Proactive IP Mobility Management for Context-aware all-IP Wireless Access Networks

Theo Pagtzis [‡] Rajiv Chakravorty[†] Jon Crowcroft[†]

[‡] Department of Computer Science, University College London, Gower Street, London WC1E 6BT, U.K. t.pagtzis@cs.ucl.ac.uk

[†]Computer Laboratory, University of Cambridge , William Gates Bldg., JJ Thompson Avenue, Cambridge CB3 0FD, U.K. {firstname.lastname}@cl.cam.ac.uk

***Abstract*— A major challenge in building 'all-IP' wireless access networks, besides the use of IP as the unifying layer, relates to transparency of the IP handoff process as the mobile node (MN) transits across heterogeneous wireless network domains in IP Mobility Management (IPMM). Transparency in IP handoffs, however, must be effected in two separate contexts: IP-addressing and (re-)connection latency. Excessive delays during an IP handoff degrades the seamlessness of IP transmission between the MN and its peers.**

In this paper, we assess experimentally how transport protocols such as TCP and subsequently applications can under-perform as the MN performs an IP handoff between two heterogeneous wireless networks, namely WLAN and GPRS. We discuss why dynamic establishment of IP context-state can help address these limitations that seem inherent in heterogeneous environments. To this end, we evaluate by means of simulations, the Proactive IP Mobility Management model as a means of effecting proactive context-state establishment, between candidate points of attachment, for the purposes of reducing/eliminating IP handoff delay.

***Index Terms*— Experimentation with real networks, System design, Simulations**

I. Introduction

INTERNET has been experiencing a paradigm shift in access practices, largely due to the growing acceptance of emerging wireless technologies under the unifying layer of IP. Despite this, the case for truly seamless, mobile, all-IP wireless access remains to be seen, for deployable novel data services over multiple heterogeneous wireless Internet service providers (WISPs) [1]–[3].

To achieve mobile networking of truly ubiquitous dimensions, architecture and implementation of all-IP wireless access networks requires more than just the next generation IP protocol [4] or its mobility extensions [5] as a unifier; it calls for an elaborate set of IP signalling interactions in different **contexts** so as to render all-IP wireless access viable from both commercial and performance perspective [6]. Examples of such contexts are authentication, authorisation and accounting (AAA) or Quality in provisioned service (QoS) to name but a few. Standard IP Mobility Management (IPMM) [5] effects a transparent mapping between the home address of a mobile node (MN) and the IP address (care-of address) acquired at the visited point of attachment by performing an ***IP handoff***; during an IP handoff the MN is required to attain/configure some IP connectivity state, comprising of multiple context-specific state components, before resuming communications with its peers.

IP addressing state is only one of the types of context required to admit an MN and its active IP flows into a visited subnet/domain; for instance, security context may have to be re-established prior to any packet flow; in the case of resource reservation for the purposes of Integrated Services QoS provisioning, the entire end-to-end path needs to be re-established at the new network link. For AAA-based admission control, credential verification must be effected with the home network before the MN can resume communications over the visited network link. Each context-specific state to be (re-)established requires one or more round trip times (RTT) in terms of protocol interactions. As a result, the total transmission delay becomes large enough to impede any notion of transmission delay transparency in the communications of the MN during an IP handoff. While significant transmission delay imposed by the IP handoff process affects all data applications the effect is exacerbated for data services with interactive real-time constraints such as navigation/locator services, interactive audio/video or network games. Hence, for true seamlessness, IP handoff transparency must be effected in two separate contexts: IP addressing and (re-)connection delay. Standard IPMM lacks of connection delay transparency during an IP handoff as experienced by an on-going IP flow of the MN; this is owed to the fundamental reactivity in control signalling in multi-context establishment effected in the current IPMM standard, during an IP handoff.

At the same time, the maturity of the underlying wireless IP technologies unveils evolutional dilemmas between advancing current cellular systems (PCS) and enhancement of 802.11x Wireless LANs, towards novel data-services. *Deployment scale* arises as the underlying factor of these technological crossroads [7], towards a ubiquitous Wireless Internet; cellular networks enable typically long transmission ranges with nominal low-bandwidth connectivity; 802.11 WLANs offer high-bandwidth access for a short transmission range. Each wireless technology provides a different performance profile of IP connectivity.

To this end, there has been a growing body of research

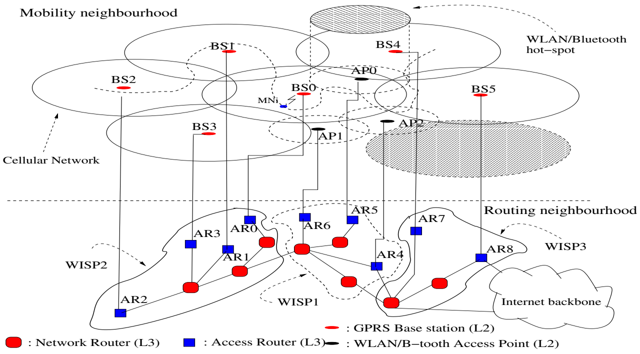

Fig. 1. Advanced multi-WISP (hot-spot) deployment scenarios and Mobility-Routing Neighbourhood mapping in Proactive IPMM

focusing on hybrid multi-network/access solutions [8]. These solutions envisage competing WISPs provisioning IPMM services over different wireless technologies (GPRS or WLAN) with disparate service/performance characteristics for a range of data services. For instance, in hot-spot deployment [9], the MN may be offered increased bandwidth via 802.11x WLAN WISP domain deployed in high concentration areas (hot-spots) with restricted mobility patterns (e.g. malls, city neighbourhoods, airports), while still subscribing to wide-area cellular WISP connectivity, as shown in Figure 1.

II. Motivation

From the perspective of the MN, WISP performance or data service capabilities, appear ad-hoc in its transit path, as it roves - to its ignorance - into new coverage areas/IP-points-of-attachment towards its destination. MNs and subsequently users are expected to avail of such provisioning by seeking to IP-attach/handoff horizontally or vertically [10] to the WISP that matches best their instantaneous needs in terms of IP service/performance capabilities.

Current IPMM standards make no provision for dynamically enabling the MN with context-specific state pertaining to IP service/performance characteristics of its current or future points of attachment; instead they base IP handoffs on *reactive* physical-layer properties such as signal-to-noise ratio (SNR), or bit-error-rate (BER) typically manifested as router advertisements [11] at the network layer of IP networks [5]. Hence, awareness over such context changes require human intervention introducing yet one more factor of disconnectivity, preventing ubiquity in wireless Internet from materialising.

Transmission delay transparency clearly prohibits reactive IPMM context-establishment mechanisms due to the excessive transmission delay incurred during an IP handoff. MNs will best ascertain WISP service/performance capabilities, **in advance** of their IP-cell transitions so as to afford context-aware IP handoff control over available WISP/cell candidates while preserving *the seamlessness principle* towards a truly ubiquitous wireless Internet. This calls for an intelligent IPMM model that **pro-acts** to enable the MN with context-aware IP handoff control.

In the remainder of this paper we assess performance limitation arising from cross-network (horizontal or vertical) IP handoffs between heterogeneous WISP networks by means of experiments. These issues motivate further our subsequent simulation study on a devised **context-aware** Proactive IPMM architecture [12], [13]. We assess the performance of the underlying Proactive IPMM protocol mechanisms designed to resolve performance limitations such as excessive transmission delay during IP handoffs, thus preserving the seamlessness principle towards truly ubiquitous mobile networking.

The rest of this paper is structured as follows: Section III provides a short experimental assessment of IP handoffs effected between real-world heterogeneous WISP networks. Section IV provides a description of the simulation model devised for the purposes of comparative analysis between standard IPMM and the proposed Proactive IPMM model. Section V presents our simulation results. We conclude with a summary of our findings and subsequent work in Section VII.

