

Multiple Interface Management and Flow Mobility in Next Generation Networks

Prepared by:
Gareth Roy Abrey

Supervised by:
Neco Ventura

Department of Electrical Engineering
University of Cape Town
2007

This dissertation is submitted to the University of Cape Town
in fulfilment of the academic requirements
for the Degree of Master of Science in Engineering

February 2007

The copyright of this thesis vests in the author. No quotation from it or information derived from it is to be published without full acknowledgement of the source. The thesis is to be used for private study or non-commercial research purposes only.

Published by the University of Cape Town (UCT) in terms of the non-exclusive license granted to UCT by the author.

Declaration

I declare that this thesis is my own work. Where collaboration with other people has taken place, or material generated by other researchers is included, the parties and/or material are indicated in the acknowledgements or references as appropriate.

This work is being submitted for the Master of Science Degree in Electrical Engineering at the University of Cape Town. It has not been submitted to any other university for any other degree or examination.

Gareth Roy Abrey

Date

Acknowledgements

I would like to express my sincere gratitude to the following individuals and organisations for their assistance during the course of this project.

Telkom and the Centre of Excellence program, for financing my Master's studies.

Mr. Neco Ventura, for his supervision and guidance throughout the project.

David Waiting and Andre Bergh, for your knowledge, encouragement and friendship.

My parents and my sister for their love, encouragement and constant support.

Synopsis

Next Generation networks will consist of a number of different access networks interconnected to provide ubiquitous access to the global resources available on the Internet. The coverage of these access networks will also overlap, allowing users a choice of access networks.

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 different radio interfaces proves most beneficial if all these interfaces can connect with the service provider and carry data in collaboration or individually. This means that a control system is needed to route the correct traffic over each different interface, depending on the requirements of that traffic.

Having multiple interfaces available provides the opportunity to aggregate two or more interfaces for faster transfer speeds and can provide redundancy. If one interface is experiencing high packet loss or no coverage an alternate interface will be available.

Multiple interface schemes aim to enable traditional networks to support devices with more than one interface. This is usually achieved by introducing a new agent into the network architecture that acts as the packet redirection point. Incoming packet flows are routed to the different interfaces of the mobile device by this agent according to the traffic types of each packet flow.

In this thesis an evaluation platform is developed to investigate whether the possible functionality of a multiple interfaced device provides useful traffic routing options. The evaluation platform consists of three key components evident in schemes from the literature, namely a Corresponding Node, Mobile Node and Router. The Router is emulated with a script-based routing software and configured as the packet redirection point in the evaluation platform.

Four test scenarios emulate traffic travelling over two interfaces of a practical mobile node. A mid-flow handover from one interface to the other is investigated to determine that this process can be seamless under certain conditions. Dual Interface Aggregation shows good performance when the limits of each interface are not exceeded. Distinct improvement in combined packet loss of two lossy links carrying duplicate packet streams shows that two interfaces can provide a reliable link in critical situations where both interfaces have poor performance when used separately. Finally, a Bandwidth-on-Demand scenario shows that

having two interfaces can allow automatic bandwidth allocation when data-rate is increased beyond the limits of one interface.

The results of these tests show significant benefits in transfer speed, link reliability and bandwidth-on-demand under most conditions. However, certain test parameters show that problems do occur when packet flows undergo a handover from one interface to another. The worst problem is out-of-order packet arrival and this was observed in certain scenarios. Despite this, it is shown that the distinct benefits of using a multiple interfaced device can outweigh the complexity of a multiple interface management scheme and its inherent problems.

Contents

Declaration	i
Acknowledgements	ii
Synopsis	iii
List of Figures	ix
List of Tables	xi
Glossary	xii
1 Introduction	1
1.1 Background Information	1
1.2 Problem Description	4
1.3 Thesis Objectives	5
1.4 Scope and Limitations	6
1.5 Thesis Outline	7
2 Literature Review	9
2.1 Introduction	9
2.2 Mobility Management	9
2.3 Multiple Interface Management Schemes	12

2.3.1	Mobile Multi-Access IP	12
2.3.2	Multiple Access Interface Management and Flow Mobility	14
	Filters for Mobile IP	15
2.4	Issues with Mobile IP	16
2.5	Discussion	17
3	Design Considerations	19
3.1	The Evolution from First to Third Generation Networks	19
3.2	Beyond the Third Generation	20
3.3	Challenges faced by 4G Networks	21
3.3.1	The Mobile Device	21
3.3.2	Access Network Discovery	22
3.3.3	Access Network Selection	22
3.3.4	Device Mobility	23
3.3.5	Network Infrastructure	24
3.3.6	Security and Privacy	24
3.3.7	Fault Tolerance and Reliability	25
3.3.8	Billing and Accounting	25
3.3.9	Profile System	26
3.4	Multiple Interface Functionality	28
3.5	Requirements of the Evaluation Platform	30
3.5.1	Real time Traffic Constraints	30
3.5.2	Packet Sequencing Requirements	31
3.6	Chapter Discussion	32

4	Design and Implementation of Evaluation Platform	33
4.1	Multiple Interface testbed Components	34
4.2	Testbed Hardware	36
4.2.1	Mobile Node	36
4.2.2	Router and Access Networks	36
4.2.3	Corresponding Node	37
4.2.4	Ethernet and IP Address Configuration	37
4.3	Operating System	38
4.4	Router Software	39
4.4.1	Packet Modification	39
4.5	Packet Generation	41
4.6	Data Collection	41
4.7	Test Scenarios and Configuration Scripts	42
4.7.1	Mid-Flow Interface Transfer	43
4.7.2	Dual Interface Aggregation	45
4.7.3	Lossy Link Performance	47
4.7.4	Bandwidth On Demand	49
4.8	Chapter Discussion	51
5	Performance Evaluation	53
5.1	Results	53
5.1.1	Mid-Flow Interface Transfer	53
	Results of Test A	55
	Results of Test B	57
	Discussion	60
5.1.2	Dual Interface Aggregation	61
5.1.3	Lossy Link	64

5.1.4	Bandwidth on Demand	68
	Test 1 - One interface active	68
	Test 2 - Both interfaces active	69
5.2	Discussion of Scenarios	71
6	Conclusions	75
6.1	Conclusions	75
6.2	Recommendations and Future Work	77
	Bibliography	79

