Evaluation of Multiple Interface Functionality and Flow Mobility

Gareth Abrey {abreyg@crg.ee.uct.ac.za}, Neco Ventura {neco@crg.ee.uct.ac.za} Communications Research Group, Department of Electrical Engineering, University of Cape Town, South Africa.

***Abstract*— The Internet and mobile communications are converging towards a single, globally connected network. The latest mobile devices are increasingly equipped with multiple access technologies. These interfaces will prove most useful if used in collaboration. The functionality possible with multiple interfaces is evaluated on a test-bed in this paper to determine the success of a typical multiple interface scheme. Our results show that there are distinct advantages in using multiple interfaces for data flows, however some scenarios illustrate potential performance problems under certain conditions.**

***Index Terms*— multiple interfaces,flow mobility,mobile device**

I. Introduction

NEXT Generation networks will consist of a number of different access networks interconnected to provide ubiquitous access to the global resource of the Internet. The coverage of these networks will also overlap, allowing users the choice of access networks for their communication needs.

Increasingly, mobile devices have more than one type of radio access interface built-in. In current mobile devices, a single primary radio interface performs all communications with the service provider. The availability of multiple radio interfaces proves most beneficial if all these interfaces can connect with the service provider simultaneously to carry data in collaboration or individually.

It is expected that manufacturers will continue this trend and equip their portable devices with more wireless standards and capabilities such as Wireless LAN and WiMAX. But these devices tend to use their multiple radio access capabilities in isolation. The user would specifically activate the Bluetooth radio to exchange a photograph with a peer for example, then deactivate the link. The GSM or 3G radio would still be used for all primary communication with the cellular service provider. If the portable device had a Wireless LAN interface, this would be used to exchange data with an office wireless network, for example, but not for connecting to a service provider's network to make a voice call. The obvious question is how the user could benefit should the mobile device be capable of using its multiple access technologies in synergy and collaboration to complement the primary access technology. While the applications of such a capability would greatly expand as users' demands increased, some possible benefits include:

- Increased data transmission speeds by combining multiple interfaces.
- Overlapping radio coverage for wider access.
- Redundancy with multiple radio options in event of failures.
- Cost effective access network selection suiting each application.
- Geographically appropriate access network deployment .
- Radio spectrum diversification due to different frequencies and modulation techniques.

A mobile device with multiple access technologies needs mobility support. In order to support mobility, the device would need a protocol capable of handling the changing network addresses assigned to it as it roams through various access networks using its multiple interfaces. With the increasing adoption of the Mobile IPv6 standard in some isolated networks connected to the Internet, it is expected that Mobile IPv6 will become the dominant protocol for transporting data across the Internet, specifically to these mobile devices. It is therefore encouraging to see that a number of solutions proposed to implement multiple interface management in mobile networks have focused on Mobile IPv6 as the enabling protocol. But modifications to the protocol are necessary, since the basic MobileIPv6 protocol encounters a number of problems if the mobile device has more than one interface [6].

II. Previous Work

Dutta et al propose that a good mobility management scheme will support various features [3] including:

- Personal, service and terminal mobility
- Global roaming - the mobile device should enable communications to be independent of the access technology being used.
- Real-time and non-real-time communication should be possible over any access network with comparable pricing, Quality of Service and authentication methods.
- Support for TCP based applications.
- Multicast and anycast capabilities despite the mobile device's changing network attachment point.

Their SIP based testbed establishes a multimedia SIP session with real-time traffic between two 802.11 wireless LAN clients. It was found that 150ms in total was needed to complete the entire re-registration process necessary to inform the stationary node of the mobile node's new location. During this time, packet loss occurred, interrupting the multimedia session.

Although this Mobility Management scheme pertains to a device with a single interface, two observations can be made that are relevant to a multiple interfaced device. Firstly, the Mobility Management takes place at the Application layer. Our software router based solution is effectively an application, but operates at various OSI layers. Secondly, the Mobility Management scheme above experiences packet loss during the handover process. Our scheme aims to eliminate packet loss during handover.

NTT DoCoMo is developing IP-based mobility management technology for 4G system[7]. In their paper, the authors establish that there are three important requirements for mobility management.