III. Experimental assessment of heterogeneous IP handoffs

An IP handoff is a special case of re-routing. Re-routing occurs typically deep inside the network; thus, end-systems are (largely) impervious to its effects. On the contrary an IP hand-off not only occurs at the 'edge' , but is also potentially visible to end-hosts. Normally API extensions [14] to IP mask this, by unifying link-layer transmission effects onto the IP layer – and unavoidably on the IP handoff process –, introducing performance issues above IP for end-to-end protocols such as TCP. Previous work [15] has concentrated on mitigating link-layer transmission characteristics of a wireless network such as GPRS in TCP. This section describes performance issues induced solely by the IP handoff process as the MN transits between two heterogeneous wireless networks.

A. *Disparity in Link-layer Characteristics*

Each wireless technology imposes a disparity in wireless link characteristics in all-IP wireless networks. This is the case for two widely-deployed wireless networks, WLANs and GPRS, in terms of round trip times (RTT). Highly variable RTTs between communicating peers induce significant transmission delay during an IP handoff and in particular when the MN updates its bindings (BU signal) with its peers.

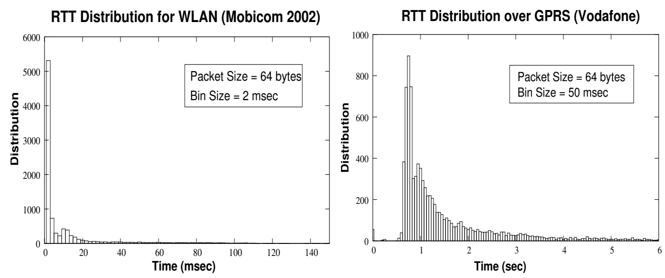

Fig. 2. Histograms showing RTT (ping) distribution (using 64 bytes size ping packets) for 10,000 ping packets over: (a) a congested WLAN network (Mobicom 2002), (b) a commercial GPRS network.

Figure 2 shows the round trip time (RTT) distribution for 10,000-packet ping traces taken over a congested WLAN and an operational GPRS network. Over WLAN, ping RTTs are 1-hop away from the destination. Over GPRS, the edge routers of both the GPRS network (GPRS IPv6 AR) and lab's network (WLAN IPv6 AR) are connected through a well-provisioned IPSec VPN (see Figure 3). The RTT between these two edge routers is negligible ($\leq 20ms$) compared to total RTTs experienced by pinging through the GPRS network an in-lab destination. It can be seen that RTTs for GPRS links are distributed an order magnitude higher than those for WLANs. In addition, RTTs over WLANs, are small (e.g., 5ms-100ms) and relatively stable, whereas for GPRS, RTTs are typically larger *but significantly variable* (e.g., 800ms-2s).

B. *IP handoff behaviour*

To investigate further the performance of inter-network handoffs we have devised an experimental wireless WLAN-GPRS testbed [16], [17] as shown in Figure 3. The testbed facilitates a truly heterogeneous all-IP wireless environment where MNs use multi-access mechanisms to handoff across wireless IP overlays [18].

The inter-network IP handoff experiments comprised of a series file downloads of 25MB size, initiated by the MN over WLAN from a web-server acting as the corresponding node (CN). Inter-network handoffs were then forced on the MN toggling between the GPRS and WLAN networks. IP handoff behaviour was monitored through `tcpdump` traces collected on the AR of the home (WLAN) network. Experiments were conducted during minimal load conditions on the GPRS network.

Figure 4 shows one of the IP handoff traces collected at the AR of the home network; the remaining traces are omitted as they exhibit near-identical behaviour. An IP handoff from WLAN→GPRS takes around 4s. Given the highly variable RTT on the GPRS network, the BU signal experiences a significant delay before it reaches the CN. By that time the TCP session at the CN-source has backed-off exponentially, retransmitting several times before the receipt of bytes is restored at the MN over GPRS network.

In the reverse direction (GPRS→WLAN) the IP handoff takes significantly longer, about 7s with no retransmissions. The behaviour is explained by the amount of excess buffering offered by the GPRS network per MN (around 120Kb); as it hands-off onto the WLAN network the MN ceases to receive

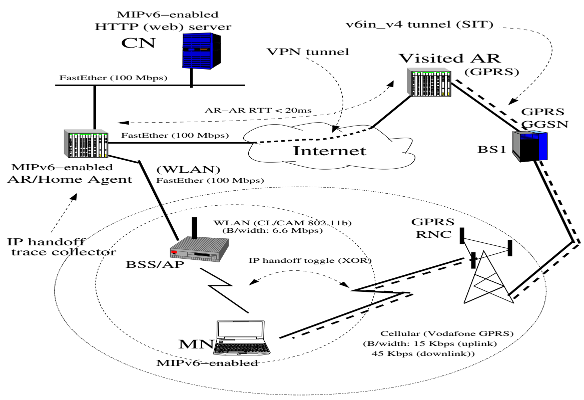

Fig. 3. Simplified view of the experimental WLAN-GPRS testbed.

the packets buffered at the GPRS network (GGSN); after associating with the WLAN BSS it sends a BU to the CN and awaits TCP traffic from the CN source. However, until that time, all TCP packets sent by the CN-source remain buffered at the GPRS network. Because of this the retransmit timeout (RTO) and subsequently the RTT experienced over the WLAN network after the IP handoff become *inflated*. Thus, the BU signal by itself fails to restore receipt of TCP packets at the MN over the WLAN network, until the RTO timer (inflated) is triggered at the CN-source. Table I provides a breakdown of the transmission delay incurred during an IP handoff between GPRS and WLAN networks, averaged over 10 consecutive runs. The RTT variability is manifested in both detection t_d and registration time t_r of the total IP handoff delay t_h by equal variability on the frequency of router advertisements and binding update delay respectively. An IP handoff in the GPRS→WLAN direction, incurs a significantly smaller t_d than in the reverse direction, since router advertisements experience a significantly smaller one way-delay on the WLAN network. Similarly, an IP handoff in the WLAN→GPRS direction experiences significantly higher t_r than in the reverse direction due to inflated RTT and subsequently RTOs at the CN-source. In both cases, a significant number of packets are lost; the effect is exacerbated for an increasing rate of cross-network IP handoffs.

While TCP can severely under-perform as the MN transits between heterogeneous points of IP attachment, the aforementioned IP handoff performance may affect any transport or application protocol. In fact, the disparity in link-layer characteristics manifested as high RTT variability impacts with greater severity interactive (UDP) applications that typically require stringent delay bounds in the delivery of IP traffic. Clearly under the current IPMM model, the seamlessness principle cannot be preserved towards a truly ubiquitous, next generation (v6) wireless Internet. Existing IPMM standards must be reconsidered so as to address the diversity, and high variability in the underlying wireless link characteristics towards a unifying all-IP mobility management model.

C. *The need for context-awareness in IPMM*

To alleviate the performance limitations witnessed during experimentation with heterogeneous IP handoffs it is necessary to exploit knowledge of the wireless link conditions to inform the IP handoff process; that is to say, exploit context-specific state pertinent to the performance of an IP handoff to effect two fundamental functions in a heterogeneous wireless internetwork: *adaptation* of active IP flows and *selection* over choice, such as WISP domain capabilities.

Given the stringent requirement for seamlessness and the existing high RTT variability over heterogeneous WISP domains, *reactive* manipulations of context in an IPMM model sound already obsolete. The observation is rather intuitive since, typically, integrity and transparency over any set of tasks requires planning; planning by definition can never be reactive. Exploiting context-specific state for the purposes of a heterogeneous IP handoff, requires proactive signalling deliberations, well in advance of the MN's IP handoff transition.