List of Figures

2.1	Topology of Mobile Multi-Access IP showing Multi-Access Agent.	13
4.1	Critical Components of a Multiple Interface Scheme	34
4.2	testbed Hardware and Ethernet/IP Address Assignments	38
4.3	Packet modification diagram for testbed.	40
4.4	Click Element Flowchart for Mid-Flow Interface Transfer Scenario.	45
4.5	Click Element Flowchart for Dual Interface Aggregation Scenario.	47
4.6	Click Element Flowchart for Lossy Link Scenario	49
4.7	Click element Flowchart for Bandwidth-on-Demand scenario.	51
5.1	Packet flow - Mid-flow Interface Transfer Scenario Test A	54
5.2	Test 1 - 20ms (left) Test 2 - 2ms (right)	56
5.3	Test 3 - 1ms (left) Test 4 - 500 μ s (right)	56
5.4	Test 5 - 200 μ s	57
5.5	Packet flow - Mid-flow Interface Transfer Scenario Test B	57
5.6	Test 1 - 20ms (left) and Test 2 - 2ms (right)	58
5.7	Test 3 - 1ms (left) and close-up of handover point (right)	59
5.8	Test 4 - 500 μ s (left) and Test 5 - 200 μ s (right)	59
5.9	Packet flow diagram - Dual Interface Aggregation Scenario	61
5.10	Test 1 - 20ms (left) Test 2 - 2ms (right)	62
5.11	Test 3 - 1ms (left) Test 4 - 500 μ s (right)	62

5.12	Test 5 - 200 μs	63
5.13	Packet flow diagram - Lossy Link Scenario	65
5.14	Test 1 - 5% packet loss (left) Test 2 - 10% packet loss (right)	66
5.15	Test 3 - 15% packet loss (left) Test 4 - 20% packet loss (right)	66
5.16	Packet losses for the four test runs	67
5.17	Link and Combined Packet Loss	67
5.18	Packet flow diagram - Bandwidth on Demand Scenario Test 1	68
5.19	Ethernet link with 64 kilobit/sec stream	69
5.20	Wireless LAN link with 64 kilobit/sec stream	70
5.21	Packet flow Diagram - Bandwidth on Demand Test 2	70
5.22	Ethernet and Wireless LAN link with 640 kilobit/sec stream	71
5.23	Wireless LAN and Ethernet link with 640 kilobits/sec	71

List of Tables

4.1	Example output of dgram-rec.c	42
5.1	UDP Packet Parameters from Corresponding Node.	54
5.2	UDP Packet Parameters - Dual Interface Aggregation	61
5.3	UDP Packet parameters - Lossy Link Scenario	64
5.4	Theoretical Combined Packet Loss	65

Glossary

This section defines some of the commonly used terms and abbreviations that appear throughout this document.

1G The first generation of cellular communication standards (including AMPS).

2G The second generation of cellular communication standards (including GSM). Digital transmission is used to improve voice quality of the first generation analogue systems.

3G Third generation cellular communication standards covered by the ITU IMT-2000 family. Provides higher bandwidth than 2G systems.

ADSL Assymmetric Digital Subscriber Lines carry broadband over analogue copper telephone wires using higher frequency bands than voice.

AMPS Analogue Mobile Phone System using analogue radio technology to carry voice calls to mobile phones.

GPRS General Packet Radio Services is a packet-based data delivery system for GSM.

GSM Global System for Mobile Communications. Uses TDMA for up to eight calls per radio frequency.

IP Internet protocol transports data across networks using a packet header specifying the source and destination addresses.

ISDN Integrated Services Digital Network provides digital data and voice services over a standard copper pair.

LAN Local Area Network. A company or home network of computers and peripherals.

PDA Portable Digital Assistants are handheld computers with software aimed at business and organising.

PSTN Public Switched Telephone network refers to the telephone network carrying voice from customer to customer and provides copper line access to voice and data services.

TCP Transmission Control Protocol provides sequenced, reliable transmission of datagrams over IP networks.

VoIP Voice over IP is the two-way transmission of voice information over a packet-switched TCP/IP network.

Chapter 1

Introduction

1.1 Background Information

The Internet and mobile communications are converging towards a single, globally connected network. In the past, users connected to various networks via one specific type of access network, such as GSM/GPRS, 3G and IEEE 802.11 Wireless LAN or using wired access networks such as the PSTN, ISDN and ADSL. Each type of access network has its own advantages and disadvantages and they differ in terms of access speed, jitter, latency, link reliability, access charges and device capabilities. As access networks expand and increase in capability, users increase their expectations, demanding better quality access, larger coverage areas and lower cost. The latest mobile devices, mobile phones, PDA's and laptop computers are no longer based around a single access technology - a cellular phone would previously be equipped with one of the variants of the GSM mobile standard or the more modern 3G standard as its primary radio access technology. These mobile devices then evolved to include other radio access protocols. Bluetooth radios are prevalent in many handsets to allow interaction and exchange of data with other handsets and peripherals in the vicinity, however Bluetooth is not used as a primary radio access method for the handset's voice and data transmission requirements in terms of accessing the cellular service provider. More recently, some devices such as laptops and PDA's are equipped with 802.11 Wireless LAN, but this technology is still rarely found in the traditional cellular phone.

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 what performance benefits are possible, should the mobile device be capable of using all its multiple access technologies in synergy and collaboration. While the applications of such a capability would greatly expand as users' demands increased, some possible benefits include [1]:

- Data transmission speed could be increased should more than one access network be used simultaneously.
- Effective radio network coverage would include the radio footprints of all the access networks.
- Redundancy is provided should an access network be unavailable, unreliable or congested.
- The most cost effective network could be utilised for the required data session.
- The access network could be chosen based on the type of application required by the user.
- Service providers could deploy geographically optimised and cost effective additions to their networks choosing the best access technology for each area.
- More scarce radio spectrum would be available due to the different frequency bands and modulation techniques of each access technology.

This list is expandable and will be discussed in later chapters.

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 network using its multiple interfaces. With the increasing adoption of the Mobile IPv6 [2] 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 [3]. Mobile IPv6 was designed with a single interfaced device in mind, with a unique address assigned to each mobile device. This unique address is called the Home Address and is globally reachable by any other node on the Internet. If the mobile node is connected to its Home Network, then packets destined for the mobile node are routed to this Home Address. When the Mobile IPv6 mobile node is away from its Home network and connected to a Foreign Network, it is assigned a Care-of Address which is routable on that Foreign Network only. Packets destined for the mobile node are still sent to its Home Address, but are routed from the Home Network to the new Care-of Address by the Home Agent. In this way, a tunnel is set up between the Home Agent and the mobile node on the Foreign Network. But this situation only applies to a mobile node with one interface. In order to support a device with multiple interfaces, Multi-Homing capability is needed within the Mobile IPv6 protocol. This allows each interface on the mobile node to have its own unique Home Address and Care-of Address. While this may seem to solve the problem of a multiple interface device, issues remain such as the inability to separate individual packet flows transmitted on each interface. Mobile IPv6 is also not widely deployed on the Internet, so as an enabling protocol for multiple interfaced devices, it is still in its infancy. Future networks are expected to support IPv6 and IPv4 simultaneously and later migrate to only IPv6.

Despite this increasing adoption of MobileIPv6, the Internet and service provider networks currently still make use of the IPv4 protocol [4] and therefore it is important to attempt to enable these networks with the capability of managing devices with multiple interfaces. The IPv4 protocol was not designed to natively support any form of mobility, so upper layers in the OSI Model [5] need to be used in order to manage the IP flows travelling in a multiple interface enabled network. This means that control of IP flows will be handled by an intelligent agent at a packet level rather than by routing according to IPv4 packet headers only, since these headers were designed to support point-to-point transmission in a fixed network rather than a network of mobile devices with multiple interfaces.