- High packet transmission quality - essential for applications to operate efficiently. This refers to minimal packet transmission delays, low packet loss and low packet jitter.
- Cost control over wireless links - due to wireless links having less capacity than wired links, signalling overhead for mobility management needs to be minimal.
- Seamless mobility - hand-overs and access network roaming need to be seamless on all types of access networks.

In our evaluation platform two of these three conditions are met. High packet transmission quality is achieved and is one of the determining factors when conclusions are made about each test scenario performed on the evaluation platform. No signalling overhead is needed, since intelligence is pre-configured into the routing software but signalling would be used in a real world implementation, so this condition is not achieved. Hand-overs are seamless when the multiple interfaces are operating below their transfer limits and are seamless due to both interfaces being ready to transmit before the handover takes place.

Deciding at which layer to implement mobility management needs to take into account the above three requirements. If mobility management is implemented at the IP layer (network layer), then all three requirements are met, since the IP layer is common to the link layer, transport layer and application layer. In our evaluation platform, mobility management is effectively performed at the network layer, however, the routing software is a cross-layer solution.

A few multiple interface management schemes have been published, each with their own solutions to a large and complex problem. Mobile IP lacks any multi-access capabilities and so the authors' proposed protocol, Mobile Multi-Access-IP (MMA-IP), enables users to connect to multiple access networks and switch between different access domains[1]. MMA-IP does make use of Mobile IP and assumes that this protocol exists in the backbone networks connecting the different access domains to enable inter-domain mobility. Any mobility protocol is supported in each access domain, such as Cellular IP or Hierarchical Mobile IPv6 (HMIPv6), as long as the access network is connected to the Mobile IP backbone via an Access Domain Foreign Agent (ADFA). In order to facilitate multi-access capabilities, a new mobility agent is introduced into the access domain, called a Multi-Access Agent (MAA). This agent is the incoming point in the network topology for all packets destined for the mobile nodes. Packets are redirected according to preferences sent to the MAA by the mobile nodes in the form of preference messages.

In Figure 1 a single packet flow from a Corresponding Node (CN) is received by the MAA and split into two packet flows that are routed to the Mobile Node (MN) via two different access networks. This is an example of the functionality provided by having an MAA connected to all possible access networks that the mobile node can use.

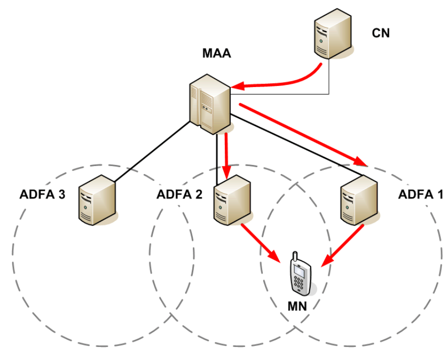

Fig. 1. Mobile Multi-Access IP (MMA-IP) Topology

Having a new agent in the network topology acting as a packet redirection point is useful in that it allows different access domains to connect at one point - this enables the MAA agent to redirect incoming packets to each access domain from a central point, while leaving each access domain to handle packet transport to the mobile node itself. A drawback of this scheme is the fact that a single point of failure is now introduced into the network. To overcome this drawback, the authors propose a hierarchical distribution of MAA agents, with multiple levels of agents in the network topology.

Fikouras et al show that using multiple wireless interfaces simultaneously provides the mobile user the most benefits over a single interfaced device if active packet flows can be distributed across the available wireless interfaces and be seamlessly transfered from one interface to another during mid-flow, without interruption[4]. Key benefits of such capabilities include the ability to aggregate networks for better performance, matching the multimedia flow requirements to the appropriate network and achieving hand-overs with no packet loss.

Mobile IP is criticised as it is designed for a single interfaced device. For this reason, agents in the Mobile IP architecture responsible for packet redirection are unable to distinguish between individual packet flows travelling to and from the mobile device. In order to equip the protocol with functionality needed for a multiple interfaced device, an extension called *Filters for Mobile IP* was developed.

This messaging protocol allows the mobile node to inform the packet redirection agent (either a Home Agent or Hierarchical Agent in the Mobile IP architecture) of its preferences for packet flow redirection. These preferences are referred to as filters. While *Filters for Mobile IP* is shown to perform flow mobility, the protocol extension needs to be deployed in a Mobile IP network and these networks are scarce on the Internet.

