthermore, the user should perceive this connectivity to be seamless from connectivity to the wider Internet.

This motivates the present paper. We propose a macro mobility management solution across GPRS networks that leverages off GPRS infrastructure and signaling but retains the public Internet as a payload transport fabric. This will enhance the efficiency of data sent between two GPRS peers. It is important to highlight that our scheme is a macro mobility solution, handling routing between hosts terminating on different GPRS networks. In contrast, the GPRS network itself handles micro mobility routing to a host within a single GPRS network. In the present paper, we are concerned with handovers between networks as opposed to between cells. This will become increasingly common as specialist service providers are founded to provide hotspot service coverage, leaving wide area network coverage to the traditional carriers. In this environment, a handover between networks will be incurred due to moving out of a facility, campus, or central business district.

Our solution complements Mobile IP, which attempts to provide seamless mobility across heterogeneous networks by optimizing communications between two GPRS terminated hosts, i.e., both peers are mobile. Mobile IP would still need to be employed to handle connections terminating out of the scope of the global GPRS network.

II. MOBILE IP BACKGROUND AND MOTIVATION

In order to highlight the optimizations provided by our solution, we consider an implementation of the existing Mobile IP protocol over an example GPRS network. Fig. 1 depicts the use of Mobile IP to provide communication between two peers located on different GPRS networks. The Mobile Station (MS) is assigned a home network, in which there exists a HA. The HA serves a dual purpose: (1) to maintain a Permanent IP Address (PA) which all data to the MS is directed, and (2) to intercept any data, and tunnel it to the MS current Care of Address (CoA), which is the MS current point of attachment. When an MS attaches to foreign network, a Foreign Agent (FA) provides the MS with a new temporary CoA, and a tunnel termination point. Each time the MS attaches to a network, it promptly informs its HA of its new CoA in order to ensure that data can be directed to it.

A Correspondent Node (CN) always sends IP packets to the MS Home Address. As a matter of convention, in the present paper, we shall refer to the mobile host sending data as the CN. The HA then intercepts these packets and tunnels them to the MS's current CoA. This leads to the establishment of a triangle route between the CN and MS, a characteristic of Mobile IP, which is depicted in Fig. 1.

When a CN requests the resolution of an MS domain name, Mobile IP requires that a permanent HA address is returned. The most predominant reason for this is to eliminate the need to deal with the problem of the MS CoA changing mid session. If this was to happen, each active CN would have to be notified [5] requiring the MS to maintain significant state information, which would only be possible with significant software modifications to both the MS and CN IP stack. Mobility within a GPRS network (micro-mobility) does not lead to the same inefficiencies since the GPRS network simply updates routing information within the network to correspond to host mobility. This is practical within a GPRS network, since it has been specifically designed to maintain mobility information.

Another limitation of Mobile IP is its independence from lower layer signaling. Although this makes the protocol independent of underlying link layer implementations, it comes at a cost of increased signaling latency, since IP connectivity needs to be established before any signaling can take place. Fig. 2 illustrates the signaling required for location updates of an MS attached to a GPRS network employing Mobile IP for wide area mobility management, notice that Mobile IP signaling only occurs after all GPRS signaling has been completed. In Section V, we will argue that a direct consequence of this is that during handover, significant packet loss or reordering can occur due to the latency attributed to waiting for an IP connection before signaling can take place. This problem is exacerbated by the Mobile IP routing inefficiency, as it leads to an increase in the number of in-flight packets that are vulnerable to loss or redirection during a handover. We propose to address this problem by using existing GPRS signaling to update Mobile IP mobility information.

III. A NEW MOBILITY MANAGEMENT ARCHITECTURE

A. System Architecture

The fundamental approach of our solution is to eliminate the triangle routing problem by using the Home Location Register (HLR) to maintain mobility information, whilst spreading the tunneling function of the HA over a network of new elements, which we call the Mobile Proxy Server (MPS). We envisage that existing SGSNs would provide MPS functionality by a software upgrade. Our solution only requires that changes be made to the GPRS network, and does not require any modifications to either the MS or CN.

Using the resources of the HLR, MPS nodes form a virtual

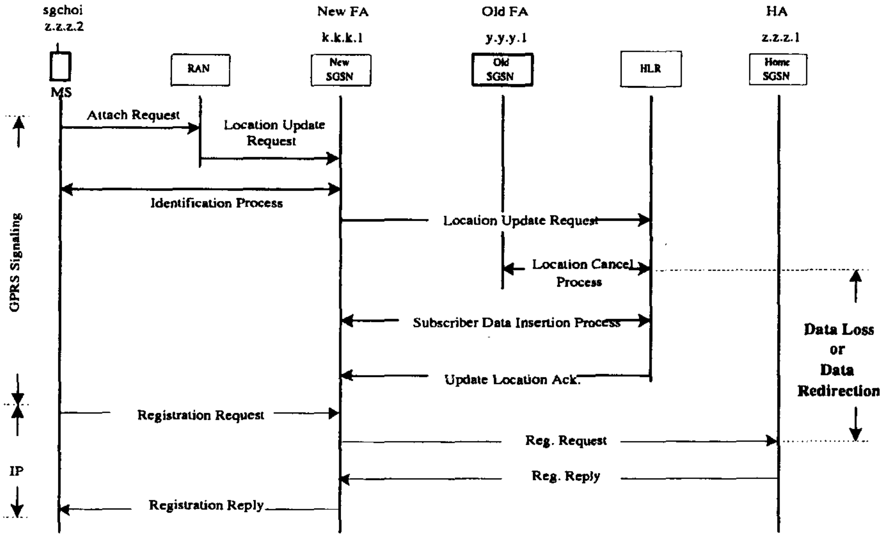

Fig. 2. Signaling procedure between an MS and a HA during handover.