The evaluation part of this research focuses on proving that the features and benefits made possible by having multiple interfaces are actually useful and do provide the user

with added functionality that warrants the extra complexity required in the network. For example, streaming a single packet flow over two interfaces is investigated, allowing for faster downloads if both interfaces are used simultaneously. Since the literature focuses on developing network architectures to support multiple interfaces and solving the problems of managing multiple interfaces rather than proving the advantages of such a scheme, this will be the focus of this thesis. It is therefore important to explore this question by developing a generic multiple interface evaluation scheme and using this scheme to determine whether it is in fact beneficial to have multiple interface capability in a network. These questions will be answered in the following chapters.

1.2 Problem Description

Each type of radio access protocol has been designed into mobile devices to act in isolation based on the application the user chooses. For example, to make a voice call, a cellular phone uses the GSM or 3G radio. If this cellular phone were also equipped with a Wireless LAN radio, no provision is made to carry the voice call over this connection, unless a separate Voice over IP application is used. Ideally, the user should have a choice as to which interface carries each type of traffic, be it voice, video, data or other form of multimedia. A single point of entry into a service provider's network through a single radio access interface is prone to certain problems, namely geographical limitations, network congestion, bandwidth restrictions and lack of redundancy and this list is extendable. The user of a single interfaced device is forced to use one type of radio access network for all their communication. Having a multiple interfaced device with a choice of wireless access networks would allow diverse access not currently possible with single interfaced devices. Users could even supplement network operators' networks with their own home or office networks by installing their own infrastructure, such as Wireless LAN access points in certain practical cases. All these benefits are however not possible unless the portable devices and networks are designed to support and manage multiple interfaces. Current networks are designed by operators with a single type of radio access solution in mind, for example, cellular network providers do not allow their users to connect to them via a Wireless LAN link should the user's device be equipped with this interface.

As network operators aim to provide better services, higher bandwidth, wider coverage and service reliability to customers and end-users, networks are becoming larger and more costly to build and maintain. A scheme that would allow network operators to develop

overlapping, redundant and scalable access networks would greatly enhance their single access-type networks through diversification of hardware vendors and access technologies, allow them to overcome certain practical limitations of each access technology and allow geographical-area related design flexibility. For example, in a densely populated airport or stadium, a cellular service provider could deploy Wireless LAN access points in close proximity to supplement the neighbourhood cellular base station if user's owned cellular phones equipped with GSM/3G and Wireless LAN interfaces that could both be used to carry voice calls. Customers are constantly demanding more choices, better functionality and a better network experience and service providers need to enhance their current networks to meet this demand while keeping costs low and using the most cost-effective technologies available.

Portable devices are already being produced with more than one type of network access technology and this trend is expected to continue as devices become more capable. Yet taking full advantage of these multiple interfaces is not currently possible without changes to network architectures and routing protocols. It is therefore important to add functionality to these portable devices in a simple manner without complicating the communication experience for novice users, nor hiding the complexities from the expert user or network operators. Networks thus need to cater for all types of users to meet the increasing expectations of all levels of users, while allowing cost effective implementation for network providers and operators.

1.3 Thesis Objectives

This thesis aims to show the need for mobile devices to have multiple interfaces and a suitable packet routing scheme to enable useful, efficient and cost-effective use of those interfaces. The possible advantages and disadvantages of a multiple-interfaced device will be investigated to evaluate whether having more than one interface in a mobile device will in fact provide added functionality that benefits the user, such as increased bandwidth, redundancy, link reliability and reduced packet loss. This study aims to investigate multiple access schemes proposed in the literature and from these schemes a simple evaluation platform was developed to implement multiple interfaces on a testbed. The aim is not to develop a new management scheme, but rather to take elements from existing schemes and use these elements in the design of the evaluation platform.

Through studying the available proposals for multiple interface architectures, this study will show how feasible such a management scheme will be and the practical hurdles of implementing such schemes. A software-router based scheme will be implemented on a testbed to obtain performance metrics using a simple software-based per-packet approach to routing packets over multiple interfaces. Important performance metrics are packet jitter, transfer speed, packet loss and handover delays. These measurements are important in determining whether a quality multimedia experience is possible in a multiple interface environment. Four test scenarios will be implemented, each addressing a specific proposed advantage of having multiple interfaces available. The strict requirement of Voice over IP traffic are used to benchmark the performance of each of the four test scenarios studied on the evaluation platform. The results of each scenario will indicate whether the proposed advantage under investigation is in fact useful and practical and show how it would prove useful to the user of a multiple-interfaced device.

1.4 Scope and Limitations

This study briefly discusses some implementation issues arising in 4G network when supporting a multiple-interfaced device. The intricacies of each level will not be addressed in detail as this is beyond the scope of this thesis. Research suggests important elements required in designing a multiple interface architecture. This thesis uses these elements to design an evaluation scheme based on components evident in schemes from the literature. The three important network components, Corresponding Node, Router and Mobile Node, are discussed and implemented as an indication of the functionality possible with simple yet effective multiple interface scenarios. Four scenarios were chosen to demonstrate the type of functionality possible with a mobile device equipped with two interfaces. Mobility of the Mobile Node was not studied as this is not part of the evaluation process.

The proposed evaluation scheme is intended to be a hardware-based simulation of typical components of such a multiple interface scheme. In the evaluation testbed, a number of possible benefits of multiple interface management are investigated using a script-based software router. This software router is designed as a simple IP network routing entity and is not intended to represent the real functionality of a enterprise sized network management system. Despite this, important observations were made in each of the four test scenarios, indicative of the issues that would arise in a real-world scheme.

This thesis does not deal with the practical, geographical or hardware challenges of implementing a multiple interface management scheme on a large scale. In addition, performance metrics obtained from the testbed are symbolic of the type of performance and benefits to be expected from a real-world scheme and do not indicate absolute figures obtainable in large-scale implementation. The details of radio networks are not considered, as this is outside the scope of a network architecture study at this level.

1.5 Thesis Outline

The remaining sections of this thesis are structured as follows:

- Chapter 2 reviews and investigates the multiple interface management schemes and proposals available in the literature to determine what key components are necessary to design an evaluation scheme. The evaluation scheme will show how multiple interfaces will benefit the user. The problems experienced by Mobile IPv6 when encountering a multiple interfaced node will be discussed briefly and lead to the conclusion that a Mobile IPv6 node is not necessarily the best way to support multiple interfaces. From these observations, it is clear that a multiple interface scheme is a complex problem to solve, so proving that multiple interface functionality is beneficial using the evaluation platform will justify the complexity of such a scheme.
- Chapter 3 focuses on briefly discussing the evolution of networks from the First Generation (1G) to the Fourth generation (4G). Issues arising from each part of the 4G network will then be examined and related to how a multiple interfaced device will complicate the implementation of next generation networks. Attention will also be given to the mobile device itself and some basic proposals of functionality will show how the user could manage the use of the multiple interfaces of the device through a simple profile system. This profile system is what ultimately controls the use of each interface and what type of traffic will be carried over each interface. The parameters used to evaluate the testbed data will then be discussed, as these parameters will determine the success of each scenario investigated on the testbed.
- Chapter 4 presents the configuration and implementation of the multiple interface testbed. Hardware and software resources are discussed. The setup of each test scenario is described in detail. Each test scenario focuses on one type of functionality

that a multiple interfaced node could possibly provide. The parameters of each test scenario and data recorded will be mentioned and result expectations are presented.