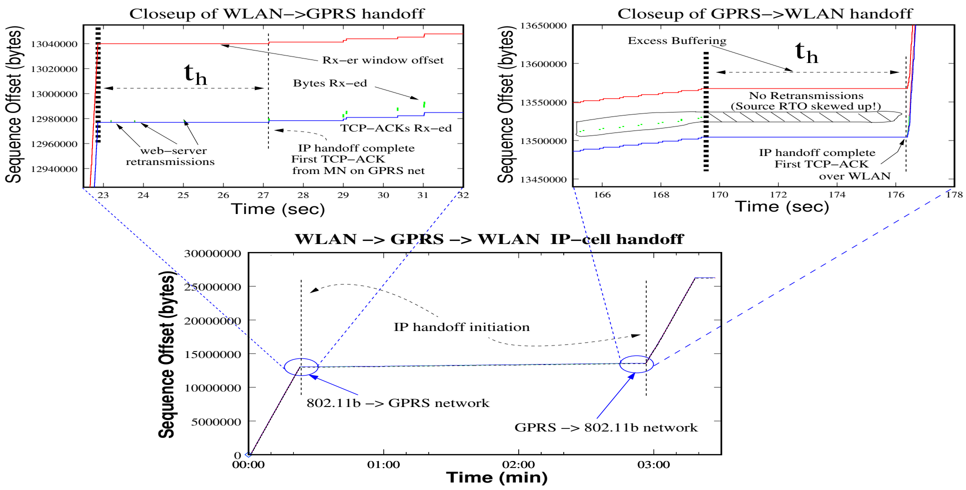

Fig. 4. Time-sequence plots, during a 25MB file transfer for inter-network MN handoffs between GPRS and WLAN. Top-left provides a zoom-in on WLAN→GPRS IP handoff; top-right focuses on GPRS→WLAN IP handoff.

WLAN↔GPRS Handoff	WLAN→GPRS				GPRS→WLAN			
(units = msec)	Min	Mean	Max	Std. Dev.	Min	Mean	Max	Std. Dev.
Detection Time (t_d)	200	808	1148	304	739	2241	3803	919
Configuration Time (t_c)	0.853	0.870	0.890	0.009	0.380	1.062	1.186	0.233
Registration Time (t_r)	2339	2997	3649	395	2585	4654	7639	1611
Total Handoff Delay (t_h)	3323	3806	4438	310	5322	6896	8833	1118

TABLE I

DELAY COMPONENTS (IN MS) DURING AN WLAN↔GPRS IP HANDOFF

D. *Beyond Mobile IP*

To this end, previous work of [12], [19] and its extensions for IP Radio Access Networks (RANs) [13], has proposed an intelligent IPMM architecture that effects *proactivity* in IP mobility management; this is achieved by establishing forward mappings between the underlying physical *Mobility Neighbourhood* of the MN and the corresponding set of access routers, abstracting in that manner *mobility-hop* routing; the latter encompasses discovery and identification of the *Routing Neighbourhood Vector* (RNV) of the MN illustrated in Figure 1. By tracking dynamically the mapping between the Mobility and its corresponding Routing Neighbourhood, the proactive IPMM model can convey forward service/performance context-specific state in a forward manner; this caters for pre-configured IP connectivity state to allow seamless IP handoffs, as well as support controlled selection over the available candidate points of attachment. Seamless IP handoffs are achieved by abstracting the pre-configured Care-of Addressing of the MN onto a multicast Handoff Care-of Address as shown in Figure 5, for the purposes of the IP handoff. In this manner the MN can *bounce* freely between neighbouring points of attachment while receiving traffic over the multicast handoff identifier. Use of the multicast routing identifier abstraction during the IP handoff, removes the reactive reliance of the handoff's success from major handoff delay components, namely detection and registration time.

Support of the above IPMM features enables IP handoff triggers to be based on proactive availability of specific service/performance characteristics in the form of *capabilities*, rather than on mere reactive link-layer signal fluctuations. These capabilities delineate different types of context regarding IP connectivity, such as link-layer capability, bandwidth, billing/authorisation, security state and so on. Currently IPMM standards [5] provide no support towards context-state establishment/relocation or support for selection over candidate points of IP attachment.

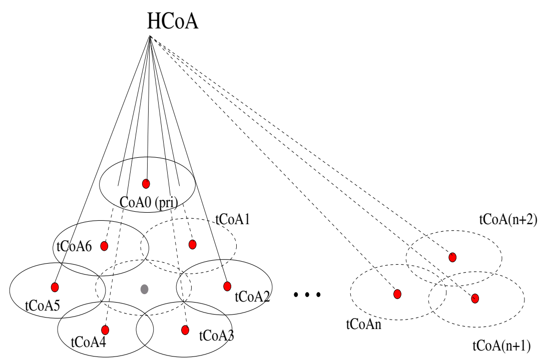

Fig. 5. Abstracting pre-configured IPv6 CoA addressing state onto a multicast address

The following sections present a simulation study of the devised proactive IPMM model contrasted with current IPMM standards based on Mobile IPv6. To investigate the IP handoff effects on the proposed IPMM model, the simulation experiments employ UDP streams as the application workload instead of TCP flows. The particular experimental workload was chosen to emphasize on the importance of the seamlessness principle in IPMM mechanisms, given the stringent delay requirements for interactive real-time IP traffic.

IV. Experimental Setup

Key experimental objectives in the course of this simulation study is to show that, proactive IP addressing state establishment, as a form of context-transferable state, exhibits significant reductions to the overall latency experienced by the MN over current IPMM standards based on Mobile IPv6. That is to say, pre-configuration of such IP connectivity context in advance of the MN's IP handoff, can provide significant reductions in reconnection delay occurring during the IP handoff of the MN. Such reductions argue towards seamless re-routing of the active IP (v6) flows communicated between the MN and its peers. Furthermore, our simulations attempt to demonstrate the effect of dynamic selection of candidate access routers (CAR) in proactive IPMM during real-time communications between the MN and its peers; ultimately investigate lossy conditions and varying AR topologies.

To establish seamlessness over real-time IP traffic delivery, we investigate the amount of IP handoffs performed within a Mobility Neighbourhood with no loss of packets. To measure this we introduce the *Handoff Efficency* (HE) metric. HE of an IPMM mechanism is defined as the number of IP handoffs with no packet loss effected at a single Access Router (AR) over the total number of IP handoffs performed at that router. This metric encompasses the factor of seamlessness from the perspective of absolute packet loss during the transmission of an interactive real-time IP flow between the MN and its peers. Due to space limitations, description of investigations on perceptible packet loss are omitted.

A. Simulation Implementation

For the purposes the comparative analysis, we carried out a series of simulation experiments employing the mobility model of NS-2 [20] as per the MobiWan NS-2 implementation effort [21]. The particular model was further extended, first to encompass key parts of the proposed Proactive IPMM mechanism. These parts implemented in particular:

- indirect and direct Routing Neighbourhood vector (RNV) updates; these messages are effected between the MN and the AR (indirect) and between ARs (direct) respectively within the Routing Neighbourhood of the MN for the purposes of establishing the Routing Neighbourhood surrounding the MN.
- state-full in-advance generation of multiple Care of Addresses (CoAs) per MN. Each address is derived by the address prefix of the neighbouring AR from the current AR[1] given the MAC address of the MN.
- extend multicast join for all CoAs proactively allocated for a single MN. The mapping assumed a single fixed multicast IP address (HCoA) per MN, for the entirety of a simulation run, namely all joins/leaves of tentative CoAs allocated proactively.
- pruning and grafting of CoAs from the multicast tree (HCoA). Currently pruning and grafting were assumed to consume negligible time since they are expected to happen proactively.
- CS-Push and HCoA-initiate/suspend.