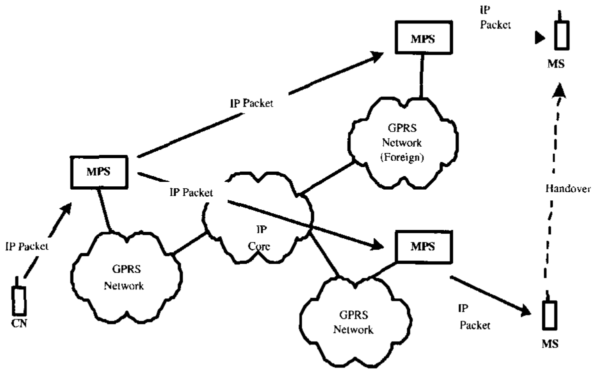

Fig. 3. Route optimization using the proposed solution.

network that maintains routing information for active CN-MS pairs. IP packets sent by the CN to the MS are intercepted by the MPS local to the CN and then tunneled to the MPS that the MS is currently attached to. In this way, any changes describing the MS's point of attachment is confined to the MPS network, neither the MS nor the CN need maintain any mobility management information.

As is currently implemented, the HLR maintains up-to-date mobility information for all local hosts. The MPS network only maintains routing information for active connections, providing a scalable mobility management solution between disjoint GPRS networks.

An MPS provides efficient IP services such as domain name resolution and IP tunneling using the IP fabric for payload data exchange, whilst minimizing the delay in obtaining up-to-date mobility information by using low latency inter-operator GSM signaling for the exchange of mobility information with other MPS nodes and other GPRS network elements.

The advantages of this solution are twofold: triangle routing between GPRS peers is eliminated, and packet latency due to handover is minimized. Fig. 3 illustrates route optimization as the MS hands-off between two GPRS networks. It is important to note that if the peer originates from a network that does not have mobility support then triangular routing is still inevitable.

Table 1. Modified HLR entry table.

MS GSM ID #	MS Domain Name	PA	CoA	PPA		
				1	2	....
IMSIz	sgchoi@icu.gprs.kr	x.x.x.4	y.y.y.1	k.k.k.1	j.j.j.1	....

B. HLR Modifications

We propose that the GSM location registers are modified, so that in addition to their current functionality, they also maintain IP layer addresses and provide Domain Name Server (DNS) hosting. Table 1 illustrates the required information that the HLR must maintain for each user, and is explained below.

The Permanent Address (PA) is a permanent IP address assigned to an MS. This is similar to the home address provided by a mobile IP HA. The CoA is the MS current point of attachment to the GPRS network. The Peer Permanent Address (PPA) fields store the PA of all peer hosts that are currently connected to the MS. This is similar to the way a HLR maintains state information for currently active circuit switched connections in GSM.

When a host transitions into the idle state, or completes a handover, GSM signaling is used to update the appropriate entries in each HLR. GSM signaling is not vulnerable to the unpredictable performance of an IP network, additionally, it can occur before IP connectivity is achieved. In contrast, Mobile IP location updates rely on a best effort network for signaling transport. Moreover, an MS must register with a foreign agent before any location updates can occur, incurring additional delay.

When a HLR receives a request to resolve a domain name, it determines whether the request was made by an MPS (it does this by explicitly querying the DNS server). If an MPS made the request, then it returns the PA as well as the current CoA. However, if an arbitrary DNS server made the request, then only the PA is returned. This provides compatibility with networks that do not provide mobility support, as shall be described in the

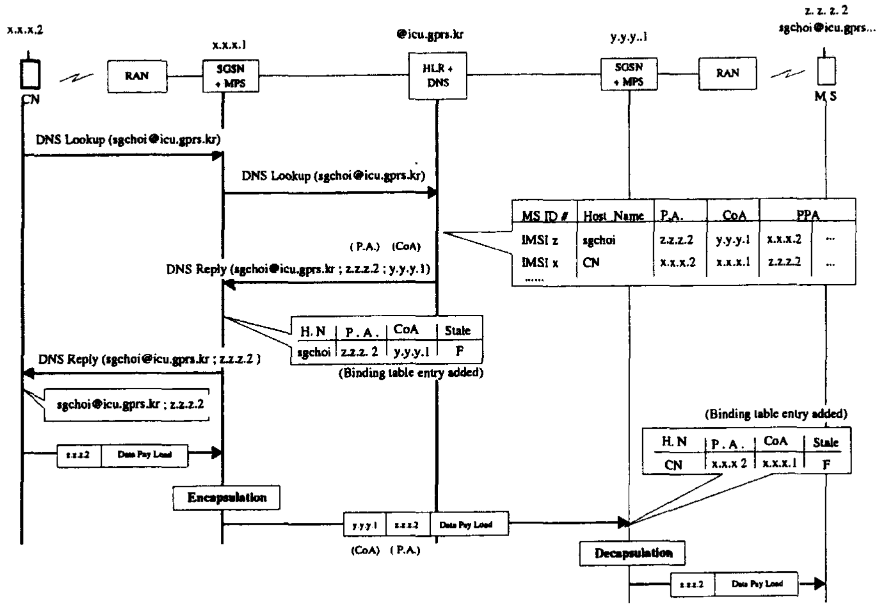

Fig. 4. DNS lookup, data transmission and corresponding binding table updates.

next sections.

IV. PROPOSED OPERATION

A. MPS Functionality

An MPS provides the following functions: Permanent Address hosting, Care of Address Hosting, Domain Name Service, and Egress filtering and IP Tunneling.

What follows is a description of each of these functions and how they support efficient wide area mobility management.

A.1 Permanent Address Hosting

The PA is a permanent, globally routable, IP address, which uniquely identifies each MS. The PA also enables an MS to be contactable by CNs that are out of the GPRS network scope.

