Prediction Assisted Fast Handovers for Mobile IPv6

Andre E. Bergh and Neco Ventura
Department of Electrical Engineering, University of Cape Town, South Africa
{andre, neco}@crg.ee.uct.ac.za

Abstract— **Achieving seamless mobility is a major challenge for future wireless networks. Presently, low latency handover protocols such as FMIPv6 depend heavily on link-layer triggers to facilitate proactive handovers. Differences in link-layer technologies and the absence of a standardised trigger model makes it difficult to achieve fast vertical handovers in a multi-access environment. This paper introduces a prediction assisted fast handover protocol (PA-FMIP) that does not rely on pre-triggers. A data mining approach is used to implement the prediction algorithm based on the user's mobility history between wireless subnets. The performance of the proposed handover scheme is compared to MIPv6, proactive FMIPv6, reactive FMIPv6 and Simultaneous Bindings. Simulation results show a 43% decrease in handover packet loss over reactive FMIPv6 and an improvement in TCP goodput.**

I. INTRODUCTION

Fixed mobile convergence is a key trend in the telecommunications industry. New developments in technology are extending mobile services from frequency regulated cellular networks to include the unlicensed spectrum. This is allowing service delivery platforms such as the IP Multimedia Subsystem to reach users in their homes and offices. Users can connect to cellular networks intermediately using a number of short range wireless mediums such as 802.11x and Bluetooth. The expanding service delivery together with an increasing demand for seamless mobility and Quality of Service (QoS) create new challenges for IP networks. The growing number of wireless mobile devices is testing the philosophy of being always connected, anywhere and anytime. Mobile IPv6 has been declared unsuitable for managing the mobility of end users running real-time applications [13]. Seamless IP mobility is looking toward hierarchical network topologies and fast handover protocols for support. A common feature among researchers is the idea of *proactive* service establishment. Transferring context, pre-negotiating network addresses and forwarding packets prior to a subnet change significantly reduces service disruption.

Fast Handovers for Mobile IPv6 (FMIPv6) [13] and other related schemes require link-layer (L2) triggers to effect low latency handovers. FMIPv6 with proactive handover (proactive FMIPv6) uses a L2 pre-trigger to notify the network layer (L3) of an impending L2 handover, so that it may pro-actively obtain a new care-of-address (CoA) and start forwarding data to the next access router (nAR) before the handover. Some access technologies require modification to support pre-triggers and implementations are also not standardised, often requiring an accurate configuration of trigger timing to achieve low latency handovers. The work presented in this paper aims to provide fast vertical handovers over heterogeneous networks (e.g 802.11x, Bluetooth, UMTS, CDMA2000) by reducing the handover requirements on the link layer, and improving vendor interoperability.

In this paper, a fast handover scheme (PA-FMIP) is proposed. It is based on the reactive FMIPv6 handover protocol (reactive FMIPv6). In PA-FMIP, the function of a pre-trigger is replaced with a simple (modular) mobility prediction application used by the mobile node (MN) to predict its next AR attachment. The prediction is based on a record of its network-layer previous mobility. The current AR tunnels the MNs incoming data to the predicted AR, in a manner similar to proactive FMIPv6, resulting in a significantly smoother handover as compared to reactive FMIPv6. The efficiency of PA-FMIP is affected by the number of required predictions per hop (M). The moderate increase in network bandwidth utilization due to packet forwarding, is considered to be an acceptable trade-off for a faster and smoother vertical handover.

The rest of the paper is organised as follows: A review of literature on some important aspects related to this work is given in Section II. Section III presents a broader description of the PA-FMIP proposal. Section IV describes the ns2 simulation experiments and discusses the important results. Section V concludes.

II. RELATED WORK

This research consists of two sections, namely the prediction of user mobility at the network level, and fast handovers based on reactive FMIPv6.

A. Mobility Prediction

Network-layer mobility prediction schemes that are designed to facilitate pro-activity in the handover process[1] also aid in resource allocation, flow control, call admission control, congestion control and QoS provisioning. Lui and Maguire [14] are pioneers in mobility prediction through their Mobile Motion Prediction (MMP) algorithms used to predict future locations of a mobile user according to the user's movement history patterns. This was one of the first of many techniques in the literature used to pro-actively connect services at the new location before the user's arrival. The MMP algorithms are based on the fact that human movement generally consists of regular and random movement. These algorithms use correlation analysis to match movement sequences in a movement database. Results show that the MMP algorithm is highly accurate for regular movements but decreases linearly with increasing random movement.