Each AR was extended to map/control multiple access points (APs). Only limited IPv6 protocol functionality [4], [11] was required for the purposes of the devised simulation experiments; it encompassed implementation of state-full generation of multiple CoAs (addresses) per MN over AR neighbours, as well as IPv6 router solicitation/advertisement from Neighbour Discovery (ND) protocol standard[2] [11]. Other functions of IPv6 Neighbour Discovery, such as neighbour unreachability detection (NUD), address resolution or Duplicate Address Detection were not essential for the validity of the experiment and thus were not implemented/employed;

For the purposes of the general experimental scenario, the network topology assumed implementation of an access point (AP) node, forwarding traffic between an wire-line correspondent node (CN) and each wireless MN. Each AP transmitted data on a unique wireless channel; periodic beacons were being broadcast over a common channel available to all MNs as per the 802.11 MAC protocol. While the MN moves between different coverage areas, it associates with a different AP and hence communicates at a different channel, although it maintains link switching/re-association over the beacon frequency.

All ARs participating in the simulation experiments integrated the implementation of Protocol Independent Multicast (PIM) dense mode (DM) multicast routing function to abstract the set of pre-allocated CoAs under a single MN routing and addressing identifier for the period of the IP handoff.

A handoff algorithm was implemented for the transition between cells of the Mobility Neighbourhood vector; the algorithm was based on the signal strength of received beacons at the MN from the associated AP. Three levels were used to effect the handoff process; H_{min} as the minimum signal level essential to consider that member of the current MNV to be a handoff candidate. H_{res} as the signal level where the MN commences selection amongst candidate AR neighbours. We note that at this signal level the MN is also required in proactive IPMM, to signal a HCoA-initiate message to the current AR; the probability of such signal effectively inversely proportional to the received signal strength threshold H_{res}. This is to inform the attached AR that IP multicast is enabled at the Routing Neighbourhood. When the signal strength at the MN is sampled to have reached H_{off}, the MN must effect the IP handoff, i.e. the selection of the correct CoA from the set of CoAs that have been proactively allocated at attachment time with the current AR. Base IP Mobility is triggered at H_{off} in the simulator reflecting the reactive character of the lost link (no beacon transmission by the current BaseStationNode).

Candidate AR selection is performed as follows: All APs within a Mobility Neighbourhood vector (MNV) mapping to a respective Routing Neighbourhood vector (RNV) and with signal strength higher than H_{min} in the travel path of the

[1]Currently NS-2 allows only a single, static address to be assigned to each node

[2]Part of the router advertisement was ported from old version of NS-2 (v6)

MN are ignored; from the remaining list of candidate ARs, the respective AP with the highest signal strength is selected as the immediate candidate. The list also includes unresolved ARs (i.e. ones where the corresponding APs maintain an indeterminate signal strength with respect to the MN); this is done pessimistically to avoid ignoring potential AR candidates of the RNV vector that simply have not been sensed (through respective signal strength readings).

B. Modelling Assumptions

The simulation model implemented, employed the WLAN 802.11 protocol [22], with Distributed Coordination (DCF) at the medium access control (MAC) sub-layer; no dynamic rate shifting was employed. With respect to link switching latency as investigated from the IEEE 802.11 specification [22] as well as research literature [23], [24] we have derived detailed delay matrices for the purposes of our simulations; these matrices are presented in [25].

Coverage areas for WLAN cells assumed a maximum range of 35m. At this range, the respective signal level thresholds considered for the purposes of the handoff algorithm were configured at 25, 30 and 35 meters respectively; these figures were derived heuristically with approximations based on the 802.11b MAC rate shifting mechanism which varies the modulation scheme as the mobile node reaches certain BER threshold; if the MAC sub-layer effects such change in the modulation scheme essential (due to higher loss rate probability) then the IP layer should best initiate the proactive handoff determination mechanism at that time, in advance of the physical handoff. Note that, the simulations reported in this paper do not account for optimisations made in the proposed architecture over ping-pong effects.

Ten sets of different AP topologies and MN mobility patterns were generated through the Stanford Graph Base interfacing to GT-ITM. The transit-stub model was employed as the generating algorithm for the topology since it produces a more realistic view of the Internet[3]. Each AP topology covered a 500m square grid. Access points were placed randomly in the grid effecting a continuous IP mobility-enabled area with no gaps between coverage areas (WLAN cells).

Each simulated topology encompassed a maximum of 30 ARs (Routing Agents in NS-2) comprising the Routing Neighbourhood layer. Different permutations of the mobility scenarios involved each AR having a different set of AP connecting to it, comprising thus the Mobility Neighbourhood layer; the number of AP per AR was varied between 1 and 3 APs.

The mobility pattern of the MNs employed the random way-point model[4] using a pause time of 10 seconds and maximum speeds of 1, 4 and 8 m/s. Simulation results over random-way point models are likely to have diminished significance for hot-spot deployment models. Hot-spot models typically assume WLANs as a high concentration coverage area whereas GPRS is virtually pervasive. As a result movement from WLAN→GPRS or vice versa is much less likely than movement *within* the WLAN/GPRS cell. This calls for a two-level probability in the movement/pause and velocity distributions which are currently under investigation. For the purposes of this simulation study we adopt the random way-point model in the mobility pattern of the MN to encourage comparability of our results with other simulations studies.

Simulations were conducted for a maximum number of up to 120 MNs. Each simulation run had a duration of 10 minutes; this is a realistic figure given the above coverage area range and MN speed. Each Routing Neighbourhood Cache entry has a lifetime 150 sec.

All coverage areas maintained a homogeneous range (i.e. no obstructions were considered that could incur a change in the effective coverage area of a cell).

H_{min}, H_{res} and H_{off} are also signal levels (approximate distances) configured at the Routing Agent (controlling respective APs) to trigger pruning and grafting of CoAs onto the multicast tree that enables transmission over the Handoff CoA (HCoA).

Duplicate packets experienced were not considered as loss but were simply discarded by the MN.

Transmission of real time flows were assumed and configured to be unidirectional with a bandwidth of 3.2 Kbps (6.7 packets/s) with a packet size of 60 bytes. The size of real-time IP flow was extremely small compared to typical ranges in real-world interactive IP transmissions of 10 Kbps to 64 Kbps. This is because of the size of the simulation (maximum number of nodes) in NS-2. When running simulations with high bandwidth IP flows, MNs are often not even able to get any connectivity although they are within the service range of the AP. We observed that control signalling frequently gets lost if IP flows are configured to send at much higher rates than the one used in our experiments. Erroneous loss of signalling prevents the MN from performing an IP handoff to the next AP or refresh its bindings to maintain live IP connectivity.

The simulations assume that MN capacity does not impede the performance for both mobility protocols on the AR agent.

V. Simulation Results

From our simulations we show the observed performance gain of Proactive IPMM over standard IP Mobility. This is shown by comparing and contrasting the observed handoff efficiency (HE) of proactive IPMM and standard Mobile IP (v6) over the instantaneous handoff rate of MN as experienced at the AR. We note that the instantaneous handoff rate does **not** describe the speed of the MN although it is affected by it; it describes the number of handoffs effected at the AR from multiple MNs. Hence handoff rate does not convey a measure of the MN speed.

It is for this purpose that each handoff had to be traced at the AR not only by means of its *effect* but also by means of its *affect* on loss during the transmission of the IP flow onto the MN (unidirectional). For this reason, an extra signal was implemented (orthogonal to the Proactive IP mobility model) to report to the previous AR (agent) whether the handoff effected was *clean* (i.e. loss free) or not. The implementation

[3]In fact the realism pertains to better representation of the ARPANET - the ancestor of Internet

[4]random way-point allows lowest bound performance over a mobile mechanism during the mobility pattern of an MN

of such signal encompassed a check of packet sequence numbers from the moment of HCoA initiation for Proactive IPMM and H_{off} trigger for standard IPMM until the dispatch of the binding update signal.