- Chapter 5 presents the results of each of the four test scenarios carried out on the multiple interface testbed. The results are discussed and analysis and observations are presented. This study uses these results and observations to prove the feasibility as well as the advantages and disadvantages of a typical multiple interface scheme. It will be shown that while there are obvious benefits to using the four types of functionality chosen for analysis, problems with each scenario are evident.
- Chapter 6 presents the conclusions obtained from this study and gives recommendations for further research to be done around multiple interface management architectures and scenarios.

Chapter 2

Literature Review

2.1 Introduction

In this chapter, literature related to the study of mobility management and multiple interface schemes is investigated. It is intended that from this research, a framework will be developed that will form the basis of the evaluation platform. Certain multiple interface management schemes make use of the Mobile IPv6 protocol. Mobile IPv6 has not been implemented on the Internet on a wide scale yet, so other mechanisms for mobility support are also studied. This chapter will therefore show the advantages and disadvantages of this protocol outlined in the literature, particularly in terms of supporting a mobile device equipped with multiple interfaces.

2.2 Mobility Management

Mobility management is a term used to describe the process of enabling mobile devices to roam seamlessly through different access networks while maintaining the devices' ongoing communications and ensuring that the device is reachable by other nodes on the Internet at all times, despite not being connected to a permanent network point. Dutta et al propose that a good mobility management scheme will support various features [6] including:

- Personal, service and terminal mobility, allowing the user to access services wherever they are, using a mobile terminal that can connect to a number of access networks

and allow multimedia sessions to continue uninterrupted while roaming from one network to the next.

- 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, to simplify the user's communication experience.
- Support for TCP based applications.
- Multicast and anycast capabilities despite the mobile device's changing network attachment point.

An Application layer Mobility Management scheme is proposed using Session Initiation Protocol (SIP) signalling to perform mobility operations such as enabling real-time and non-real-time connection management that keeps track of TCP connections and packet flows. SIP signalling is also used to perform resource allocation and authentication functions to ensure that the user is able to access any services they have subscribed to, regardless of what network they may be attached to.

Their SIP based testbed was created to investigate terminal mobility for real-time traffic. A multimedia SIP session with real-time traffic was established between two 802.11 wireless LAN clients. One client was mobile and was assigned a new IP address as soon as it discovered it was connecting to a different subnet. The testbed was used to record the size of SIP messages and signalling necessary to continue the multimedia session despite the one client changing subnets. 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, certain observations can be made that are relevant to a multiple interfaced device. Firstly, the Mobility Management takes place at the Application layer. In this research, the evaluation platform uses intelligence built into a software based router, not an Application layer based control system. Since the software router runs at the kernel level in Linux, it is a cross-layer management system as it can process link-layer, network-layer and transport-layer packet headers as well as the packet data itself. Using a cross-layer software router

allows changes to be made to the layers in the OSI Model and allows the evaluation platform to modify IP and UDP headers to perform multiple interface functionality. Secondly, the Mobility Management scheme above experiences packet loss during the handover process, since the device is equipped with a single wireless interface only. In our evaluation platform, a multiple interfaced device experiences no packet loss during the handover process using the software-based router. Their Mobility Management scheme makes use of the SIP protocol for signalling, whereas the evaluation platform in this thesis does not implement a signalling system since intelligence is pre-built into the routing software. However, SIP-based signalling would be a possible way to exchange signalling between our mobile node and the router.

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. 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 - handovers and access network roaming needs 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, however, in a real-world scheme, signalling will be needed to exchange routing configurations between the mobile node and the router. This signalling has not been explicitly implemented and so this condition for an effective mobility management system is not achieved. Handovers 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 link layer, then each link layer for each access network would have a different method of handling

mobility. This makes seamless mobility problematic, since interworking these different link layer mobility methods is a challenge. If mobility management is implemented at the transport layer or application layer, signalling traffic would be needed for each different protocol and application. This doesn't meet the requirement of reducing signalling traffic. 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. Since Mobile IP is a network layer mobility protocol, a network layer mobility scheme has already been developed for the network layer, confirming the authors' recommendation of network layer implementation. In our evaluation platform, mobility management is effectively performed at the network layer, however, it has already been shown that the routing software is a cross-layer solution. The network layer is suitable, as it is the convergence point between the different interfaces of the mobile device (link layers) and the applications running on the mobile node (transport and application layers).

2.3 Multiple Interface Management Schemes

A few multiple interface management schemes have been published, each with their own solutions to a large and complex problem. It is important to evaluate these schemes in terms of their practicality and determine if any of the components used in the scheme will in fact be useful in the process of designing an evaluation platform for this thesis. A brief overview of the functionality of the scheme will be given, followed by discussion on the relevance of the scheme to the design of our evaluation platform.

2.3.1 Mobile Multi-Access IP

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 [8]. 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. The MAA sits in the network at a level above all the access domains that connect to it, allowing packets to be sent to each access domain for each interface of the mobile node. Packets are redirected according to preferences sent to the MAA by the mobile nodes in the form of preference messages. These preference messages are stored on the MAA for each mobile node connected to it. In Figure 2.1, 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.

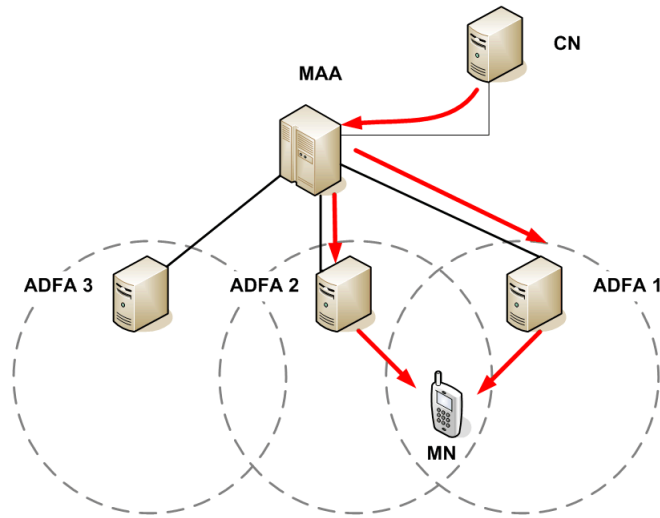


Figure 2.1: Topology of Mobile Multi-Access IP showing Multi-Access Agent.