Yavas proposes a novel data mining approach for the prediction of user mobility in mobile environments [6]. He outlines a three stage prediction algorithm based on the

---

[1] Unless specified otherwise, a handover refers to both layer-2 and layer-3

Apriori algorithm [1][2] and the web-prefetching algorithm by Nanopoulos in [17][18], to predict the mobility of a user travelling between the cells of a PCS (Personal Communication System) network. Simulation results reveal optimal prediction parameters for the PCS topology. Moderate prediction accuracy was achieved, decreasing only minimally with an increase in random mobility. The authors focus primarily on prediction recall and precision results, and make no practical use of the movement predictions. We build upon this work and integrate it into our fast handover proposal.

B. Fast Handovers

Fast Handovers for Mobile IPv6 [13] defines the protocol signalling, L2 trigger requirements and packet forwarding for proactive and reactive handovers. The decoupling of L2 and L3 handovers is a useful means of improving the inherent service disruption latency present in their combination. Although the proactive FMIPv6 handover has certain implementation issues, as stated by Ivov [12], its main advantage over reactive FMIPv6 is the ability to forward data to the nAR during the L2 handover period. When implemented correctly, this significantly reduces packet loss and service disruption, producing more seamless inter-subnet mobility than reactive FMIPv6.

Shim et. al. [21] propose NeighborCasting, a pre-trigger MIPv4 based low latency handover scheme that simply multicasts the MNs incoming data streams to *all* its neighboring ARs during a handover. Performance results indeed show a low latency handover but with significant overhead due to the aggressive multicasting.

Hsieh [11] performs a comparison of five current fast handover techniques, including an original Seamless handover (S-MIP) proposal. He evaluates the performance of each handover[2] scheme through simulation, and discusses their impact on end-to-end TCP applications. S-MIP shows the best results, however it requires a network entity called a Decision Engine to determine when and how the MN is to handover, depending on network conditions and movement patterns.

C. Combined Mobility Prediction and Fast Handovers

Feng et. al. [7] propose a prediction scheme based on the MMP algorithms [14] that uses actual mobility traces taken from the campus-wide 802.11b wireless network. Data streams are duplicated and forwarded to predicted subnets resulting in a network-layer handover latency that is close to a link-layer handover. Results show a reduced handover latency and packet loss rate compared to NeighborCasting.

Blondia [5] proposes a prediction mechanism that learns the mobility patterns of a mobile node according to an urban mobility model. The model attempts to capture realistic node movement in an urban environment, characterised by the MN's speed, direction, pause time and street coordinates. A weighted road selection process uses these parameters to predict the node's next hop, pre-emptively setting up tunnels and estimating tunnel activation times, consequently eliminating the need for a pre-trigger. This approach achieves 100% prediction accuracy only after 3000 seconds of mobility over a small Manhattan-style street topology.

III. PREDICTION ASSISTED FAST HANDOVERS

A. Data Mining background

Data Mining can be defined as the process of analysing data to identify patterns, associations, significant structures or relationships, from information stored in a data repository or database [6].

Sequential Pattern Mining (SPM), a subset of Frequent Itemset Mining, is the process of extracting certain *sequential* patterns whose support exceed a predefined minimal support threshold (min_{supp}) [22]. A minimum support threshold is a pruning mechanism used to reduce the number of sequences to a certain level of interest and make the process more efficient. SPM is used in business to study customer behaviour, in telecommunication networks to analyze system performance, stock market trend analysis and even in DNA sequencing.