A. Performance gain of Proactive over standard IPMM

Figure 6 illustrates the performance benefit during proactive IP handoff versus standard IP handoff management. Proactive IP handoff management yields significant handoff efficiency gains over the standard reactive IPMM mechanism, for handoff rates as low as 0.5 h/sec and for a topology of 30 ARs. As the handoff rate increases to 2 h/sec, around 0.8 of handoffs attach successfully to the new AR without packet loss during the real-time transmission. HE reaches over 0.92 as the instantaneous handoff rate climbs to 6 h/sec. On the contrary, base IP Mobility exhibits a much lower handoff efficiency of around 0.4 (fitted).

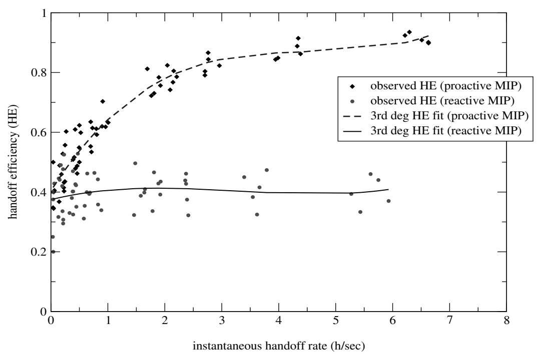

Fig. 6. Proactive versus reactive handoff efficiency in IP Mobility

This is justified simply, by the reactive character of the IP handoff: a movement of the mobile node is detected by means of lack of beacon frames from the current AP at the link-layer; multiple nodes are competing for access to effect a router solicitation so as to **initiate** re-establishment of IP addressing/routing state (CoA+BU) at the new point of attachment, in standard Mobile IPv6. On the contrary the Proactive IP mobility model has already configured its IP address/routing; this requires minimal interaction on the part of the MN; it simply senses at H_{min} an upcoming IP handoff and signals the HCoA-initiate to the current AR, to effect traffic delivery in the Mobility Neighbourhood over the HCoA address. From this point onwards (around H_{res}) the MN attempts to resolve the new AP attachment point by means of match-up of IP roaming state with incoming IPv6 router advertisements with no dependence on their transmission frequency.

The clustering of data is observed due to the fact that the instantaneous handoff rate is affected by the measure of speed of the mobile node, for which we have selected discrete values (i.e. 1/4/8 m/s); this is also due to averaging of multiple simulation runs for each permutation of the pair (topology,speed). Observations were drawn from a collective plot of all permutations on a single graph.

We should note that during the simulation runs for each (Access Point topology[5], speed) tuple permutations, we observed that the density of APs per AR and the speed of the MN did not have a significant influence on the performance of either proactive or standard IPMM models; this allowed averaging of all iterations for each permutation while focusing on the behaviour of the permutations with regards to their effect on the performance trend of either proactive or standard IPMM mechanisms.

The above derivation was arrived at by performing a goodness of fit test by means of regression. A third degree (cubic) regressive curve fitting was derived as shown in figure 6. The curve fitting with $r^2 = 0.1245$ is defined by the following function:

$$y = 0.403 + 0.297x - 0.0645x^2 + 0.00474x^3 \qquad (1)$$

Best fit was evaluated by means of calculating the residual[6] (r^2), which was found to be minimum for this curve. We remind that, the residual describes the standard error (regressor) from the mean (i.e. curve fitting).

To understand the relation between handoff efficiency of the protocol and the rate of MN handoffs at the AR we focus on the effect of RNV updates (indirect+direct) over the respective cache (RNV Cache) at the each neighbouring AR. An RNV Cache located in each AR within a Routing Neighbourhood, fills with adjacency references for their neighbouring ARs. The RNV Cache is populated by state supplied by MNs as they transit over the cell within their respective Mobility Neighbourhood (i.e. opportunistically).

As more IP handoffs occur (i.e. the handoff rate increases) at the AR, each of the RNV neighbour's cache converges closer to completeness (i.e. captures the entirety of its own routing neighbourhood), and thus MNs have more context state (in this case CoA roaming state) that allows a more robust handoff when the Handoff CoA (HCoA) mapping is activated over the set of CoAs, each allocated from the individual AR neighbours. The increase of handoff efficiency is thus, owed to the convergence of RNV state at the AR, the rate of which is tracked by the instantaneous handoff rate of MNs experienced at the AR.

To provide a measure of convergence on a random AR neighbour belonging to the same Mobility Neighbourhood in line with the observed performance of the Proactive IP mobility model, we track the ratio of the number of RNV Cache entries in an AR neighbour over the total number of AR neighbours mapping to the same mobility neighbourhood; this is the RNV Cache fill ratio metric measured during the mobility pattern of 15, 50 and 120 MNs; this is shown in Figure 7. The entire set of permutations is run for a maximum MN speed of 8 m/s (or 25 mph).

From the graph of Figure 7, we observe that the RNV Cache fill ratio encounters an upper bound in convergence that is tracked by distinct factors:

[5]we have stated in [19] that consecutive AP cells are not necessarily connected to ARs that are directly-reachable

[6]square to avoid negative values as seen at the graph

- the number of mobile nodes contributing to this convergence by means of indirect RNV-updates.
- the lifetime of the RNV Cache entry.

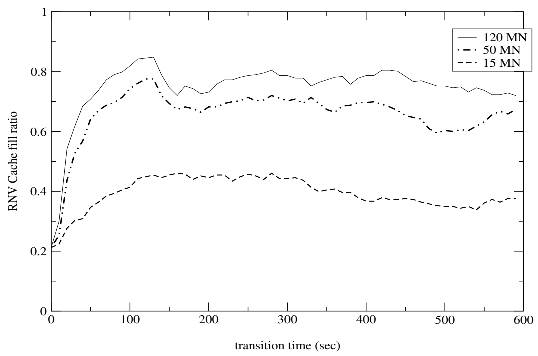

Fig. 7. RNV Cache fill ratio of a Routing Neighbourhood mapping to a single Mobility Neighbourhood

The smaller the number of MNs the lower the upper bound of the converged RNV Cache fill ratio for a random AR in the same mobility neighbourhood; for instance, for 15 MNs this upper bound is considerably lower than in the case of 50 or 120 MNs. This is in line with the slow growth of HE for small handoff rates, as seen in figure 6. Furthermore, the cutoff point on this upper bound in the convergence of the RNV Cache fill ratio is regulated by the lifetime of the RNV Cache entry. As the lifetime of the RNV Cache entry is being reached, the entry gets removed; this decreases the cache fill ratio since each AR neighbour has a narrower view of its mobility neighbourhood. At the same time, the small handoff rate incurred by low number of transiting MNs does not yield a sufficient number of RNV-updates so as to recuperate from lost state due to expired RNV entries and subsequently increase the RNV Cache fill ratio so that it reaches convergence (i.e. obtains a full view of its routing neighbourhood). This effect is diminished as the handoff rate at the AR increases (i.e an increasing number of MN cross-cell transitions). We remind that the lifetime of the RNV Cache entry is 150 sec; this lifetime size justifies the limited drop in the RNV Cache fill ratio observed around this time period of the simulation.

B. Protocol performance over varying RN size

Our experiments explored further the effect of variations of Routing Neighbourhood (RN) density over handoff efficiency; to do that, simulations were repeated for varying RNV size with range in the interval set [10,30]. Our initial expectation was that an increase of the RNV size should not affect negatively the handoff efficiency of the proactive IP mobility model while it should be orthogonal for standard Mobile IPv6.