It is now appropriate to discuss the types of functionality possible when a mobile device has more than one interface. In order to achieve maximum benefit, two interfaces of different access technologies are used and could potentially provide the following functionality:

- Mid-flow Interface Transfer - During a packet stream transmission, this packet stream could be transferred from interface one to interface two. This might be due to a forced handover from the user or one of the interfaces moving out of radio coverage. In order to minimise handover latency, it would be best to have the second interface ready to transmit and receive packets before the first interface hands over the packet stream.
- Dual Interface Aggregation - A single incoming packet flow that is destined for the mobile node is split according to some efficient scheduling mechanism so that some packets travel over interface one while the rest travel over interface two. This would in theory allow for the single packet flow to be transmitted much faster to the mobile node than over one interface.
- Redundancy - When one interface is experiencing high packet loss, a weak signal or coverage unavailability, the second interface can provide an alternate connection if it has coverage available.
- Packet Duplication - When both interfaces are experiencing packet loss, duplicating packets and transmitting them on both interfaces should reduce the overall packet loss by a large factor.
- Bandwidth-on-Demand - When using an application, data transfer rate can vary as the user performs different tasks. While interface one may be sufficient for a certain amount of data, if its transfer limits are reached, interface two could be automatically activated to carry the surplus data.

III. Methodology

To properly evaluate the functionality of a multiple interface scheme, a set of test parameters and standards is used. The evaluation platform is intended to simulate the functional components of a multiple interface scheme and carry out test scenarios on this simulated scheme. The emulated mobile device is equipped with two different interfaces. This mobile node connects with the router which provides the two access networks for the two different interfaces. This router then connects directly to a corresponding node that generates traffic.

Voice over IP has strict requirements for packet latency and packet jitter [2]. Latency should remain below 150ms for an intelligible voice conversation (100ms is optimum). Beyond this, echoing becomes problematic and the two talking parties will start to talk over each other due to the delay. Packet jitter should remain below 75ms (40ms optimum). The third parameter under investigation is packet loss. Voice over IP can tolerate up to 3% packet loss, but 1% or less is optimum. When examining the evaluation platform's test data, the latency, jitter and packet loss maxima outlined here will be used to determine if the scenario would support quality Voice over IP transmission over multiple interfaces. Packet sequencing also plays a crucial role in enabling quality multimedia streaming. Packets need to be delivered in order to the mobile node and a protocol like RTP can perform this task. Any out-of-order packets are held in the buffer until the next packet in the sequence arrives. But these buffers add further latency and jitter to the packet stream. For this reason, the evaluation platform specifically records out-of-order packets.

A clear trend emerges in terms of the hardware topology that a multiple interface scheme would require. The typical components of a multiple interface scheme are shown in Figure 2

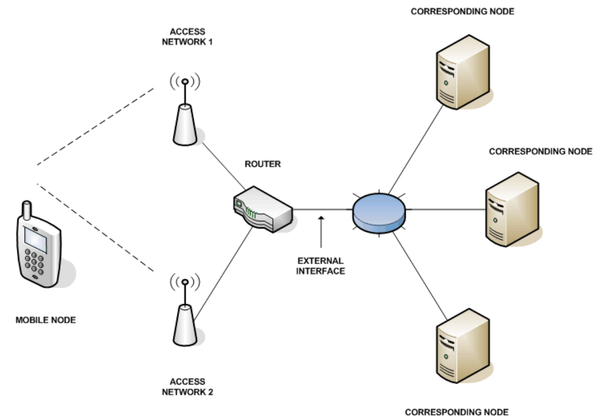

Fig. 2. Typical components of a multiple interface scheme

- The mobile node.
- Access network connections. These may take the form of a wired connection such as Ethernet or ADSL, or a wireless connection such as GSM, 3G or Wireless LAN.
- Router. The router acts as the redirection point for packet flows that are destined for the mobile node.
- External Interface on the Router. This is the interface which receives all incoming traffic from Corresponding Nodes that have packets destined for the Mobile Node.
- Corresponding Nodes. Any node sending traffic to the Mobile Node or receiving traffic from the Mobile Node.