Association Rule Mining (ARM) [1] is one of the most important and well researched techniques of data mining. It aims to extract interesting correlations, frequent patterns or associations among sets of items in the transaction databases. It is noted here that sequential patterns indicate the correlation between transactions, while association rules represent intra-transaction relationships. Furthermore, ARM does not consider the ordering of the items in a transaction. A common example of ARM is found on Amazon.com. While browsing for a specific product, the site often displays information similar to "customers who bought this product also bought ..." This association makes no use of which product was bought first, however it implies that these products were bought during the same transaction (session). Most algorithms used for sequential pattern and association rule mining are all variations based on the Apriori Algorithm [1]. Each attempts to reduce the number of times the database is scanned, and explore different candidate pruning techniques so as to minimise computational time, space and memory. The Apriori algorithm, in [3], is a level wise algorithm and makes multiple passes over the data to discover large (frequent in terms of support) sequences. These sequences, called candidates, range from length *1* to length *k*. Supports for these candidates are counted at each pass. The largest candidates of length(*k-1*) are then used to determine the candidate set for length(*k*) during the next pass. This process repeats until no new large sequences are found. Note that random items distributed within the database are essentially filtered out as they do not earn enough support, granted the value of min_{supp} is not too low. Our approach to mobility prediction is based on the work by Yavas [6] and a simple implementation of the Apriori Algorithm by Bodon [3].

B. Prediction Algorithm

Consider a transactional database, *D*, containing a log of a node's mobility history. A mobility history is created whenever a node performs a handover to a new access router (nAR). The ID of this nAR is appended to the history (*D*), gradually building up a set of mobility transactions. Consecutive IDs in a transaction form a trajectory, and represent the mobility between neighbouring cells in a network. The respective ARs may not necessarily be spatial neighbours or even routing neighbours; we describe them as handover neighbours. Tracking and predicting logical network-layer

---

[2] Hsieh only evaluates proactive handovers

mobility can be argued to be simpler and more predictable than geographic mobility. Geographic mobility depends on many restricting factors such as the physical terrain, obstacles, and the placement of radio access points, making mobility prediction reliant on external hardware.

A new transaction is started every T seconds to indicate the beginning of a new mobility session. The last entry in the database identifies the node's current position on the network, say p_n (see Table I). Applying the Apriori algorithm to D, with a min_{supp} of say 10%, a set of the most frequent sequential patterns is obtained and ranked in descending order of support. All such frequent patterns exceed the min_{supp}value of 10%. The next step in the prediction algorithm requires the generation of association rules from these sequential patterns. First, the rule mining algorithm in [9] is modified to take the ordering of items inside a transaction into account. Applying this modified algorithm to the frequent pattern set, one is left with *mobility rules* [6]. Each rule is coupled with a confidence value (*conf*), which is the conditional probability of the term's support values [8], e.g. the confidence of the rule X=>Y is $P(Y|X) = \frac{support(X \cup Y)}{support(X)}$. The minimum confidence of this set is limited by a pruning parameter min_{conf}.

TABLE I

EXAMPLE ILLUSTRATING THE STRUCTURE OF THE PREDICTION PROCESS.

Database D	Frequent Patterns	*Supp.*	Mobility Rules	*Conf.*
1,2,3,6,7	<2,3,6>	w	1,2=>3	95%
2,3	<1,2,3>	x	2,3=>6	91%
6,7,8,1	<6,7>	y	1=>2	90%
...	...	...	...	...
$..p_{n-3},p_{n-2},p_{n-1},p_n$	$<..,c_{m-2},c_{m-1},c_m>$	$>min_{supp}$	$..,r_{i-2},r_{i-1},r_i=>r_{i+1}$	$>min_{conf}$

The final step in the algorithm searches the set of rules for items immediately before the arrow that matches the condition $r_i = p_n$. These matching rules are related to the node's current position and are collected and ranked according to confidence. A user-defined M number predictions can be outputted; in this case, the M most confident r_{i+1}values are selected as the mostly likely next-hop predictions.

C. Faster Reactive Handovers

Handover latency is the primary cause of packet loss and performance degradation for mobile nodes, especially for end-to-end TCP [11]. The common denominator in most fast handover solutions in the literature is FMIPv6. FMIPv6 makes full use of cross-layer triggering to mitigate the negative effect of a handover. As shown in Fig. 1, the mobile trigger, just one of three types of pre-triggers, is activated some period $t_{\text{proactive}}$ before the L2 handover [13]. During $t_{\text{proactive}}$ the node's L2 scans for a new AR, forms a new CoA and begins the fast registration process.