Figure verifies 8 our intuition on variation of RNV densities as anticipated; the size of the Routing Neighbourhood, does not impact negatively the handoff efficiency for both methods, Proactive and standard reactive IPMM, with different effect for each mobility protocol.

For Proactive IPMM, Handoff Efficiency (HE) gains diminish for a sparse RNV population, for a decreasing handoff rate; Mobile IPv6 observes an indifferent, near-constant behaviour (roughly constant $HE \approx 0.4$) that is marginally affected by an increasing RNV population or handoff rate; this is due to the fact that there is no notion of Routing Neighbourhood in standard Mobile IPv6. On the contrary, for an increasing RNV population and handoff rate, the Proactive IP Mobility approach yields an increase in handoff efficiency.

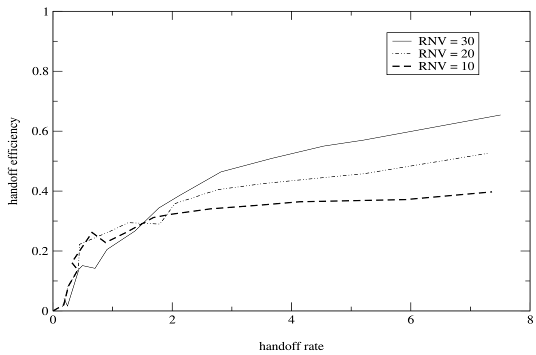

Fig. 8. Handoff Efficiency between Proactive and Reactive IP Mobility for increasing RN vector size

C. The effect of lifetime size on RNV Cache entries

For the duration of this set of simulations runs, the RNV Cache entries maintain soft-state through expiry times; however, as derived above, the lifetime of an RNV Cache entry each entry can affect the RNV Cache fill ratio incurred for a given handoff rate.

Intuitively, keeping the lifetime period of the cache entry small improves robustness in the face of reconfiguration, failures, or even malicious entries; on the other hand, setting the lifetime size of a cache entry to a long value (or possibly make it permanent), requires fewer handoffs to maintain relatively complete RNV state. It is, thus, essential to explore the effect that lifetime size can have on handoff efficiency, as observed in the performance gains attained over the Proactive IP mobility model; for this purpose, a second set of simulation runs was performed with lifetime size ranging from 30 to 350 sec, while MN were transiting at maximum speed (8 m/s); this was done to focus on potential variability of the performance gains in terms of handoff efficiency, attained by the Proactive IP Mobility model. In this manner we can assess factors that can potentially influence HE performance observed so far for Proactive IPMM over standard Mobile IPv6.

Figure 9 demonstrates the dependence of Handoff Efficiency on RNV Cache entry duration and the number of MNs transiting between AR neighbours. We can observe the steep increase of Handoff efficiency for any lifetime size greater than 150 while the number of MNs increases up to around 45 MNs (HE is around 0.8). Although the handoff efficiency improves continually with an increasing lifetime size for RNV entries, the improvement becomes marginal for lifetime size $t_l = 200$sec, irrespective of any further increase of MNs (HE improvement differential only 0.1-.15).

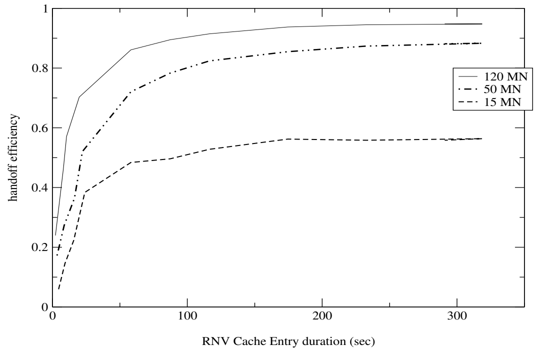

Fig. 9. Handoff efficiency and its dependence oh RNV entry lifetime

We can further observe that any attainable improvement on Handoff Efficiency is more dependent on the number of the MNs at least for up to 50 MNs and while the RNV Cache entry lifetime is up to 170 sec. We can thus confirm that large RNV Cache entry lifetimes do not yield significant HE gains unless the size of MN population begins to increase.

D. Proactive Link State notification and Vertical IP Handoffs

To demonstrate the applicability of proactivity in conveying forward information relating to service/performance characteristics of multiple IP handoff candidate ARs, small simulation extensions were devised for AR (agent) neighbours. The extensions pertained to the exchange of link-layer capability state of heterogeneous APs attached to their respective AR as a member of the Routing Neighbourhood of the MN.

Heterogeneity amongst connecting APs (BaseStationNode) was introduced by implementing a capability parameter that differentiates between available signalling rates of APs in the Mobility Neighbourhood; all APs maintained a 802.11 MAC layer. Two distinct bandwidth capabilities were implemented for participating APs: 11 Mbps and 5.5 Mbps. Each AP was allocated randomly one of the two bandwidth states. Context-specific capabilities are exchanged between individual ARs through direct RNV-updates and are pushed to the MN by means of an extended Context State Push (CS-Push) message. Exchange of such mobility signals was assumed to be error-free.

The mobile node implemented further a second 802.11 WLAN interface which, based on context transfer information that matches up the received beacon (which has been extended to announce reactively the bandwidth capability of the link) can switch from the operating 802.11 WLAN interface with zero link-switching latency. For the scenario of reactive detection of the capability and the corresponding switching between interfaces our extensions inject an artificial link-switching delay of 300 ms. The selection of such delay figure is not concerned with the exact behaviour of the link layer; this simulation experiment pursues to verify that given sufficient context state that can aid selection over multiple IP handoff candidates, the notion of proactivity in IPMM can provide significant gains towards intelligent IP mobility management for the MN on the move; such gains are not possible in reactive IP mobility management models.

Intelligence in the case for this simulation experiment is effected in the form of multi-access between heterogeneous APs with no rate-shifting capabilities; such wireless points of attachment, while providing a continuous coverage span, they do not encompass the same performance characteristics. This is the case for disparate service providers which while effecting coverage over a specific wireless medium, its characteristics are so diverse that IP Mobility cannot be effected without loss of IP connectivity transparency. Intelligence can be extended to cater for adaptations in the real-time IP flows transmitted by the MN. However extensions to simulate such behaviour add significantly to the complexity of the simulated model and for this reason have been considered out of scope for the purposes of this paper.

It should further be noted that the above rationale assumes a simplified operational scenario where the notion of administrative domain boundaries does not exist; this is done so for reasons of simplifying the implementation of the simulation as well as focusing on protocol gains rather than administrative Internet externalities.

Figure 10 shows a comparison between use of proactive context transfer of link state[7] and reactive detection of such link state during an MN transition upon a proactive IP handoff, against the Handoff Efficiency for the maximum MN population (120 MNs) while travelling at maximum speed (8 m/s). This simulation scenario was iterated for 10 runs with each observation signifying a 10-run average. No error bars have been included in the graph, as the distribution of observed values did not deviate significantly[8] from the mean.

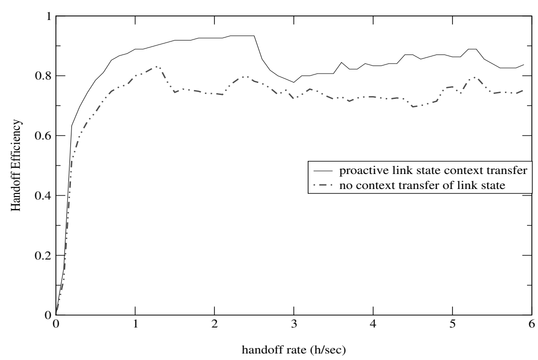

Fig. 10. Handoff Efficiency for Routing Neighbourhood supporting proactive link state notifications

We may observe that in advance transfer of link bandwidth context state does incur an increase in the observed handoff efficiency over Proactive IP mobility signalling deliberations. The observed gain between such proactive link state transfer and its reactive detection counterpart is not very big (about 0.1 maximum) while we observe sudden drops of HE that tend to stabilise as the handoff rate observed at the AR increases.