If the MS's HLR is unable to verify that the request is being made by an MPS or by another network capable of maintaining mobility information, then the PA is returned. In this case, since the CN's network does not support mobility, triangular routing will occur.

However, the MS is still contactable and routing performance will be equal to that provided by Mobile IP. Packets that are received by the home MPS are forwarded to the MS current CoA, as specified by the Mobile IP standard [1].

A.2 Care of Address Hosting

The MPS provides a CoA that can be assigned to an MS to provide a point of attachment to the network. This is similar to the functionality provided by a Foreign Agent (FA) in Mobile IP.

Table 2. MPS binding table.

H.N	P.A	CoA	Stale Flag
CN	x.x.x.1	y.y.y.1	FALSE

A.3 Domain Name Service

The MPS provides a DNS lookup service for all nodes connected to the GPRS network. If the domain name corresponds to a GPRS peer (as notified by the HLR-DNS host server), then the MPS expects both the PA and current CoA to be returned, and will accept future CoA updates.

A.4 Egress Filtering and IP Tunneling

The MS serving MPS monitors all outbound IP packets sent to the Internet. The MPS keeps track of which addresses correspond to the PA of an MS. If a packet is destined to an MS (as determined by its PA), then the MPS will encapsulate the IP packet with a header containing the MS current CoA before sending it. The MPS located at the MS point of attachment will then receive this packet, remove the IP header containing the CoA and pass the packet directly to the MS.

B. Binding Table Updates

In order to maintain low latency communications between two GPRS peers, it is essential that both the MPS currently serving the CN and the one currently serving the MS are aware of the current CoA of their respective peer. This is accomplished by employing a binding table (Table 2) at each MPS. The binding table maintains information on all currently active sessions between GPRS peers. Binding table entries are temporary.

Binding table entries will be created in response to any one

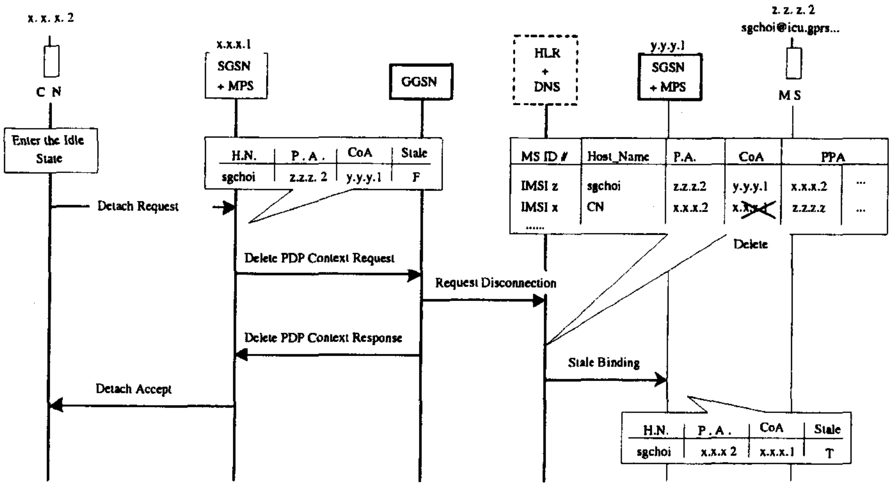

Fig. 5. Location update sequence as a result of host disconnection.

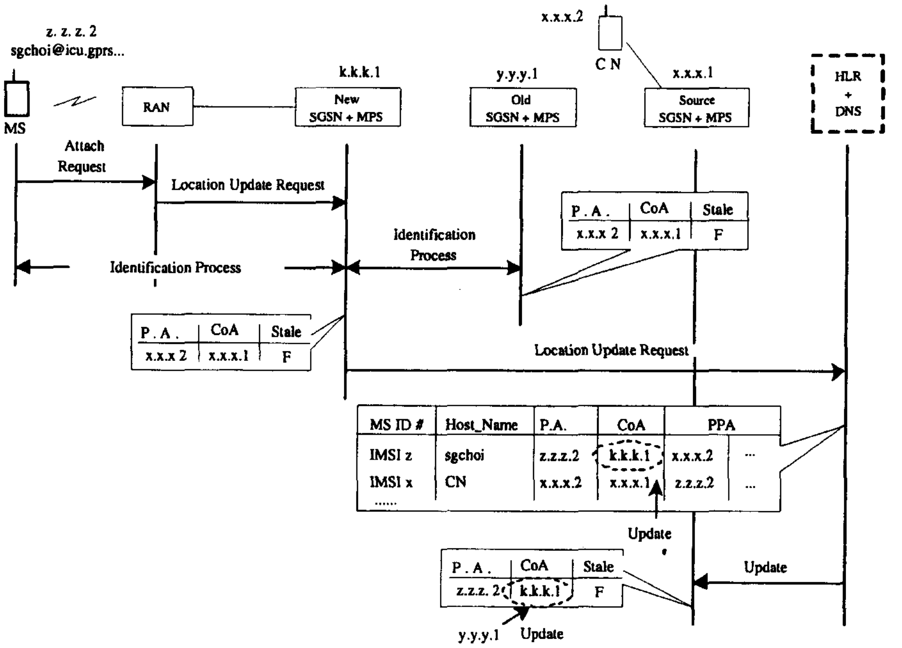

Fig. 6. Location update as a result of handover.

of the following three events: (1) A CN initiates a connection using an MS's host name, e.g., sgchoi@icu.gprs.kr. As a result, a DNS lookup will return both the PA and CoA of the MS (as illustrated in Fig. 4). (2) A tunneled packet is received from another MPS. (3) A CN explicitly uses a MS's PA, without having to initiate a DNS lookup, e.g., as may be the case if the user explicitly uses the IP address to initiate a data transfer session. In this situation there may not be a binding table entry for this PA in the CN's serving MPS. On receiving the first data packet destined for this PA the MPS will ascertain whether a CoA is available. This is achieved by performing a reverse DNS lookup on this PA. This will return the host name, which can then be used to request the CoA (if one exists for that PA).