PA-FMIP uses the prediction(s) from the prediction algorithm to essentially replace the need for a pre-trigger. The prediction algorithm is run as an application at some time between handovers. It notifies the current AR (pAR) of its M predicted movements by sending an AR Notice message. This message is forwarded by the pAR to all predicted nARs. In Fig. 1, bi-directional tunnels are setup between points a and b, but only activated at point c. The acknowledgement of the AR Notice message completes the handover preparation. The link-down (LD) trigger indicates the start of the L2 handover and the instant the pAR begins forwarding the node's incoming data stream to the notified nARs. A link-up (LU) trigger indicates that the node has completed its association with a new AP. The node then begins the fast registration process by sending a Fast Neighbour Advertisement (FNA) to the nAR. Once the nAR receives a Fast Binding Acknowledgement (FBACK) from the pAR, the nAR delivers any en-route or buffered packets to the node. The node continues to receive its forwarded data with no interruption until it completes its home registration of its new CoA, upon which, the tunnel between pAR and nAR is deactivated. Context transfer usually occurs at point d in reactive handovers. Any delay here would directly affect the overall handover latency. PA-FMIP solves this by proactively transferring any context in the AR Notice messages.

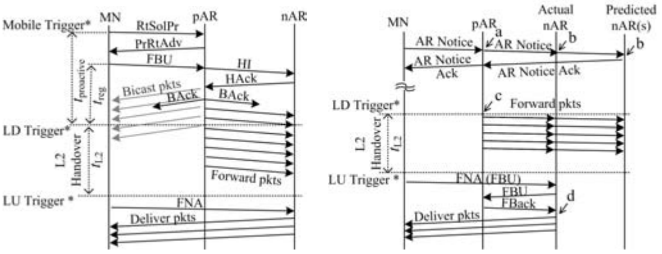

Fig. 1. Handover timing diagrams of 1) Proactive FMIPv6 and 2) PA-FMIP proposal.

IV. PERFORMANCE EVALUATION

In this section the performance of PA-FMIP is evaluated through two simulation experiments. Firstly, the accuracy of the prediction algorithm is assessed in order to derive an optimal value for the number of predictions per hop, M. And secondly, PA-FMIP is compared to similar fast handover protocols in UDP (User Datagram Protocol) and TCP (Transfer Control Protocol) application scenarios. All simulations are performed using ns2 [16].

A. Prediction Accuracy

Mobility models such as the the random waypoint model are unable to emulate regular mobility patterns. For this reason we only use the restricted random waypoint (RRWP) model [19] over a city section. Restricting mobility paths of the mobile node to a finite number of city streets creates a degree of regularity that we exploit. RRWP models are commonly used to model vehicular ad-hoc networks and vehicular mobility [20].

In this experiment, the mobility model generator in [19] is used to generate ns2 compatible RRWP mobility patterns. The model generator uses the TIGER (Topologically Integrated Geographic Encoding and Referencing) [4] database which contains selected geographic and cartographic information on road maps in the USA. This information is typically used to provide the digital map base for Geographic Information Systems or mapping software [20]. 36 access routers (ARs) are positioned across a 1200m by 1200m city section of West University Area in Houston Texas, which consists of 383 intersections and 594 road segments as shown in Fig. 2. These routers are connected with 100Mb

Ethernet links to form a large IPv6 network. Mobile IPv6 maintains the mobile node's connectivity as it moves. For simplicity, 802.11b is used with a uniform transmission range of 140m; however, we acknowledge that a multi-access type environment would certainly be more appropriate for evaluating vertical handovers.

The prediction algorithm proposed in section III-B is executed by the mobile node as an application at some period before each handover. To determine the accuracy of each prediction, the prediction result is recorded after each handover and compared to the ID of the new AR. The following assumptions are made in order to simplify this experiment and ensure repeatability:

- Only the mobility of a single MN is considered.
- Both AR and 802.11b AP are co-located.
- Each subnet has only one AR.
- The radio coverage areas of each AR are overlapping.
- A handover only occurs between neighbouring subnets.
- The MN has an initial mobility history totalling 30000 seconds of mobility over the topology.

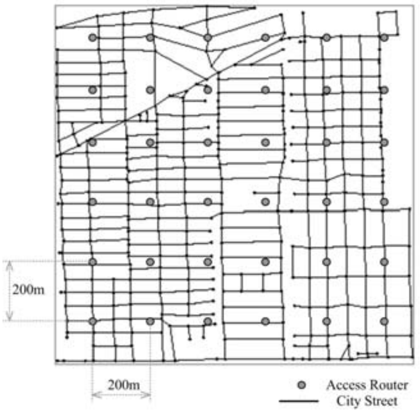

Fig. 2. Simulation topology - 1200x1200m city section of West University Area covered by 36 ARs.