The drop in handoff efficiency is incurred due to the lifetime size (150 sec). Nevertheless, despite the expiry of RNV Cache entries, in-advance link state notifications do

[7]for this experiment simply the type of link layer
[8]if required they can be included

exhibit a performance gain over reactive detection of such capability. The size of RNV Cache entry lifetime is not expected to affect adversely handoff efficiency since as the handoff rate increases, more MN tend to refresh RNV update to the attaching AR at a higher rate respectively.

E. Protocol performance over signalling losses

While the above provide a performance view of the Proactive IP Mobility approach under error-free conditions, the efficiency of any IP mobility protocol is tracked by its performance over signalling losses; this is so because the reliability of the wireless medium is inherently a function of the signal to noise ratio (SNR) experienced by the MN in communications with its AP. For this reason the performance of the proposed Proactive IPMM model was evaluated over lossy signalling for Context-Push and HCoA-initiate messages.

To capture such behaviour, we drop randomly Context-Push and HCoA-initiate messages with a probability inversely proportional to the radius r from the edge of coverage $P(x) = \frac{\alpha}{r^2}$; α is the random variable affecting the prescribed loss probability, given a fixed AP cell radius. The loss probability values were explored and then plotted against the observed handoff efficiency for Proactive IP mobility and for a maximum handoff rate of 8 h/sec. Figure 11, focuses on the observed handoff efficiency for discrete loss probability values of $P(x) = 0.2, 0.4, 0.5$. It can be seen that high error rates affect negatively the handoff efficiency of the Proactive IPMM mechanism for an increasing handoff rate. This is owed to an increase in the *apparent* speed of the MN which limits the time-window of lost message retransmission before handoff becomes critical.

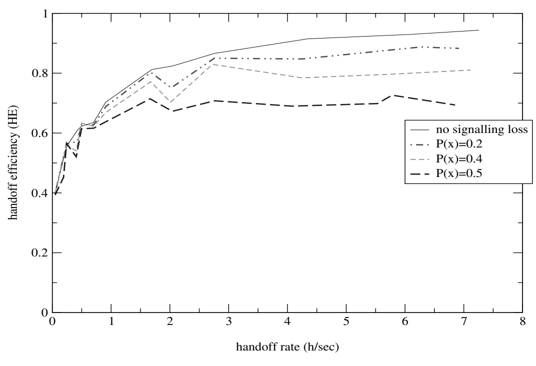

Fig. 11. 2-d view of Handoff efficiency for lossy RNV-updates

We emphasise on *apparent* speed, since, loss of signalling while on the move, presents the MN to any IPMM mechanism as travelling at much higher speed than its physical one. For it we argue that while the speed of the MN may be within bounds of the performance of IPMM, signalling losses of the IPMM protocol can introduce artificial fluctuations in the physical speed of the MN effecting an increased apparent speed over which the performance of the proactive IP mobility protocol appears to be degraded. While the degradation in performance becomes at worst equal to the performance experienced in standard IPMM (due to lack of proactive stimuli), it becomes essential that control signalling that effects proactivity in IP mobility management for the purposes of IP handoff management must be assured between the MN and the current Access Router.

VI. Related work

Recently, several protocols have been proposed for next generation (IPv6) wireless access networks in support of seamless mobility. From the perspective of Internet specifications, Fast Handoffs [26] (FMIPv6) is an extension to the standard IPMM model based on Mobile IPv6 and is currently under development. The FMIPv6 model maintains some similarities with Proactive IPMM with respect to the concept of pre-registration of the MN with other ARs that may be candidates for an IP handoff. Our proactive mobility model is significantly different from the FMIPv6 proposal for a number of reasons; our mobility model does not employ proxy neighbour discovery messaging since tasks like duplicate address detection (DAD) mandated by IPv6 Neighbour Discovery standard, require additional signalling which also induces delays of at least one RTT between the MN and *every* candidate AR.

Furthermore, our model caters for a candidate access router discovery algorithm. The FMIPv6 proposal provides no such mechanism; this is an essential part for a seamless IP mobility model. In addition, our model provides a signalling substrate for proactive context transfers from AR neighbours; the FMIPv6 proposal is explicitly designed towards expediting IP handoff process only. It requires additional protocol mechanisms for context transfers [27] currently under development in the SEAMOBY WG of the IETF. This implies that under FMIPv6 (or standard IPMM for that purpose) two (independent) protocols are essential for context-aware IPMM[9]; the proposed Proactive IPMM investigated provides by design tight coupling of context transfer as part of its proactive signalling deliberations. What's more, the fast handoff proposal lacks of robustness with respect to cell-bouncing effects (a.k.a ping-pong); this is because it assumes existence of multiple tunnels towards the MN from candidate ARs. On the contrary, our scheme employs only a single tunnel which is effected over multicast; this ensure that the MN can *freely* bounce between any AR within the mobility neighbourhood.

Alternative approaches in Mobile IP have also been proposed in [28], [29], [30]. These schemes employ IP-Multicast for addressing and forwarding of packets to MNs on an end-to-end basis; i.e. the source is destined at either the CN or the HA while MN is the receiver for the purposes of minimal latency handoffs. The Proactive IPMM is architecturally different from all the above schemes. In the majority of the mobility proposals, IP multicast is employed either end-to-end or as a pseudo multicast mechanism where the protocol abstracts many unicast destinations as a source at the CN or HA. In the proposed mobility model IP multicast is employed **locally** with respect to the location of the mobile node (i.e. mobility neighbourhood). As such, none of the above schemes proposes a mechanism of establishing such mobility

[9]it is common place that for protocol independence lacks of performance of tightly coupled signalling mechanisms

neighbourhood. Furthermore, there is no consideration in any of the proposed mobility mechanisms for state relocation with respect to multiple mobility contexts. Our abstraction of IP Roaming state as a representative class of relocatable state context allows to populate the Context State cache with multiple contexts/capabilities that neighbour ARs may need to provide to the CURRENT AR towards the seamless movement of the MN in its mobility neighbourhood.

VII. Conclusions

This paper presented an experimental assessment on the behaviour and performance of transport protocols such as TCP during IP handoffs of the MN between two heterogeneous wireless networks, namely GPRS and WLAN. Based on real-world measurements we established wireless-link characteristics affecting the IP handoff process in terms of transmission delay as the MN transits between heterogeneous networks.

We further identified the need of dynamic establishment/relocation of IP related context-state. This would serve as a vehicle to exploit knowledge of particular context-state, so as to inform the IP handoff process in advance of its occurrence towards potential adaptations. To this end, we have investigated, by means of simulations, the Proactive IP Mobility Management model as a mechanism of enabling context-awareness to assist the IP handoff decision (vertical or horizontal). Our findings argue in favour of the hypothesis that reactive IP Mobility Management, based on Mobile IPv6, is not sufficient for handling advanced aspects of IPMM such as provisioning of context-aware information to the MN to enable proactive informed IP handoff decisions in wireless multi-access heterogeneous networks.

From the simulation results derived we have arrived at a number of conclusions with respect IP handoff performance in IPMM; in particular we have shown that the handoff efficiency for Proactive IPMM is significantly better to that of current IPMM standards based on Mobile IPv6. The AP topology does not affect significantly this result. The above claim is valid for an increasing handoff rate experienced at any AR within a Routing Neighbourhood.