These components are the minimum requirement for the Multiple Interface evaluation platform and are emulated with hardware.

In order to introduce both hardware and performance diversity, a 100BaseT Ethernet card and 802.11b Wireless card were chosen to represent the two different interfaces of a typical Mobile Node. It is expected that the wired Ethernet connection would have different performance characteristics to the Wireless connection.

The router needs three interfaces. A wired 100BaseT Ethernet connection directly to the Mobile Node (Interface One), an 802.11b Wireless LAN card connected to the wireless LAN card in the Mobile Node (Interface Two) and an External Interface (Incoming Interface) provided by a wired 100BaseT Ethernet connection directly to the Corresponding Node. Only a single interface is required for the Corresponding Node, so it is equipped with a wired 100BaseT Ethernet card, connected directly to the Router.

Each of the interfaces in the test bed was assigned a unique IP Address based on the subnet associated with each pair of communicating interfaces. Debian Linux proved an ideal platform on which to build the test bed. Each of the three PC's had Debian Linux installed, with routing tables setup according to the subnets of each pair of interfaces.

Suitable routing software was needed that would allow packets to be individually intercepted from each network card and then process and routed according to the test scenario underway. A suitable software router is the *Click Modular Router* developed by MIT and Eddie Kohler[5], [8]. This software allows the user to build complex router configurations and packet processing entities out of simpler modules called elements. These elements are connected to one another like a flow diagram.

Since packets will be routed over different interfaces during each test scenario, their headers need to be modified to suit the subnets over which they will be transmitted . This functionality is part of the router configuration scripts designed for each test scenario. This ensures that packets are topologically correct and will be accepted at the Mobile Node as if packets originated from the Router itself, rather than the Corresponding Node.

Packets are generated on the Corresponding Node as 172 byte UDP Packets, with parameters such as number of packets, inter-packet transmit delay and packet number being varied for each scenario.

A program running on the Mobile Node captures each packet received and produces a table of data showing important parameters of each received packet. These parameters include packet number, inter-packet receive time and interface identifier.

Four test scenarios were implemented.

1) A Mid-flow interface transfer simulates a handover of a single packet flow, from one interface of the MN to the other.
2) Dual Interface Aggregation investigates the splitting of a single flow into two streams, one sent over each of the MN's interfaces.
3) Lossy Link performance is investigated where both links are configured with a certain packet loss. A single packet stream is duplicated and each copy is sent to one of the links.
4) Bandwidth-on-Demand is implemented with a single link carrying all packets up to a certain data rate. Once this data rate is exceeded, the second interface is activated to carry the excess packets.

These four scenarios represent four major benefits possible in a multiple interface environment. The results of these four scenarios are discussed next.

IV. Results

The Mid-flow Interface transfer scenario was executed as follows: Two separate mid-flow hand-overs were investigated. The first handover is from the Ethernet interface to the Wireless LAN interface of the MN. Five separate tests were conducted, increasing packet send rate from the CN. This will test the handover point under increasing data rates for any anomalies that might degrade the performance of the real-time traffic flow. The second handover is from the Wireless LAN to the Ethernet interface. Again, five separate tests were conducted with the same increasing packet send rate from the CN.

Handover from Ethernet to Wireless LAN shows no obvious handover point until inter-packet transmit time was 500μs as shown in Figure 3. The Ethernet interface shows very low jitter, while after the handover point at packet 100, the higher jitter of the Wireless LAN interface is evident. Despite this increased jitter, all packets were received in order at the MN, but since jitter is well below 40ms, the Voice over IP stream would not be affected in terms of voice quality.

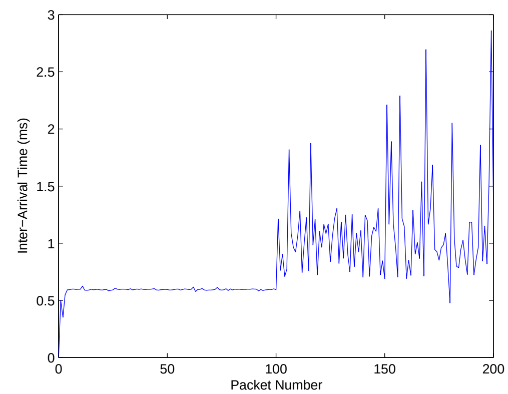

Fig. 3. Handover from Ethernet to Wireless LAN

When examining a handover in the reverse direction, from Wireless LAN to Ethernet, a different trend emerges. Out-of-order packets were observed in a number of test runs. For this reason, it is no longer appropriate to plot inter-arrival time on a graph. Instead, absolute arrival time is plotted. This will show out-of-order packets as a deviation from a straight line graph.