The MN moves according to the following characteristics: Average speed = $5ms^{-1}$, max. speed deviation = $2ms^{-1}$, av. pause time = 0s. As stated, the node initially roams around the city for 30000s recording its points of attachment[3]. In doing so, it "learns" the topology, and creates the initial database needed by the prediction algorithm. The following simulation parameters are used: Simulation time = 2000s, min_{supp}=10%, min_{conf}=90%, and T=170s. The average number of neighbouring ARs is set to 8 even though the ARs on the perimeter only have between 3 and 5 neighbours. The measured accuracy for each value of M is the average of 10 sets of recorded simulation data.

During each 2000s simulation an average of 25 handovers occurs. Simulation results (shown in Fig. 3) for this experiment reveal an expected increase in accuracy as M increases, with a minimum of 39.05% for M=1. The main objective of this experiment is to obtain an optimal value for M, as M defines the number of new AR targets and therefore affects the neighbouring link bandwidth utilisation (see Fig. 6). We find that the MN should maintain a value of $M \geq 4$ to ensure a 92.5% probability of a prediction hit.

The factors most influencing the accuracy of this system are the database and the size of the topology. A prediction on the next AR will not be accurate unless the MN has previously traversed this path, implying that movement in smaller topologies is more predictable. The prediction accuracy of this algorithm is compared to NeighborCasting [21] in Fig. 3. It follows a linear increase in accuracy as M increases, with a minimum of $\frac{1}{8}$=12.5% . Its uninformed AR selection limits its potential accuracy and produces more redundant packet forwarding.

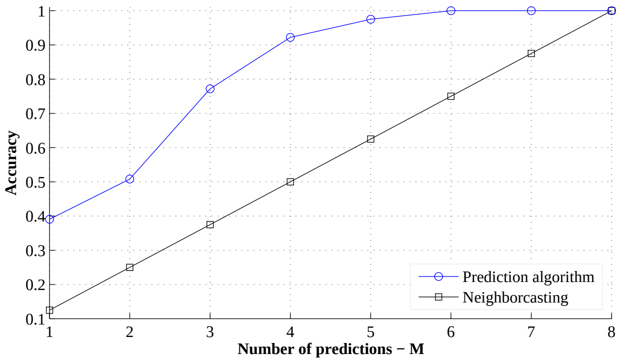

Fig. 3. Graph showing prediction accuracy vs. number of predictions M.

B. Handover Performance

This scenario is designed to investigate the behaviour of a MN as it moves between two access routers anywhere within the city section of Fig. 2. Hsieh released a Fast Hierarchical Mobile IPv6 extension [10] to ns2 which we adapted to implement the following fast handover protocols in ns2: MIPv6, reactive FMIPv6, proactive FMIPv6, the PA-FMIP proposal and Simultaneous Bindings. The handover performance of each are compared using packet loss and latency as metrics, over separate UDP and TCP applications. Fig. 4 shows the node layout, link properties and the MN's trajectory.

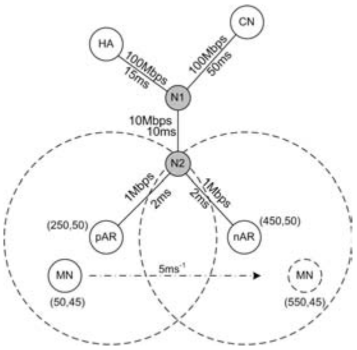

Fig. 4. Simulation topology for handover evaluations in ns2.

In an FTP (File Transfer Protocol) application between the MN and the corresponding node (CN), TCP-Tahoe is used with a packet size of 512 bytes and a window size of 32 packets. The MN begins a 3Mb file download from the CN 10 seconds after it begins to move toward nAR. At approximately mid-way between the two ARs, a handover

---

[3] 30000s is approximately equivalent to 154 transactions with an average transaction length of 4. This is long enough for the node to cover a large percentage of the topology.

occurs. The 802.11 link-layer handover is modelled as a constant 20ms delay (t_{L2}) for comparability. The handover latency is defined as the period of time between the first retransmitted packet and the last time this packet was sent [11]. In Fig. 5, this period is indicated as d_r for reactive FMIPv6 and d_{pa} for PA-FMIP. Table II summarises all the recorded measurements.