On creation of a binding table entry, the PA and the CoA fields are filled, and the stale flag is set to "FALSE".

C. Location Update Policy

Each MPS follows a location update policy that attempts to minimize signaling and data path length.

When a host detaches itself from an MPS serving itself because it either moves to another MPS service region (Fig. 6), or transitions to the GPRS "idle" state (Fig. 5), the following sequence of events will be initiated:

1. As soon as the MS disconnects from the old network, each peer that has a PPA entry in the MS's HLR table is notified that the MS's binding is now stale. By quickly informing each peer MPS that the binding information is stale, data packets routed to the previous MS MPS are minimized and only in-flight packets that have already left the CN serving MPS are sent to the previous MS MPS.

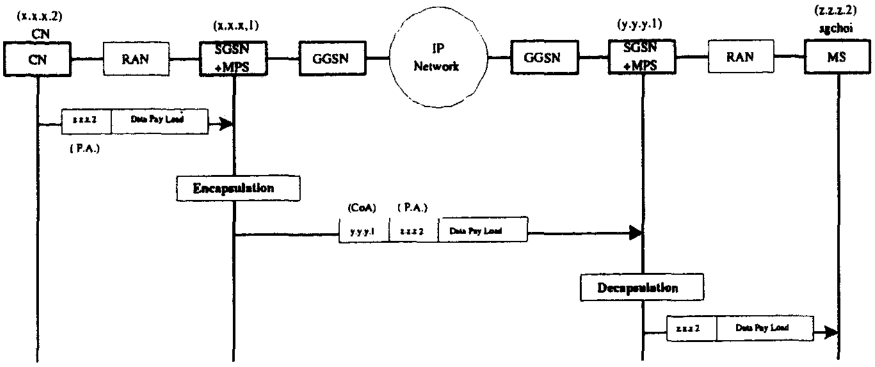

Fig. 7. Packet transmission between two hosts attached to an MPS.

2. If the MS has moved, then the new MPS will inform the HLR of its new CoA. In turn, the HLR will notify each CN serving MPS for which there is a PPA entry of the new MS CoA. Otherwise, if the MS has become "idle", then the old MS serving MPS will notify the HLR of the MS status. In this state, the HLR will use GSM paging to locate the MS before returning a CoA in response to a subsequent DNS request.
3. When an MPS serving the CN receives a packet destined to an MS with a stale binding entry, the packets are buffered at the MPS. If the MS has moved to another network, then the MS's HLR will inform the MPS serving the CN of the new CoA as soon as this becomes available. This two-stage approach minimizes packet reordering, which can be detrimental to TCP performance, as we shall show in Section V. In contrast, Mobile IP would continue to allow the CN to send packets to the old CoA until the handover is complete, this will result in significant packet loss or re-ordering.

D. Data Transmission

Once the MS host name has been resolved to a PA, and a binding table entry has been created at the MPS serving the CN, data transmission can begin. Referring to Fig. 7, a packet transmission encompasses the following sequence of events: (1) The CN sends an IP packet addressed to the MS's PA. (2) MPS serving the CN intercepts the packet, and recognizes that the PA has a valid binding table entry. (3) The packet is encapsulated with current CoA of the MS and tunneled to MPS currently serving the MS. (4) The MPS serving MS receives the tunneled packet, de-encapsulates the IP packet and delivers it directly to the MS as identified by the PA.

V. PERFORMANCE ANALYSIS

In this section, we present numerical analysis to support our claim that the proposed architecture will increase the performance as perceived by the user. In order to conduct meaningful analysis, which provides a fair comparison of performance, it is essential that we make as few assumptions about the underlying implementation as possible. Accordingly, we must avoid attributing specific costs and latencies to any particular event.

We shall focus our analysis on the performance of bulk data transfer controlled by TCP flow control assuming TCP Reno [11]. Due to the limitations of a mobile handset, it is most likely that the user will only attempt to transfer a single data object at any time, hence, we assume that only one TCP session is active at any time. Under these assumptions, two aspects of the system's operation will contribute to the perceived performance, packet loss or redirection incurred by handover latency, and packet transmission latency incurred by parallelogram routing.

We shall now show that when compared to Mobile IP, by reducing the routes path length, our proposed scheme is more immune to handover as well as reducing buffering requirements at the wireless-wire line interface.

A. The Model

For the purpose of performance evaluation, we abstract the following simple topological representation for both existing Mobile IP and the proposed enhancement. Referring to Fig. 8(a): the CN is a greedy source sending the data, the MS is currently receiving data. For simplicity, but without loss of generality, assume all packet sizes to be equal. The HA attributed to both the MS and CN are labeled accordingly. We make the realistic assumption that the wire-line network is highly over provisioned, and the wireless link is the predominant bottleneck. Accordingly, in order to make efficient use of the precious wireless link, there is a buffer at the interface between the wire-line and wireless network that buffers packets before passing them to the wireless link layer. This is illustrated in the Fig. 8 by a queue of capacity B (packets) and occupancy Q (packets) served at a mean rate R_C (packets/sec), which corresponds to the mean rate of the wireless channel. We argue that since the network bandwidth is highly heterogeneous, the majority of queuing occurs at this bottleneck. Note that this assumes that the CN's uplink has a mean rate that is greater or equal to R_C. Furthermore, the location of the queue as shown relative to the HA and MS is important as it will determine the effect of handover. Similarly, Fig. 8(b) illustrates the respective topology for our proposed scheme, where the HA and FA are replaced by MPSs. Our intention is not to present a detailed exact analysis, but rather to present an approximate mean value analysis that provides insight into comparing TCP performance over both schemes.