Two interesting observations from this scheme need to be made. Firstly, 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. This allows different access domains with various access domain transport protocols to deliver packets from the central MAA agent. A drawback of this scheme is the fact that a single point of failure is now introduced into the network. Should the MAA fail, be congested or over-subscribed, the entire multi-access scheme is affected, since this central point is the controller of all multi-access capabilities. To overcome this drawback, the authors propose a hierarchical distribution of MAA agents, with multiple levels of agents in the network

topology. Together with multiple levels of agents, a backup scheme is also discussed, allowing agents on the same hierarchical level to act as backups for one another. Should an MAA agent fail, a nearby agent will assume its duties.

The second observation is the need for preference settings in a multi-access scheme. Preferences are the means by which a user can tailor the packet redirection decisions made by the multi-access agent. Without preferences, control is taken away from the user and the scheme becomes autonomous and less suited to the diverse needs of each individual user. Storing preferences on the MAA agents and allowing these preferences to be exchanged with neighbouring MAA agents in the case of failure is important to provisioning a robust scheme.

2.3.2 Multiple Access Interface Management and Flow Mobility

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 transferred from one interface to another during mid-flow, without interruption [1]. The increasing deployment of wireless technologies with packet services points towards the development of wireless overlay networks. A wireless overlay network is one where in any specific geographical location, a number of wireless technologies have coverage available and overlap one another. For example, an office may be covered by the company Wireless LAN, a GSM and 3G base station as well as a Satellite. In order to best take advantage of these overlapping networks, the user should be able to selectively and intelligently distribute their active multimedia packet flows over the available networks according to which is best suited to the flow. 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 handovers with no packet loss. In order to achieve these objectives, portable devices need to be manufactured with multiple interfaces and access networks need to be interconnected to allow vertical mobility (allowing the mobile device to handover flows between different access network types). This means that resources on the Internet need to be globally available to all access networks and associated with a unique identity for each mobile user. A control system is necessary to manage the multiple interfaces of the device, so a policy based database system will keep track of interface selection according to the type of multimedia flows being transported.

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. During a hand-off, all active IP flows will be redirected to the newly acquired access point. This does not meet the requirements of associating individual packet flows with different access networks according to which best suits the packet flow's requirements. It is possible to acquire multiple IP addresses for the mobile node's interfaces and use route optimisation to distribute active IP flows among the available interfaces by sending updates to corresponding nodes and informing them of the new interface's IP address. However this still does not allow individual IP flows from one corresponding node to be separated for redistribution to another interface. Since Mobile IP is not yet deployed on a wide scale over the Internet, the abundance of IPv6 addresses does not help in the present IPv4 address space, where IP addresses are scarce. So Mobile IP in its raw form is not suitable for a multiple interfaced device. In order to equip the protocol with functionality needed for a multiple interfaced device, an extension called *Filters for Mobile IP* was developed.

Filters for Mobile IP

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. Filters specify packet properties like source and destination addresses, port numbers and packet protocol. These properties are used by the packet redirection agent to make routing decisions for each packet flow destined for the mobile node. In order to reduce signalling overhead that would be necessary to transmit these filter properties to the packet redirection agent, filter information is attached to existing registration signalling. Filtering is transparent to transport layer protocols like TCP and no changes or updates are required for corresponding nodes. Only the mobile node and Mobile IP components (Home Agents and Foreign Agents) need to be aware of the filters. *Filters for Mobile IP* is also compatible with multi-homing for every separate home address associated with the mobile node's interfaces.

To test the performance of Filters for Mobile IP, Fikouras et al constructed a testbed. The testbed consists of mobile node with two Wireless LAN interfaces, a Home Agent, two corresponding nodes and two Foreign Agents that provide the wireless LAN access points. The two corresponding nodes initiate one TCP flow each to the mobile node.

Initially, both TCP flows are routed over only one of the wireless interfaces. After a 15 second timeout, the second wireless interface is activated. Once the required registration signalling is sent to the packet redirection agent (Home Agent in this testbed), the second TCP flow is tunneled over the newly activated wireless interface using filtering criteria sent to the Home Agent. The Home Agent is running an implementation of *Filters for Mobile IP*. In their results, the authors show that the transmission speed increased by a significant amount when the second interface was activated. This proves that two interfaces carrying one TCP stream each are capable of higher throughput than a single interface carrying both TCP streams. These results show how two interfaces can carry one stream each. In this thesis a different scenario is investigated on our evaluation platform. Two interfaces streaming a single split UDP packet flow are measured and it is clear that the packet flow is transmitted successfully to the mobile node. This process achieves higher throughput as the single stream is afforded two separate paths to the mobile node and is referred to as Dual Interface Aggregation.

While *Filters for Mobile IP* is shown to perform flow mobility, the protocol extension needs to be deployed in a Mobile IP network. Since these networks only exist in small isolated areas on the Internet, this solution to flow mobility is limited in use for today's mobile users. Our evaluation scheme uses an IPv4 network and software router without the need for Mobile IP in the network. It is expected the Mobile IP will become prolific on the Internet, so *Filters for Mobile IP* has potential in the near future.

2.4 Issues with Mobile IP

Mobile IP experiences a number of challenges when it is faced with a multiple-interfaced device. Multiple interface management schemes tend to use Mobile IP at some point in their architectures, but basic Mobile IP needs to be modified to accommodate multiple interfaced devices. The question remains as to whether it is in fact necessary to use Mobile IP for multiple interface architectures. Sun et al discuss the applicability of Mobile IP and whether it is in fact necessary, making two key observations [9]. The authors propose that Mobile IP lacks flexibility in providing mobility support to various applications and in the case of a multiple interfaced device, flow level traffic control is not provided by Mobile IP - an additional management system is needed. It is also shown that mobility management using Mobile IP is not needed in all types of access networks - IEEE 802.11, GPRS and UMTS all implement their own mobility management functionality with better

performance in terms of handoff delays, so there is no need to implement Mobile IP in these access networks. Also, application-layer protocols such as SIP can perform mobility management in normal IP networks. Even applications support device mobility - POP3 based email can be accessed from anywhere, meaning Mobile IP functionality is unnecessary in this case. So in all these cases, Mobile IP is overlapping the functionality already provided in certain access networks, applications and protocols.

Mobile IP also has issues in a multi-homing environment. In a single interfaced device, a single Home Address and Home Agent are assigned. The Home Address is the unique identifier for the node on the Internet. In the case of a multiple interfaced device, each interface will need its own Home Address. A problem arises - the mobile device now has multiple identities and it is not clear which Home Address refers to the mobile device itself, rather than an interface of the mobile device. Mobile IP is in effect a device level mobility control agent. Higher level protocols are more suitable for packet level traffic control.

Extra network entities like Home Agents and Foreign Agents need to be deployed. This increases costs and introduces points of failure in the network. While other literature has already shown the need for centralised packet redirection points in the network to facilitate multiple interfaces, in many cases, there is no obvious Home Agent for a device. The user may not have a permanent location for a Home Agent (such as a home computer) and so this entity would need to be provided by a service provider. Operating systems will also need protocol stack updates and Mobile IP will need to be added to devices communicating with the Mobile IP protocol. These compatibility issues will require wide-scale changes on the Internet and so it is impossible for these changes to occur in the near future.