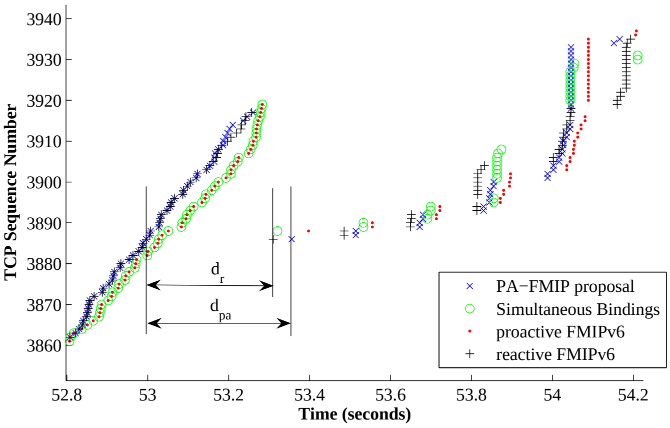

Fig. 5. A comparison of handover characteristics during a FTP application. Only TCP source profile shown.

An interesting result is immediately visible. PA-FMIP (d_{pa}) actually has a 41ms longer handover than reactive FMIPv6 (d_r), and 9ms longer than proactive FMIPv6. This is a direct result of the congestion effect of the large packet size and the packet forwarding between the pAR and the nAR on the wired links that decreases the arrival rate of the FMIPv6 control messages (e.g. FBU, FBACK, etc). A positive result of the packet forwarding is a significant decrease in the number of lost packets by PA-FMIP. A 43% decrease in packet loss is observed compared to reactive FMIPv6. The download results indicate that goodput performance is influenced more by packet loss than latency. Besides the slightly slower TCP handover and buffer usage due to the forwarding, PA-FMIP shows competitive goodput and packet loss figures next to proactive FMIPv6.

Proactive FMIPv6 however was expected to significantly outperform reactive FMIPv6. Its poor performance is attributed to the timing ambiguity problem [15]. Referring to Fig. 1, $t_{\text{proactive}}$ occurs 202ms before the link-layer handover. This is the time required by the MN to complete the CoA negotiation and home registration. Once the registration process has begun, the MN does not change links for another 62ms (t_{reg}) and it cannot receive packets from the pAR. This results in the increase in packet loss, each lost packet causing the TCP-Tahoe congestion control mechanism to halve the source's transmission rate, subsequently reducing the packet loss rate. Simultaneous Bindings [15] mitigates the timing ambiguity problem by bi-casting packets onto the pAR's wireless link as well as to the nAR. Thus allowing the MN to receive packets until the LD event. Bi-casting has a similar effect on the arrival rate of signalling messages [11] like PA-FMIP; however it is not as noticeable.

TABLE II

COMPARISON OF TCP GOODPUT FOR DIFFERENT FAST HANDOVER PROTOCOLS.

Handover Method	UDP		TCP				
	Handover Latency (ms)	Lost pkts	Handover Latency (ms)	Lost pkts	Buffer usage (pkts)	av. Goodput (Kbytes/s)	Download time (s)
MIPv6	257.8	45	1260	28	3	57.1158	53.7854
Reactive FMIPv6	41.92	6	314	16	10	58.4177	52.5868
Proactive FMIPv6	78.6	13	346	10	16	58.4989	52.5138
Prediction Assisted	33.8	3	355	9	18	58.4714	52.5385
Simultaneous Bindings	25.9	5	269	2	7	58.4929	52.5192
No Handover						60.4945	50.7815

In the next experiment, the FTP application is replaced with a constant bit rate (CBR) application to investigate the effect of handovers on *real-time* traffic. The CBR application is configured with a constant data rate of 300kb/s and a packet size of 210 bytes. A VoIP application typically requires a bandwidth of 64kb/s with the same packet size. We speculate that future real-time applications will require simultaneous voice, video and data, thus justifying the seemingly high data rate of 300kb/s for the simulations. The handover latency in this case is defined as the maximum period of disruption in the CBR packet stream during the handover. Thus the high data rate also provides a high resolution for measuring packet delays. PA-FMIP shows an improved performance over both proactive FMIPv6 and reactive FMIPv6 due to the reduced packet size. Again, PA-FMIP displays a 50% decrease in packet loss over reactive FMIPv6. The significantly longer handover latency of proactive FMIPv6 is again due to the timing ambiguity problem. PA-FMIP shows a handover latency close to that of the link-layer latency, with the least amount of packets lost. Bi-casting illustrates the performance that proactive FMIP should achieve, although it does increase the potential for duplicated packets.