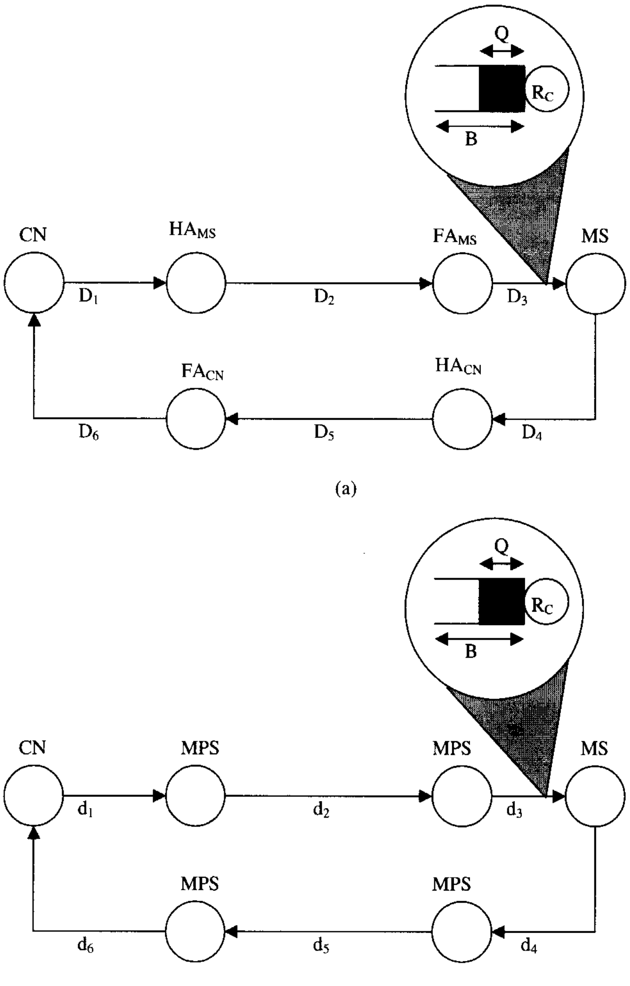

(a)

(b)

Fig. 8. Topological representation of (a) Mobile IP and (b) Enhanced GPRS mobility management.

Trip times that include propagation and packet processing delays, but not queuing delays between each node are denoted by D_1 to D_6 on Fig. 8(a) and d_1 to d_6 on Fig. 8(b) respectively. In order to make a fair comparison between the two schemes, assume that the locations of all HAs and FAs in the Mobile IP topology are identical to the locations of MPSs in the topology using our proposed scheme. This enables us to state the following relationships between the trip times:

$$D_3 = d_3, \tag{1}$$

$$D_6 = d_6, \tag{2}$$

$$D_1 \geq d_1, \tag{3}$$

$$D_4 \geq d_4, \tag{4}$$

$$D_1 + D_2 \geq d_1 + d_2, \tag{5}$$

$$D_5 + D_6 \geq d_5 + d_6. \tag{6}$$

Consider an MS handover between two GPRS networks. Let us define the following quantities:

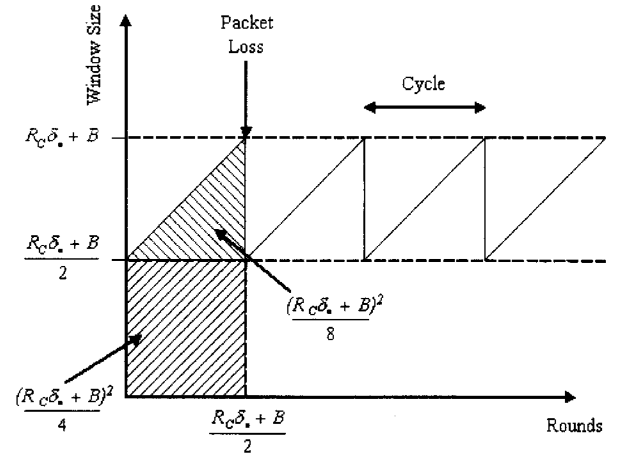

Fig. 9. Macroscopic behavior of TCP Reno in the congestion avoidance phase of operation.

h_L: the time between when the MS disconnects from one network and establishes a link layer connection to the new network.

h_{IP}: the time between when the MS establishes a link layer connection with the new network and obtains a new CoA IP address.

The macroscopic behavior of TCP Reno in the congestion avoidance mode of operation under this scenario is illustrated in Fig. 9, where, $\delta_\bullet$ is the total trip time, devoid of queuing delays and is defined as $\delta_M = \sum_{i=1}^{6} D_i$ for Mobile IP and $\delta_G = \sum_{i=1}^{6} d_i$ for the proposed GPRS scheme. We assume, as an approximation, that in-flight packets are evenly spaced apart. Referring to Fig. 9: a round is denoted by the reception of a window worth of packets, a cycle is denoted by a packet loss due to buffer overflow. Under the stated assumptions, the window size will increase by one packet size at the end of each round, and hence a single packet loss will mark the conclusion of a cycle. Under these deterministic conditions, $3/8(R_c\delta_\bullet + B)^2$ packets will be sent per round.

B. *Effect of Increased Latency*

The bandwidth of the wireless link is precious, as it is a limited resource. Hence, from the users' perspective, they would like to ensure that their allocated share is fully utilized to maximize throughput. Conditional on TCP Reno flow control, and a single active flow, it is clear from Fig. 9 that access link utilization can only be unity when the queue capacity, B, satisfies:

$$B \geq R_c\delta_\bullet. \tag{7}$$

Accordingly, the storage required for our proposal relative to Mobile IP is a linear function of the differences in trip times $(\delta_M - \delta_G)$.