With an inter-packet transmit time of 500μs from the CN, out-of-order packets are clearly shown in Figure 4. Packets are received over the Ethernet interface (top line in the graph) before the packets transmitted on the Wireless LAN interface (bottom line). An overlap of 40.5ms is shown. During this time, packets will need to be buffered.

In both handover investigations, a seamless handover was achieved, with out-of-order packets only evident when handover from Wireless LAN to Ethernet occurs. Buffering will mitigate the problem, since future out-of-order packets will be stored until all Wireless LAN packets have been received. This scenario shows that Wireless LAN and Ethernet interfaces can work seamlessly together to eliminate packet loss and handover delays during a handover, however, intelligent buffering is needed to keep the Voice over IP stream constant during handover. Mid-flow Interface transfer is therefore made possible and practical with two interfaces and can perform well when interface limits are not exceeded.

The second scenario investigated was Dual Interface Aggregation. The CN generated a single packet stream and this single stream was split at the router to travel over both outgoing interfaces. Five test runs with different packet transmission rates were executed. An inter-packet transmit rate of 1ms resulted in 21 out-of-order packets out of a total of 500, but only a single packet-sized buffer would be needed to re-order the packet stream at the MN. Average jitter was 0.7ms with a maximum of 4.8ms, well below the 40ms limit for Voice over IP. When inter-packet transmit rate was decreased to

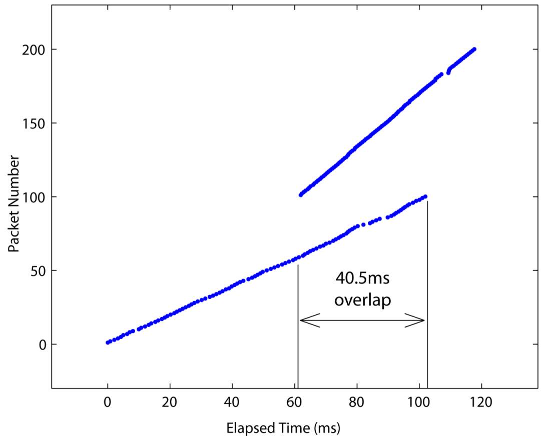

Fig. 4. Out-of-order packets during Wireless LAN to Ethernet handover

500μs, the Wireless LAN interface had significant difficulty in keeping synchronised with the Ethernet interface. 130 out of 500 packets were out of order and now a 17 packet-wide buffer would be needed. Maximum jitter was 8.1ms with an average of 1.4ms. However, these numbers are still within acceptable limits for quality Voice over IP communication.

A synchronisation problem is clearly shown in Figure 5. The two divergent lines indicate that the Wireless LAN interface is unable to maintain high enough transmission rate to synchronise with the Ethernet interface. Jitter is therefore increasing linearly and indefinitely. No buffering will help synchronise out-of-order packets, so under these test conditions (a 200μs inter-packet transmit time), Wireless LAN cannot be used for dual interface aggregation with an Ethernet interface.

These results show that Dual Interface Aggregation can work if both interfaces are closely matched in capability, but as soon as one interface cannot keep up with the other, synchronisation will no longer be possible. It is also shown that two interfaces working together will not necessarily result in increased transmission rate, since once synchronisation is lost, the quality of the stream decreases. However, when both interfaces are operating within their limits, Dual Interface Aggregation is shown to work.