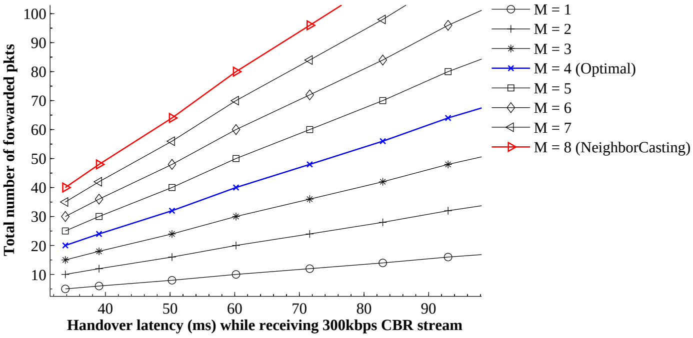

Fig. 6. Effect handover latency and number of predictions on total packets forwarded.

The simulation topology is modified to include a total of 8 neighbouring ARs, according to the grid topology in Fig. 2. The number of tunneled packets were recorded for ranging values of M. We see that for a PA-FMIP (UDP) handover of 33.8ms, M=4 (prediction hit probability 92.5%), a total of 20 packets would be forwarded with only 15 of these being redundant. With a similar latency, NeighborCasting would achieve the results equivalent to M=8 in Fig. 6: a total of 40 forwarded packets, 35 of these being redundant. Therefore PA-FMIP achieves a 50% improvement in forwarding overhead over NeighborCasting. The lowest cost of a 33.8ms PA-FMIP (UDP) handover (with M=4) can be approximated as $\frac{20}{f}$ times the cost of one proactive FMIPv6 handover, where f is the number of packets proactive FMIPv6 forwards during its handover. In this experiment f=7, therefore PA-FMIP costs a minimum of 2.86 times more than proactive FMIPv6. The simulation results in Fig. 6 are also confirmed analytically by the following equation:

$$P = \frac{\lambda}{pktsize \times 8} \times M \times t_h$$

Here the number of forwarded packets (P) is a function of the total streaming bit rate (λ) and packet size, the number of target nARs (M), and the duration of the handover t_h. This equation does not take the variable effects of congestion and link loss, nevertheless it closely models the dynamics involved in this experiment.

V. CONCLUSIONS

In this paper, a prediction assisted fast handover protocol (PA-FMIP) based on reactive FMIPv6 is proposed. Through extensive simulations and careful evaluation of the results, it is shown that PA-FMIP performs exceptionally well in a real-time application scenario. The results also indicate that PA-FMIP is comparable to proactive FMIPv6, yet it does not require a pre-trigger. This fact reduces the growing cross-layer dependencies evident in current fast handover proposals, simplifying transitions across heterogeneous access networks at a small cost.

Data mining has been shown to be a useful approach to mobility prediction, however the overall performance of PA-FMIP is restricted by the prediction accuracy. Perfectly accurate mobility prediction will remain a challenge for many years. Similar handover techniques that employ packet multicasting face the same overhead vs performance trade-offs so long as the cost of bandwidth is considered. Local mobility scenarios (e.g. in home, office or PCS networks) with finite movement boundaries may derive more benefit from this proposal, especially within low cost private or unlicensed access networks. Nevertheless, this work highlights the benefits of proactive packet forwarding in handovers, the poor handover performance of the standard MIPv6 and the important role that fast handover protocols play in providing seamless IP mobility.

ACKNOWLEDGEMENTS

The authors would like to thank Telkom SA, Siemens, the National Research Foundation (NRF) and the Department of Trade and Industry (DTI) for supporting this research project.