Handoff efficiency gains of Proactive IP Mobility are directly dependent on the degree of completeness of the Routing Neighbourhood vector for each Mobility Neighbourhood; the latter is in turn, dependent on the amount of bypassing MN as they transit between ARs within a Mobility Neighbourhood, given limited RNV entry lifetime. Handoff Efficiency gains for Proactive IPMM diminish for a decreasing size of the Routing Neighbourhood when the observed handoff rate at the AR decreases. The main reason for this is the expiry of RNV entries and the lack of quick refreshes since not very many MN attach on some neighbouring AR.

Handoff efficiency improves for an increasing RNV Cache entry lifetime; the effect is magnified for an increasing number of MN. Later expiry of RNV Cache entries prolongs their validity and hence sustains a higher RNV Cache fill ratio for longer periods.

Proactive capability context-state transfers enhance the handoff efficiency of the Proactive IP mobility model. Forward awareness of link-layer capabilities of future IP attachment points as context state of interest to the MN, demonstrated the significance of context-aware IP handoffs by means of enabling that network interface with the highest bandwidth provisioning. Given a different context such as billing, similar gains may be obtained by effecting an context-aware IP handoff to the IP attachment point with the most competitive tariff.

References

[1] H. Wang, J. Giese, and R. Katz, "Policy-Enabled Handoffs Across Heterogenous Wireless Networks," in *Proceedings of 2nd ACM/IEEE Workshop on Mobile Computing Systems and Applications (WMCSA)*, New Orleans, LA, Feb. 1999.

[2] H. Yokota, A. Idoue, T. Hasegawa, and T. Kato, "Link Layer Assisted Mobile IP Fast Handoff Method over Wireless-LAN Networks," in *Proceedings of ACM/IEEE International Conference on Mobile Computing and Networking (MOBICOM)*, Atlanta, Georgia, Sep. 2002.

[3] S. Mukkamalla and B. Raman, "Scaling and Latency Issues in Mobile-IP," http://www.cs.berkeley.edu/ãdj/cs294-1.s98/projects/MobileIP/sld001.htm, Apr. 1998.

[4] S. Deering and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification," RFC 2460, Internet Engineering Task Force (IETF), Dec. 1998.

[5] C. Perkins and D. Johnson, "Mobility Support in IPv6," Internet draft (work in progress), Internet Engineering Task Force (IETF), May 2003.

[6] C. Carter and R. Kravets, "Contact Networking: A Localized Mobility System," in *Proceedings of 1st International Conference on Mobile Systems, Applications and Services (MobiSys)*, May 2003, pp. 145–158.

[7] R. Katz and E. Brewer, "The Case for Wireless Overlay Networks," in *Proceedings of SPIE Multimedia and Networking Conference (MMNC)*, San Jose, CA, Jan. 1996.

[8] M. Ylianttila, M. Pande, J. Makela, and P. Mahonen, "Optimisation Scheme for Mobile Users performing Vertical Handoffs between IEEE 802.11 and GPRS/EDGE Networks," in *Proceedings of IEEE Global Communications Conference (GLOBECOM)*, Texas, USA, Nov. 2001.

[9] T.B. Zahariadis and N.A. Nikolaou, "Unified Wireless Access in Hot-Spot Environment," *IEEE Communications*, vol. 6, no. 6, pp. 259–261, Jun. 2002.

[10] R. Caceres and V. Padmanabhan, "Fast and Scalable Wireless Handoffs in Support of Mobile Internet Audio," *ACM Journal in Mobile Networking and Applications*, vol. 3, no. 4, pp. 351–363, Jun. 1998.

[11] T. Narten, E. Nordmark, and W. Simpson, "Neighbor Discovery for IP Version 6 (IPv6)," RFC 2461, Internet Engineering Task Force (IETF), Dec. 1998.

[12] T. Pagtzis and P. Kirstein, "A Model for Proactive Seamless IP Mobility and Mobility-Hop Routing," in *Proceedings of 10th IEEE International Conference on Networks (ICON)*, Singapore, Aug. 2002.

[13] T. Pagtzis, P. Kirstein, and S. Hailes, "Proactive Seamless Mobility Management for future IP Radio Access Networks," *Computer Communications*, vol. Special Issue on '3G Wireless and Beyond', pp. (to appear), Jul. 2003.

[14] A. Yegin, M. Tariq, A. Yokote, et al., "Mobile IP API," Internet draft (work in progress), Internet Engineering Task Force (IETF), Jun. 2003.

[15] R. Chakravorty, S. Katti, I. Pratt, and J. Crowcroft, "Flow Aggregation for Enhanced TCP over Wide-Area Wireless," in *Proceedings of IEEE Conference on Computer Communications (INFOCOM)*, San Francisco, CA, Mar. 2003, http://www.cl.cam.ac.uk/users/rc277/gprs.html.

[16] P. Vidales, Patanapongpibul L., and R. Chakravorty, "Ubiquitous Networking in Heterogeneous Environments," in *Proceedings of International Workshop and on Mobile Multimedia Communications (MoMuC)*, Oct. 2003, (to appear) http://www.cl.cam.ac.uk/users/rc277/overlay.html.

[17] R. Chakravorty, P. Vidales, L. Patanapongpibul, I. Pratt, and J. Crowcroft, "Vertical Handovers with Mobile IPv6," Technical Report CAM/CL-TR-03-05, University of Cambridge, Cambridge, UK, May 2003.

[18] M. Stemm and R. Katz, "Vertical Handoffs in Wireless Overlay Networks," *ACM Journal in Mobile Networking and Applications*, vol. 3, no. 4, pp. 335–350, 1998.

[19] T. Pagtzis and P. Kirstein, "A Framework for Seamless Proactive Mobility in Mobile IPv6," Tech. Rep., Internet Engineering Task Force (IETF), Jun. 2001, http://www.cs.ucl.ac.uk/staff/t.pagtzis/papers/draft-pagtzis-mobileip-proactivev6-00.ps.

[20] "The Network Simulator - NS-2," http://www.isi.edu/nsnam/ns/, 1989.

[21] "MobiWan: NS-2 extensions to study Mobility in Wide-Area IPv6 Networks," http://www.inrialpes.fr/planete/pub/mobiwan/, 2002.

[22] 802.11 WG, *IEEE 802.11b High Rate Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specification Standard*, IEEE, Oct. 2000.

[23] G. Bianchi, "Performance Analysis of the IEEE 802.11 Distributed Coordination Function," *Journal on Selected Areas on Communications*, vol. 18, no. 3, pp. 535–547, Mar. 2000.

[24] G. Anastasi and L. Lenzini, "QoS provided by the IEEE 802.11 Wireless LAN to Advanced Data Applications: a Simulation Analysis," *ACM/Baltzer Wireless Networks Journal*, vol. 6, no. 2, pp. 99–108, Jun. 2000.

[25] T. Pagtzis, "Link-switching characteristics for IEEE802.11 Wireless LANs," Technical Report UCL/CS-TR-02-165, University College London, London, UK, Oct. 2002.

[26] R. Koodli, A. Yegin, C. Perkins, G. Dommety, and K. El-Malki, "Fast Handovers for Mobile IPv6," Internet draft (work in progress), Internet Engineering Task Force (IETF), Oct. 2002.

[27] J. (Editor) Kempf, "Reasons For Doing Context Transfers Between Nodes in an IP Access Network," RFC 3374, Internet Engineering Task Force (IETF), Sep. 2002.

[28] S. Seshan, H. Balakrishnan, and R. Katz, "Handoffs in cellular wireless networks: The Daedalus implementation and experience," 1996.

[29] A. Helmy, "A Multicast-based Protocol for IP Mobility Support," in *ACM 2nd International Workshop on Networked Group Communication (NGC)*, Nov. 2000.

[30] J. Mysore and V. Bharghavan, "A New Multicasting-based Architecture for Internet Host Mobility," in *Proceedings of ACM/IEEE International Conference on Mobile Computing and Networking (MOBICOM)*, Sep. 1997.