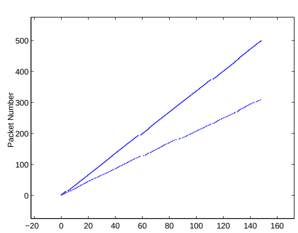

Fig. 5. Synchronisation problem evident during Dual Interface Aggregation

The third scenario investigated was packet duplication over two lossy links. In theory, transmitting duplicated packets over two identically lossy links will result in reduced combined packet loss at the receiving MN. For this scenario, both interfaces were scripted to drop 5%,10%,15% and 20% of all packets for each respective test run. The theoretical combined packet loss of the two duplicate packet streams is L x L, where L is the packet loss as a fraction of 1. In the case of 5% packet loss per link, this equates to 0.05 x .0.05 = 0.0025 ie.: 0.25% packet loss for the combined links. For each of the four test runs, 1000 packets were streamed to the MN over each link. A total of 2000 packets is thus transmitted. Two statistics are recorded. 'Packets dropped ' states the total number of packets dropped by both links. 'Packets missing' states the number of packets that were missing from the two recombined streams at the MN ie. both links dropped the same packet. This number is therefore the critical statistic for each test run, since missing packets will disrupt the Voice over IP stream in terms of quality. Figure 6 shows these two statistics for all four test runs.

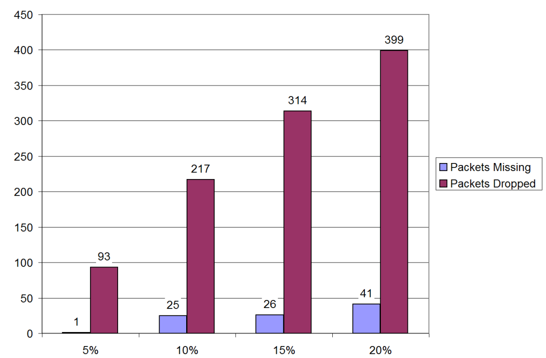

Fig. 6. Packets dropped and missing for the four test runs.

Voice over IP can tolerate a combined packet loss of 1% before voice quality deteriorates beyond acceptable limits. It is clear that in the test run with 5% packet loss per link, neither link is capable of carrying a quality Voice over IP session individually. However, when these two links are combined and stream the duplicate packets, the recombined stream was found to have 0.1% packet loss which meets the <1% packet loss requirement. Effectively, two unusable links have been combined into a single link that is significantly more reliable. For the remaining three test runs, the combined link packet losses were in excess of 1%, but as can be seen in Figure 7,the combined packet loss is many times less than the individual links' packet loss.

This scenario shows that lossy links can be significantly improved if they are combined into a single channel. For this reason, multiple interfaced devices would be extremely useful in situations where link quality is poor or coverage and network congestion are unsuitable for sensitive applications like Voice over IP.

The final scenario is a Bandwidth-on-Demand demonstration using multiple interfaces. The CN will stream packets to the router at 640 kbits/sec. Both interfaces are limited to 384 kbits/sec to simulate 3G links. The router script sends packets over only one interface if the incoming packet stream is below 384 kbits/sec. Once the incoming packet stream exceeds 384 kbits/sec, the second link will automatically route the excess

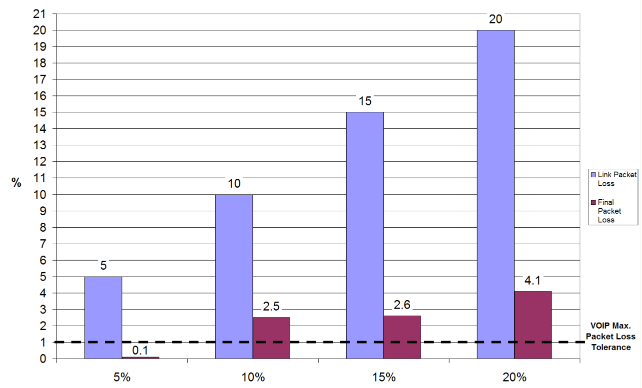

Fig. 7. Individual and Combined link packet loss

packets. Both links are in an active state before the test run begins. Two test runs were executed. In test one, the Wireless LAN link acts as the standby for the Ethernet link. In test two their roles are reversed. In both tests, only 3 packets out of 500 were out-of-order. Figure 8 shows the Wireless LAN and Ethernet links carrying 640 kbits/sec of incoming traffic. The Ethernet link is the on-demand link in this test run. It is clear that packets are being delivered consistently and both links are able to support the 640 kbits/sec stream. The triggering of the on-demand Ethernet link happens automatically when the router detects an incoming packet flow of greater than 384 kbits/sec.