Given the setting of B according to (7), we can calculate the following quantities: Mean queue size, $\overline{Q}$ size:

$$\overline{Q}_\bullet = \frac{\sum_{i=0}^{B}\left(\delta_\bullet + \frac{i}{R_c}\right) i}{\sum_{i=0}^{B}\left(\delta_\bullet + \frac{i}{R_c}\right)} = \frac{(2B + 3\delta_\bullet R_c + 1)B}{3(B + 2\delta_\bullet R_c)}, \tag{8}$$

the mean round trip time, RTT$_\bullet$:

$$RTT_\bullet = \delta_\bullet + \frac{(2B + 3\delta_\bullet R_c + 1)B}{3R_c(B + 2\delta_\bullet R_c)}. \tag{9}$$

For the remainder of the present paper we set, $B = R_c\delta_\bullet$, which is the minimum value of B that will provide an access link utilization of unity. Substituting the optimal value of B into (8) and (9) we respectively obtain:

$$\overline{Q}_\bullet = \frac{5}{9}\delta_\bullet R_c + \frac{1}{9}, \tag{10}$$

and

$$RTT_\bullet = \frac{14\delta_\bullet}{9} + \frac{1}{R_c}. \tag{11}$$

From (11), we can conclude that if the mean RTT is almost double of the travel time, $\delta_\bullet$, irrespective of the mean channel rate, then high utilization can be maintained. From (1)–(6), it is clear that $(\delta_M \geq \delta_O)$ emphasizing the importance of the routing optimization provided by our scheme.

C. Effect of Handover

In the previous section, we demonstrated the effect of increasing the travel time as a result of Mobile IP's routing inefficiency on RTT. In this section, we demonstrate that increasing the product $R_c\delta_\bullet$ as a result of Mobile IP's routing inefficiency coupled with the constraints of the locations of the HA and FA exacerbates the negative impact of handovers.

Both packet loss and latency will result during a handover. Both of these will reduce the throughput of TCP before and for a period after the handover occurs. We wish to characterize how long it takes TCP to re-obtain the full channel rate after a handover occurs.

This will depend on two factors, the number of packets lost, and how many packets are received out of order.

Our analysis is based on the following assumptions:

1. Packets that have already left the FA$_{\text{MS}}$ or the MPS serving the MS cannot be redirected and are lost.
2. FA$_{\text{MS}}$ and MPS serving the MS are notified immediately when a MS disconnects from the network.
3. The HA is only notified when a new CoA is obtained $h_L + h_{IP}$ seconds after the MS disconnected from the previous network.

Referring to Fig. 8, we derive the following quantities:

Mobile IP (M)	Proposed GPRS Scheme (G)
Number of packets lost:	Number of packets lost:
$L_M = \overline{Q}_M + D_3 R_G$	$L_{\text{GPRS}} = \overline{Q}_G + d_3 R_C$
Number of packets redirected:	Number of packets redirected:
$\text{PR}_\text{M} = R_C(h_{\text{IP}} + h_L) + D_2 R_C$	$\text{PR}_\text{G} = d_2 R_C$

The number of packets received out of order will depend on the topology and relative location of the mobility management elements to the MS. However, we can assume that PR$_\bullet$ is an upper bound to the number of packets received out of order. Notice that PR$_\text{M}$ is a function of the time required to complete a handover. In contrast, under the proposed scheme, the number of packets that are received out of order is independent of the handover time. This is a due to the fact that GPRS signaling provides full information to the MPS nodes as soon as link layer connectivity is lost.

As a result of packet loss, TCP will either go into a prolonged state of fast retransmit/fast recovery or will time out. Which event occurs will depend on the statistics of the RTT and the network topology. Hence, it is not possible to determine which one is more likely without making arbitrary assumptions about the network topology. However, we will present the recovery times after each of these events occur to demonstrate that the performance is always worse when Mobile IP is employed.

Fast Retransmit/Fast Recovery

Here we only consider the effect of lost packets, as the effect of re-ordered packets will depend on several factors that cannot be quantified without detailed knowledge of the network topology. Below is a description of the TCP Reno [11] fast retransmit/fast recovery definition that has been simplified to highlight the important features.

1. When a third duplicate acknowledgment is received, set the congestion window size to halve its previous size, but no less than two, and retransmit the lost packet.
2. For each additional acknowledgement received, transmit a new data packet.
3. When an acknowledgment arrives that acknowledges new data (this should occur one RTT after step 1), return to normal mode of operation i.e., Slow Start/Congestion avoidance.

Hence, all of the lost packets (lost in the handover) will be retransmitted over a period of $L_\bullet$ RTTs after the first packet loss was detected. At the end of this period, the window size will be approximately equal to $\max\left(2, \left(\overline{Q}_\bullet + \delta_\bullet R_C\right)/2^{L_\bullet}\right) = 2$, for all $\delta_\bullet R_c \geq 0$, and TCP will be in the congestion avoidance mode of operation. For simplicity, assume R_c and $\delta_\bullet$ remain unchanged after the handover. This assumption is not necessary, but will simplify the analysis for the sake of the present argument. Furthermore, assume $\delta_\bullet R_C \geq 2$ segments, such that $\max\left(2, \left(\overline{Q}_\bullet + \delta_\bullet R_C\right)/2^{L_\bullet}\right) \leq \delta_\bullet R_C$. Hence, RTT$_\bullet \approx \delta_\bullet$ during this recovery period, and we can show that the total time between packet loss due to handover being detected and re-obtaining the channel rate can be approximated by:

$$t_{\text{recovery}} \approx \frac{(L_\bullet + (\delta_\bullet R_c - 2))\delta_\bullet}{C_1}, \tag{12}$$

where C_I is defined as 1 (packet/round) in the TCP Reno congestion avoidance specification. It is clear from (12) that the t_{recovery} is proportional to the square of $\delta_\bullet$. We plot t_{recovery} versus $\delta_\bullet$ for various ratios of $D_3/\delta_\bullet$ (which is equivalent to $d_3/\delta_\bullet$) and R_C = 96 (packets/sec), the results are shown in Fig. 10. It is clear from the figure that the recovery time is reduced the closer the MS is to the FA (or the serving foreign MPS).