In summary, Mobile IP clearly needs a substantial amount of adaptation to suit multiple interfaced devices, while practical limitations mean that implementation on a wide scale remains implausible.

2.5 Discussion

When considering mobility of a device and any scheme that attempts to manage this, it is clear that there are certain requirements for a successful scheme. High packet transmission quality, low signalling overhead and handover with minimum packet loss are necessary for quality communication. Mobility management needs to take place at specific layers within

the OSI Model. It is shown that the network layer provides ideal functionality in this regard, since in an all-IP scheme, the packet flows are routed at the network level. But a cross-layer mobility management scheme would be better suited to the overall control of individual packet flows. That is motivation for the cross-layer software router has been used for the evaluation platform in Chapter 4 and 5. This cross-layer ability means that the packet headers and contents can be modified to suit the redirection of that packet over multiple interfaces, while providing inter-operability between different access network types.

When examining some multiple interface schemes in the literature, some key components are observed. A centralised packet redirection point performs the routing of packet flows to and from the mobile node and allows different access networks to be interconnected. Our evaluation platform provides this central packet redirection point and interconnects two different access networks. The need for some sort of preferences is obvious, so that the user can control how packet flows are routed over each interface. These preferences were not implemented explicitly but the router was hard-coded to perform each test scenario.

The problems with Mobile IP as a mobility agent for multiple interfaced devices have also been outlined. While this protocol is suitable for single interfaced devices roaming through different access networks, it is not yet ready for use with multiple interfaced devices. Since certain access networks have their own mobility functionality already built in, Mobile IP is not necessary within these access networks.

The multiple interface schemes discussed here do not however show that the functionality they intend to provide is actually useful. This is the motivation behind the rest of this thesis. In order to justify the complexity of designing and implementing a multiple interface scheme, having a mobile device with two or more interfaces should show benefits that a single interfaced device cannot. The evaluation platform modelled the proposed schemes in the literature and the test scenarios were developed from functionality evident in these papers. Before the evaluation platform is described, it is important to investigate what challenges the Fourth Generation (4G) network will face in supporting multiple-interfaced devices. The following chapter outlines these challenges and attempts solutions to this ever expanding problem.

Chapter 3

Design Considerations

This chapter will provide a brief overview of the future of wireless communication networks beyond 3G and evidence that the user will no longer be restricted to using a single access network for all their communications. In conjunction with this, a brief look at the evolution of mobile terminals will show that these portable devices will also not restrict users to a single access network. The problems and challenges that multiple interfaced devices will bring to the 4G network will be mentioned as well as any possible solutions. Finally, some design considerations will be presented that will enable the evaluation platform to perform the necessary testing scenarios.

3.1 The Evolution from First to Third Generation Networks

Mobile communication networks had their origins in the early 1980's when these networks provided mobile telephony services to a select few who could afford and justify owning the technology. The First Generation (1G) mobile phone service was based on analog technology such as the AMPS system [10] used in North America. Mobile phones were bulky, expensive devices and provided only fair communication quality in good environmental conditions. The real revolution of mobile phone service came about when Second Generation (2G) devices were released onto the market - an example is the GSM standard currently used in many parts of the world. These mobile phones use digital radio technology to provide better voice quality and higher data rates with extensions like High

Speed Circuit Switched Data (HSCSD), General Packet Radio Service (GPRS) and Enhanced Data rates for GSM Evolution (EDGE). It is now practical for users to transmit and receive multimedia content such as pictures, audio and video at data rates comparable to analog computer modems and basic rate ISDN. This multimedia mobile experience has caused an increasing demand for faster mobile data rates and so Third Generation (3G) mobile phones and networks were developed. While 2G mobile phones are designed for voice and basic low-bit-rate data services, 3G devices allow true mobile broadband to reach users. With a basic data rate of 384kbps and enhancements such as HSDPA providing over 1Mbps, users can finally experience fixed line type speeds on their portable devices. Real time audio, video and multimedia are now possible while on the move and users can access the Internet and their home or office networks from anywhere with 3G coverage.

Current 3G smart phones are equipped with multiple interfaces, such as Wireless LAN and Bluetooth, but these additional interfaces cannot be used in conjunction with the 3G radio. For this reason, 4G networks aims to allow collaboration of all the interfaces on a mobile device so that they are no longer used in isolation.

3.2 Beyond the Third Generation

The evolution of 3G towards 4G is expected to produce an all-IP based heterogeneous network that will encompass many different access technologies into a unified system that users can access from anywhere at anytime [11]. Different research groups have different visions of the future 4G networks but some common components are evident. Integrated services will be key, allowing users to access anything from email to television channels through a single device. A great emphasis is placed on personalisation, allowing users to tailor their multimedia services to their individual needs and choose access networks based on the application or task running on the mobile device. The cost of data is also expected to drop drastically as users start paying for content rather than the amount of data exchanged. These three key components are of course only useful if the 4G network provides high availability and quality. In order to do this, the 4G network will need to support high bandwidth, seamless interconnection to the Internet and content provider networks, wide geographic access network coverage as well as choice of access network in overlapping areas. This means that 4G networks will be heterogeneous and this ubiquity poses some important challenges to network designers. For example, using several interfaces

simultaneously and achieving seamless handovers requires co-operation between different access networks.

3.3 Challenges faced by 4G Networks

There are three key areas in 4G Networks that require attention in order to solve the problems each faces in providing a heterogeneous network - the mobile device, the network and the services provided on the network [11]. Each of these three entities will face a number of challenges when dealing with mobile nodes and specifically mobile nodes equipped with multiple interfaces. The following section deals with these challenges, outlined by Hui et al [11].

3.3.1 The Mobile Device

Past and current mobile devices are specifically designed to operate using one type of access technology. Emerging and future devices will need to incorporate various access network standards to allow the device to connect to the 4G network and its varying coverage areas of each access network. This could be achieved in two ways. Firstly, current mobile devices tend to include a number of separate access technologies in their hardware. This of course requires more physical space for the hardware components of each access technology and more battery power is used.

A second method has been proposed to solve some of these problems - the Software radio. The physical hardware includes an antenna and the necessary detection circuits to transform the analog radio signal into digital form and the digital signal processor (DSP) that processes the digital signal according to the access network technology. Software radios, however, do not support all radio access standards at present, since each radio access standard is transmitted on different frequencies - a single antenna and detection circuit cannot currently cover all frequencies. A further obstacle is that current DSP technology does not support high enough decoding rates or bit resolution. It is clear then that although software radios are appealing, the current state of the art does not make them practical. It is for this reason that 4G devices will instead need to focus on decreasing hardware size and cost to fit multiple access technologies into small mobile devices, while battery technology will need to advance to allow higher energy densities to run multiple access technologies on the mobile device simultaneously.