REFERENCES

[1] R. Agrawal and R. Srikant. Fast Algorithms for Mining Association Rules. *Proc. 20th. Conf. Very Large Data Bases (VLDB'94)*, pages 487–499, Sept 1994.

[2] R. Agrawal and R. Srikant. Mining Sequential Patterns. *Proc. IEEE Conf. Data Eng. (IEEE ICDE'95)*, pages 3–14, March 1995.

[3] F. Bodon. Trie-based Apriori Implementation for Mining Frequent Itemsequences. *OSDM'05*, Bart Goethals and Siegfried Nijssen and Mohammed J. Zaki editors, Chicago, IL, USA, 2005.

[4] U.S. Census Bureau. Tiger/line page. http://www.census.gov/geo/www/tiger/, April 2006.

[5] N. Van den Wijngaert and C. Blondia. A Predictive Low Latency Handover Scheme for Mobile IP. *ICMU'05, Osaka, Japan*, April 2005.

[6] G. Yavas et al. A data mining approach for location prediction in mobile environments. *Data and Knowledge Engineering*, 54:121–146, 2005.

[7] F. Feng and D. S. Reeves. Explicit Proactive Handoff with Motion Prediction for Mobile IP. *Proc. of IEEE WCNC'04*, 2:855–860, March 2004.

[8] B. Goethals. Survey on Frequent Pattern Mining. Technical report, HIIT Basic Research Unit, University of Helsinki, http://www.adrem.ua.ac.be/ goethals/software/, 2003.

[9] B. Goethals. Association Rule Mining Implementation. http://www.adrem.ua.ac.be/ goethals/software/, 2006.

[10] R. Hsieh. FHMIPv6 extension for ns2. http://mobqos.ee.unsw.edu.au/ robert/nsinstall.php, 2003.

[11] R. Hsieh and A. Seneviratne. A Comparison of Mechanisms for Improving Mobile IP Handoff Latency for End-to-End TCP. *MobiCom'03*, San Diego, USA:14–19, Sept 2003.

[12] E. Ivov and T. Noel. An Experimental Performance Evaluation of the IETF FMIPv6 Protocol over IEEE 802.11 WLANs. *in Proc. IEEE WCNC '06, Las Vegas, USA*, March 2006.

[13] R. Koodli. Fast Handovers for Mobile IPv6. RFC 4068, July 2005.

[14] G. Lui and G. Maguire Jr. A Class of Mobile Motion Prediction Algorithms for Wireless Mobile Computing and Communications. *ACM/Baltzer MONET*, pages 113–121, 1996.

[15] K. Malki and H. Soliman. Simultaneous Bindings for Mobile IPv6 Fast Handovers. *IETF Internet Draft*, May 2003.

[16] S. McCanne and S. Floyd. ns2 - Network Simulator. http://www.isi.edu/nsnam/ns/, April 2006.

[17] A. Nanopoulos, D. Katsaros, and Y. Manolopoulos. Effective Prediction of Web-user Access: A Data Mining Approach. *in Proc. of the WebKDD Workshop (WebKDD'01)*, 2001.

[18] A. Nanopoulos, D. Katsaros, and Y. Manolopoulos. A data mining algorithm for generalized web prefetching. *IEEE Trans. Knowl. Data Eng.*, 15:1155–1169, 2005.

[19] S. PalChaudhuri, J. Le Boudec, and M. Vojnovic. Perfect Simulations for Random Trip Mobility Models. *Published at 38th Annual Simulation Symposium, San Diego, California*, April 2005.

[20] A. K. Saha and D. B. Johnson. Modeling Mobility for Vehicular Ad Hoc Networks. *ACM VANET'04, Philadelphia, USA*, Oct 2004.

[21] E. Shim, H. Wei, and R. D. Gitlin. Low Latency for Wireless IP QoS with NeighborCasting. *in Proc. ICC 2002*, April 2002.

[22] Q. Zhao. *Mining Deltas of Web Structure: Issues, Challenges and Solutions*. PhD thesis, Nanyang Technological University, http://www.cais.ntu.edu.sg/ qkzhao/pdf/FYR.pdf, June 2003.

Andre E. Bergh is an MSc student in the Department of Electrical Engineering at the University of Cape Town. In 2004 he received his BSc degree in electronic engineering from the University of Natal. He is also the Secretary-General for the student branch of the IEEE at UCT.