Time Out

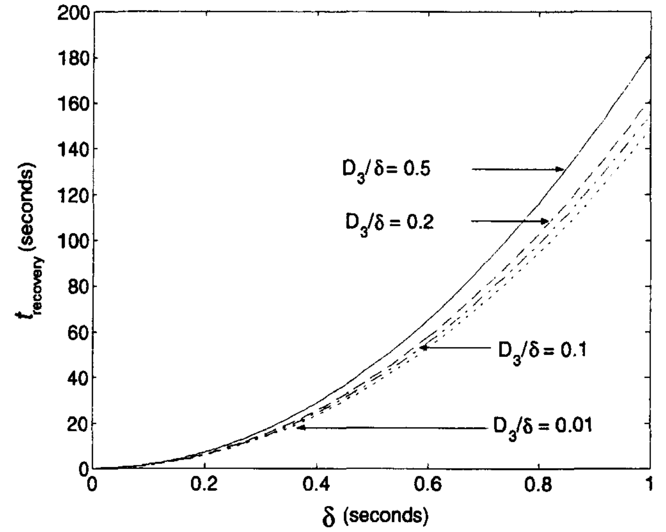

Fig. 10. Time required for TCP to recover from packet loss due to handover as a function of packet travel time.

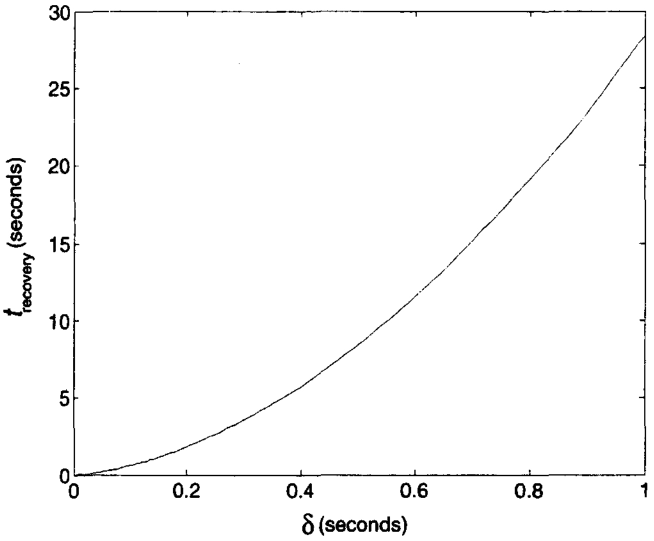

Fig. 11. Time required for TCP to recover a TCP time out due to handover as a function of packet travel time.

A TCP time out occurs if an acknowledgement is not received within an interval, RTO. The RTO interval is dynamically updated based on estimates of the mean and standard deviation of a packet transmission RTT [12]. Handovers that result in a disconnection period greater than the mean RTT are highly likely to result in a time out; however, this will depend on the statistics of the RTT. Given that a time out does occur, we wish to determine how long it takes from the time of reconnection to when TCP re-obtains the available channel rate.

When a time out occurs, the RTO value is doubled, the slow start threshold [11] is set to half of the congestion window size, and the congestion window is set to one segment size.

This causes TCP to return to the slow start mode of operation. Furthermore, since each successive time out doubles the RTO, successive failed retransmissions occur at exponentially increasing intervals. The net effect is that if successive time outs occur as a result of disconnection during an extended handover period it may take a considerable amount of time after the connection is re-established before TCP attempts to send another retransmission. Hence, the recovery time after a time out will depend on two factors, the current window size, and the handover period relative to the exponential multiples of the current RTO value.

Let $RTO_\bullet \approx \delta_\bullet + \overline{Q}/_{R_C} + S$ where S is four times the smoothed RTT deviation estimate as defined in [12]. Let $h_\bullet$ denote the handover time ($h_M = h_L + h_{IP}$ or $h_G = h_L$), and let z be the smallest integer that satisfies $(1 - 2^z)/(1 - 2)RTO_\bullet \geq h_\bullet$. Then the mean recovery time can be approximated by:

$$\begin{aligned} t_{\text{recovery}} &\approx \Delta H + (\log_2(14/9\delta_\bullet R_C + 1/9) - 1)\,\delta_\bullet \\ &\quad + (2/9\delta_\bullet R_C - 1/9)\,\delta_\bullet, \end{aligned} \tag{13}$$

where $\Delta H = \left({}^{(1-2^z)}/_{(1-2)}\, RTO_\bullet - h_\bullet\right)$. corresponds to the time for TCP to time out again after the handover is complete (this assumes the worst case, that no packets are buffered within the network). The second (bracketed) term corresponds to the average time required to complete the slow start phase of operation that will terminate when the congestion window has grown to half of the value it was immediately before the handover, the third (bracketed) term corresponds to the time required for the window to grow to the bandwidth delay product in the congestion avoidance mode of increase. We plot t_{recovery} vs. $\delta_\bullet$ for $R_C = 96$ packets/second and $\Delta H = 0$. Notice that the gradient of t_{recovery} in Fig. 10 is greater than that in Fig. 11. This is due to the fact that after a time out, TCP starts in the slow start (exponential increase) phase of increase. In contrast, if a large number of consecutive packets are lost, TCP will reduce the slow start threshold to two, and congestion avoidance (linear increase) will persist from this point. Hence, in this case a time out has less negative impact than multiple consecutive packet losses.