In the evaluation platform for this thesis, an emulated mobile device is created, since multiple interfaced devices are not currently available that would allow effective testing. Issues like battery life are not addressed since this will not affect the outcomes of the test scenarios.

3.3.2 Access Network Discovery

In the 4G network, the mobile device should be able to choose which access networks it connects to in order to perform the task that is running on the device. In the case of current 2G and 3G networks, the device is connected to a single base station based on a single messaging protocol that allows the mobile device to monitor the connection quality or signal strength. A handover is performed to a different base station with a better connection, if one is available. In the case of 4G networks, the mobile device is presented with a number of different access networks with different access protocols, coverage area and hand-off abilities. Again, this problem can be partially solved by using software radios that can reconfigure themselves to a new network access protocol when the device wants to connect to it. But it has already been established that software radios are not yet feasible with current technology. This issue thus remains unresolved. If considering a current mobile device with multiple separate access technologies built in, it would be necessary for the operating system of the mobile device to perform the scanning, connecting and disconnecting processes necessary to associate the device with changing access networks. Here it is clear that a management scheme for multiple-interfaced devices will need more than just a network protocol such as Mobile IPv6, since Mobile IPv6 controls packet flows to and from the device at the network layer, not processes like connections to access networks. This leads to the next challenge in 4G Networks: Access Network Selection.

3.3.3 Access Network Selection

Each type of access network will carry different types of traffic better than others. For example, the GSM network is optimised for reliable transmission of voice, while only offering low-bit rate data services. The 3G network provides high-bit rate data, but may be too costly for large downloads of information. A free Wireless LAN network would be ideal for very high-bit rate downloads, but provides isolated coverage at best. So it is clear that each access network would need to be used based on its advantages and disadvantages

to the user, as well as taking into account the geographical coverage available. Thus it is necessary to provide the user with choices as to which access network is used for each task required, bearing in mind that not all access networks will be available at all times. It is thus important that some form of preference system be available to the user to define the usage of each access network. This preference system is discussed later in this chapter.

Access network selection is performed explicitly in the evaluation platform to simplify the testing process. Since both interfaces of the mobile node are under lab conditions and free of charge, a system to enable selection of interfaces was not needed.

3.3.4 Device Mobility

Since user devices in 4G networks will be mobile, it is important that these devices can connect to an access network anywhere and at any time. An important consideration of device mobility is hand-off management. Since the mobile device will be entering and leaving different access networks as it changes interfaces and moves geographically, it is important to have a scheme that will assign globally routable IP addresses to the mobile device that will allow it to connect to the Internet. Mobile IPv6 is a suitable protocol for handling the task of assigning an appropriate IPv6 address to the device as it connects to a new access network, while allowing the device to be reachable at all times by other nodes on the Internet. This is achieved by assigning a global Home Address to the device that acts as its universal identifier on the Internet. An in-depth discussion on Mobile IPv6 is not needed to determine that Mobile IPv6 was initially designed to support a device with a single interface. Although some schemes have been proposed to modify Mobile IPv6 to support multiple interfaces on a single device, the issue remains that the protocol extensions themselves do not provide the intelligence for interface management, relying on the user to set rules for packet redirection. The protocol extensions simply allow the device to have one or more unique addresses assigned to each of its interfaces. With the potential to perform hand-offs within a network (horizontal hand-off) and between networks (vertical hand-off), problems arise when the user is in the middle of a multimedia session, be it voice, video or data streaming. Since the hand-off process is not instantaneous, there will be a noticeable break in the multimedia session as connections are interrupted during the hand-off process. While some proposals aim to minimise this hand-off delay, the problem still remains that some break in data transmission will occur. It is expected that in the case of a device equipped with multiple interfaces, this data transmission interruption could

be further minimised by using a hand-off scheme that would setup a new connection to a new access network on another interface before breaking the connection on the existing interface on which data is being transmitted. The evaluation platform presented in this thesis aims to achieve this seamless handover.

3.3.5 Network Infrastructure

When a mobile device is connected to a single access network, that access network is able to deliver a certain Quality of Service (QoS) guarantee if such a service is supported. In the 4G scenario, many access network types are supposed to work together to provide universal access to the user. It is then clear that there will be conflicting Quality of Service policies when different types of access networks are forced to connect together for a user's end-to-end session. In the case of non-IP based access networks like GSM, these networks are highly optimised for voice traffic and so the QoS policies of these networks are designed to maximise voice delivery. With IP based networks such as Wireless LAN and Ethernet-based networks, the QoS policies are designed with high speed data in mind. The challenge comes when trying to provide an end-to-end QoS guarantee when these diverse networks are connected together. It would be ideal if all these networks supported IP and a single QoS policy framework and much research is needed to solve the interworking issues. For this reason, an IP network with no QoS policy is used for the evaluation platform.

3.3.6 Security and Privacy

As is the case with network infrastructure, each type of access network has different implementations of security and privacy policies. As an example, GSM networks are designed to secure voice traffic against interception and interference from malicious parties. However, in a 4G network, having unique security policies for each access network is not practical since interworking becomes problematic. It is also important to realise that each security scheme will have differing hardware requirements to encode and decode data, making it necessary to build mobile devices with enough processing power to process multiple security schemes simultaneously. This increases device costs, size and battery power consumption. Proposals such as IPSecv6 have attempted to address these issues if access networks support IPv6. This proves problematic in a non-IP network like GSM. Another issue with interconnection of various access networks is managing user access to the network. The

authentication measures necessary to connect to a GSM network vary greatly from the types of authentication required in a Wireless LAN network for example. Again, it is a challenge to make these authentication systems interwork seamlessly, especially in a device with multiple interfaces, since each interface would require differing authentication policies. No authentication policy is required for the evaluation platform, since lab conditions emulate a secure network with a pre-authenticated user.

3.3.7 Fault Tolerance and Reliability

In an isolated access network, fault tolerance will be based on the topological design of that access network. Since the access network will consist of various levels such as base stations, switches and links to external networks, there are multiple levels for faults to occur. For this reason, each access network will incorporate a certain amount of redundancy and backup measures to provide the maximum uptime at any point in the network. In the case of 4G systems, when all these access networks are brought together, the potential for faults to occur is greatly increased since unavailability of one component may affect other access networks during an end-to-end session. To achieve better fault tolerance in a 4G system, it is possible to either provide some form of hierarchical network topology with redundant components, or use the very nature of different access networks able to co-exist and overlap geographically as a measure of redundancy. If these measures are implemented, then the problem of fault tolerance is transformed into an issue of over-provisioning and how much over-provisioning is necessary to increase 4G network reliability. But over-provisioning is a delicate balance between cost and benefit. Have network infrastructure sitting idle in one part of the network while another part is experiencing congestion is not ideal. Therefore fault tolerance schemes need to balance their cost against the performance benefits they provide. In this thesis, two unreliable links are combined to form a single channel for a data stream. This single combined channel experience significantly less packet loss than each separate interface. The fact that two interfaces are present on the Mobile Node means that should one be unavailable, the other interface can substitute as a backup interface.