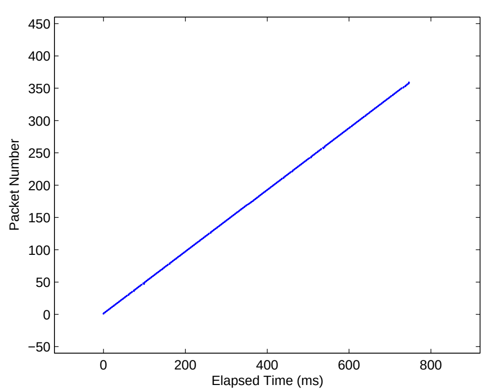

Fig. 8. Wireless LAN with Ethernet as on-demand link.

This scenario shows that a Bandwidth-on-Demand scheme is made possible by having a multiple interfaced device. Packets are routed reliably over both links with minimal out-of-order packets occurring. A single interfaced device would not be able to support any sort of Bandwidth-on-Demand scheme and so this application is shown to be both feasible and practical using two interfaces.

V. Conclusions

This paper has outlined investigations into the functionality possible with a multiple interface device in an IP network. Four possible scenarios show the type of performance possible and demonstrate problems that can occur under certain conditions. In the literature, it was found that multiple interface management schemes tend to have a key component - a central packet redirection point. This point was reproduced in a test bed, which enabled the following conclusions to be drawn:

1) Single interfaced devices are no longer sufficient to enable ubiquitous access to emerging next generation networks and mobile devices.
2) Mobile IPv6 can enable mobility with sufficient modification to support multiple interfaces.
3) An evaluation platform using the Click Modular Router software allows for packet modification with minimal delay and enables effective benchmarking of a multiple interface scheme.
4) UDP is a good choice of protocol for routing real-time traffic over multiple interfaces.
5) A Mid-Flow handover with no packet loss is possible, but both interfaces must be capable of supporting the full transfer rate of the incoming packet stream.
6) Dual Interface Aggregation is a possible method of increasing total data rate, however, out-of-order packets cause delays that affect buffer sizes and stream quality. It is therefore important that both interfaces are closely matched in ability.
7) Two lossy links with unacceptably high packet loss can be combined into a single link with significantly reduced packet loss.
8) A Bandwidth-on-Demand scheme will prove effective if both interfaces carry traffic within their transfer limits.

It has been shown that a multiple interfaced device provides significant benefits over a single interfaced device. This proved that multiple interface management schemes are needed in future mobile networks. The functionality possible with these devices will enable better communication, availability of resources and user satisfaction. For these reasons, manufacturers should continue to equip mobile devices with more interfaces and network operators should support the usage of these interfaces as outlined in this paper.

References

[1] S. Altuntas and B. Baykal, "Mobile Multi-Access IP: A Proposal for Mobile Multi-Access Management in Future Wireless IP Networks," *Computer Networks*, vol. 47, pp. 577–592, 2005.
[2] Cisco, "A Clear Understanding: What makes Voice over IP Work?" Internet - http://www.cisco.com/warp/public/779/smbiz/community/voip_basics.html. Last accessed March 2007.
[3] A. Dutta, F. Vakil, J. Chen, and M.Tauil, "Application Layer Mobility Management Scheme," *Computer Science Department, Columbia University, New York.*
[4] N. Fikouras and K. K. et al, "Multiple Access Interface Management and Flow Mobility," *University of Bremen, Germany*, 2004.
[5] E. Kohler, "The Click Modular Router," Ph.D. dissertation, MIT, February 2001.
[6] N. Montavont, R. Wakikawa, T. Ernst, C. Ng, and K. Kuladinithi, "Analysis of Multihoming in Mobile IPv6," Internet - http://www.nautilus6.org/doc/drafts/draft-montavont-mobileip-multihoming-pb-statement-05.html. Last accessed March 2007, April 2006.
[7] I. Okajima, M. Inoue, K. Omae, H. Takahashi., and N. Umeda, "IP-based Mobility Management Technology," *NTT Technical Review*.
[8] Various, "Click Modular Router Wiki," Internet - http://www.read.cs.ucla.edu/click/. Last Accessed March 2007.