VI. CONCLUSION

We have presented a solution that will enhance the performance of packet data communications between two GPRS peers. The solution does not require any changes to the end hosts, and maintains compatibility with the existing Internet. In order to implement the solution, we have specified additional state information that the HLR should maintain, and introduced a new GPRS network component - the MPS. The MPS provides the functionality of the Home Agent for communication with non-GPRS peers whilst enhancing the performance between GPRS peers.

We have demonstrated that by reducing latency, the proposed scheme improves the performance of TCP. In particular, we have shown that in order to maximize the utilization buffering required in the core constrains the mean RTT to be proportional to almost twice the propagation delay. We have also shown that after a handover, the time TCP requires to re-obtain full channel utilization is proportional to the square of the propagation delay. Hence, by reducing transmission time latency, the proposed scheme significantly improves the performance of TCP data transfers.

ACKNOWLEDGEMENTS

Manuscript received May 14, 2002. This research was supported in part by KOSEF (Korea Science and Engineering Foundation) and OMIC (Ministry of Information and Communication) of Korean government through OIRC(Optical Internet Research Center) project. Part of this work was supported by the Australian Research Council.

We would like to acknowledge the useful feedback provided by Stephan Hanly at the Univ. of Melbourne during the early stages of development.

REFERENCES

[1] C. Perkins (1996). "IP mobility Support," *IETF RFC 2002.*

[2] E. Gustavsson, A. Jonsson and C. Perkins (2000), "Mobile IP regional registration," *IETF Internet Draft.*

[3] D. B. Johnson and C. Perkins (2000), "Mobility support in IPv6," *IETF Internet Draft.*

[4] 3rd Generation Partnership Project (3GPP), "General packet radio service (GPRS)," *3G TS 23.060.*

[5] C. Perkins and D. B. Johnson (2000), "Route Optimization in Mobile IP," *IETF Internet Draft.*

[6] W. Wang and I. F. Akyildiz, "A cost-efficient signaling protocol for mobility application part (MAP) in IMT-2000 systems," *ACM SIGMOBILE 7/01* Rome, Italy 2001.

[7] A. Bertrand, "Jambala mobility gateway-convergence and inter-system roaming," *Ericsson Review*, pp. 89–93, no. 2, 1999.

[8] L. Kleinrock, *Queueing Systems Volume 1: THEORY*, John Wiley & Sons,1975.

[9] 3rd Generation Partnership Project (3GPP), "Multimedia messaging service (MMS): functional description," *3GPP TS 23.140*, Dec. 2002.

[10] Wireless Application Protocol Forum, "Wireless application protocol, multimedia messaging service, architecture overview specification," *WAP-205- MMSArchOverview-20010425-a*, Apr. 2001.

[11] W. Stevens, "TCP slow start, congestion avoidance, fast retransmit, and fast recover algorithms," *IETF Request for Comments, RFC2001*, Jan. 1997.

[12] W.R. Stevens, *TCP/IP Illustrated*, vol. 1, Addison-Wesley, 1994.

Seong Gon Choi received his BS degree in electronics from Kyung Pook National University in 1990 and MS degree from Information and Communications University (ICU) in 1999. He worked for LG Information and Communications during 1992-1998. He is currently taking a doctoral course in ICU, Daejeon, Korea. His research interests include Mobility, high-speed network architecture, and protocol.

Rami Mukhtar received a Bachelor of Engineering with first class honours from the University of Melbourne, Australia in 1998. In 1997, he worked for the Commonwealth Science Industrial Research Organization on the automatic indexing of speech in video sound tracks. During 1999-2001, he worked on a collaborative project between The University of Melbourne, Nortel Networks and Telstra Research Laboratories investigating IP Mobility Management and wireless IP transport protocols. He has written papers on the design and analysis of transport layer flow control algorithms for wireless networks, performance of wireless link layer protocols, IP mobility management and content distribution networks. He is currently completing a Bachelor of Science and a Ph.D. degree at the University of Melbourne.

Jun Kyun Choi received his BS degree in electronics from Seoul National University in 1982 and MS and PhD degrees from KAIST in 1985 and 1988. He worked for ETRI during 1986-1997. He is currently working as an Associate Professor in Information and Communications University, Daejeon, Korea. His research interests include high-speed network architecture and protocol.

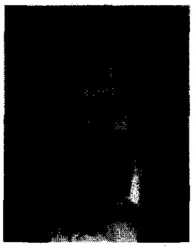

Moshe Zukerman received his B.Sc. in Industrial Engineering and Management and his M.Sc. in Operation Research from Technion-Israel Institute of Technology and a Ph.D. degree in Electrical Engineering from The University of California Los Angeles in 1985. He was an independent consultant with IRI Corporation and a post-doctoral fellow at UCLA during 1985-1986. During 1986-1997, he served in Telstra Research Laboratories (TRL), first as a research engineer and between 1988-1997 as a project leader managing a team of researchers providing expert advice to Telstra on network design and traffic engineering, and on traffic aspects of evolving telecommunications standards. He is the recipient of the Telstra Research Laboratories Outstanding Achievement Award in 1990. In 1997, he joined The University of Melbourne where he is a professor responsible for expanding telecommunications research and teaching in the Electrical and Electronic Engineering Department. Since 1990, he has also taught and supervised graduate students at Monash University. He has served as a session chair and member of technical and organizing committees of numerous national and international conferences. He gave tutorials in several major international conferences such as IEEE ICC and IEEE GLOBECOM. He served on the editorial board of the Australian Telecommunications Research Journal during 1991-1996. He also served as a Guest Editor of IEEE JSAC for two issues: Future Voice Technologies and Analysis and Synthesis of MAC Protocols. Presently, he is serving on the editorial board of the IEEE/ACM Transactions on Networking, the International Journal of Communication Systems, Computer Networks, and as a Wireless Communications Series Editor for the IEEE Communications Magazine. He has over 180 publications in scientific journals and conference proceedings.