3.3.8 Billing and Accounting

Users are accustomed to today's simple method of being billed for using access networks, like cellular networks. Users are charged based on call duration or the amount of data

transmitted. In a 4G system, each type of access network would have a different charging and billing scheme in place, or no data charge as is the case for a free Wireless LAN network for example. This is complicated further if multiple interfaced devices are introduced in a 4G system. Firstly, different interfaces would be charged different amounts depending on the access network they were connected to and secondly, the user would need to keep track of what multimedia streams are sent over each interface due to cost consideration - for example, the user wouldn't want large volumes of data to be sent over an expensive 3G link when a free Wireless LAN connection was available. In order to solve these billing issues, there are some proposed ideas that would greatly simplify the user's experience. If some form of brokerage service was available, the user could subscribe all their service providers to the brokerage service who would keep track of the various bills. The brokerage service would then send the user a single bill on behalf of the service providers. However, this does not solve the problem of keeping track of the multiple charging schemes for each access network. Ideally, service providers should implement flat rate charging for access to their network, but earn additional revenue from charging for content delivered to the users. This would largely eliminate the problem of data costing different amounts on different access networks. But since service providers would like to encourage users to use access networks that used cheaper hardware and had more abundant bandwidth available, for example a Wireless LAN as opposed to a 3G base station, there would still be different charges for each access network. It is then up to the user to make use of a proposed profile database to determine which types of multimedia are transmitted over each interface in the multiple interfaced device. This complicates the user's experience somewhat, but would provide large scope for advanced users to tailor their multimedia sessions to the prevailing access networks available. For this thesis, the evaluation platform uses two interfaces with no billing or accounting. Certain 4G networks will be free of charge such as company LAN's and Wireless LAN's, so billing and accounting will not be an issue in these cases.

3.3.9 Profile System

From the overview of the issues arising in 4G networks, it is clear that in many network components there is a need to make specific decisions relating to which access network to connect to and which multimedia streams each interface on the mobile device will carry. Clearly there is the need for a control system that will make these decisions on behalf of the user. This control system can take the form of a profile of user preferences stored

in a database. This database needs to be stored on the mobile device, since applications on the mobile device will need to initiate multimedia streams and this means interface selection for these multimedia streams is needed. When an application wishes to initiate a new multimedia stream, for example a voice over IP call, the application would consult the interface preference database. The preference database would reply with the outgoing interface that the application is to use for the multimedia stream. In the case of the voice over IP call, the profile database may select the Wireless LAN interface of the mobile device.

With the capability of storing interface preferences comes the need for the user to actually create this database during the configuration of their mobile device. While advanced users may find this simple, novice users may not want to deal with the complexities of configuring every aspect of their mobile device in such detail. Herein lies the need for a default profile database with basic settings that will still provide the benefits of a multiple interface device, but hide the complexities from novice users. To achieve this, service providers could supply mobile devices already configured with preference settings optimised for their networks. More advanced users could then tailor their mobile devices according to their specific needs. It should be noted that novice users might not need the functionality provided by a multiple interfaced device, so simple default settings would be sufficient for them. For example, a service provider could setup the mobile device to simply select any available access network will this lowest cost for whatever application the mobile device was currently running. More advanced users could tailor every application, multimedia stream and access network to their specific needs to achieve even more benefits in performance and cost.

Since mobile devices run different operating systems and software and are equipped with different types of interfaces, the profile database would be different for each type of mobile device. This means that a generic profile database would not suit all mobile devices. Since the applications would be responsible for the initiation of multimedia sessions, these applications would need to communicate with the profile database on the mobile device to determine what preferences the user has set for their multimedia stream. Many applications are designed for a large variety of mobile devices and a specific operating system so there is a need to have a set of standards with regards to communicating with the profile database stored on the mobile device. This would need collaboration from software and hardware vendors to develop a protocol for creating, storing and retrieving information for the profile database.

Preferences could be set for each application by the user. For example, an instant messaging client with voice and video capability could detect the interfaces available on the mobile device. In the preference settings, the user could configure the client to use the GSM interface for voice, the GPRS interface for text and the Wireless LAN interface for video. This eliminates the need for a separately stored database on the mobile device with all interface preferences for all interfaces and applications. A second alternative is for the operating system of the mobile device to store preference information for different types of applications and multimedia. The operating system vendor would need to create a series of Application Program Interfaces (API's) to allow applications to communicate with the profile database. Again the need for standardisation arises.

To support legacy applications which are not aware of multiple interfaces, a third option would be to include preferences for interface usage at the network layer of the mobile device. The device would need to inspect each packet for header information that would assist in routing the packet to the correct interface for its traffic type. The drawback of this solution is that non-IP interfaces like GSM would not be subject to this per-packet inspection. Such a solution would however work on a device exclusively using IP-based access technologies.

Whichever method of profile database implementation is ultimately used, the need for storing user preferences is clear. With the wide variety of applications, multimedia and access networks available on a multiple interfaces device, it is critical that the user's preferences are adhered to in order to maximise the benefits achievable in a multiple interfaces environment.

A profile system was not specifically implemented on the evaluation platform. All decisions on packet routing are pre-configured into the router, since each test scenario has specific requirements.

3.4 Multiple Interface Functionality

It is now appropriate to discuss the types of functionality possible when a mobile device has more than one interface. For simplicity, the evaluation platform deals with a mobile node equipped with two interfaces. This makes demonstrating multiple interface functionality clearer than with a device equipped with three or more interfaces. In order to achieve maximum benefit, the two interfaces must be of different access technologies. Having a

device with two of the same type of interfaces will not provide the full benefits possible with multiple interfaces. Should no coverage of that type of access network be available in a certain location, neither interface is able to connect. With two different interfaces, overlapping coverage is necessary for both interfaces to be able to connect to their respective access network, but in urban areas this should be widely available. A device with two different interfaces 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. This is the first scenario implemented in our evaluation platform.
- 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. This scenario is investigated in Chapter 4 and 5.
- 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. While this scenario is not specifically investigated in the evaluation platform, the benefits are clear.
- 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. This scenario is implemented on the evaluation platform for further analysis.
- 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. This scenario is also investigated.

These are possibly the most beneficial functions that would be possible with a multiple interface mobile device. However, investigation into the performance of each scenario will

show that while these benefits are possible, there are some problems encountered when packets are rerouted over two different interfaces. The results from Chapter 5 will illustrate these benefits and the problems they face.

3.5 Requirements of the Evaluation Platform

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. Since Voice over IP has very strict requirements, this type of traffic is used to perform the test scenarios.

3.5.1 Real time Traffic Constraints

Real-time traffic experiences the greatest performance degradation and loss of quality if packets arrive late or out-of-order at the mobile node. Therefore this type of traffic is used in the test scenarios. The two types of real time traffic that will commonly be streamed to a mobile devices are voice and video. Video will generally be buffered to some extent to minimise artifacting and quality loss due to variations in packet arrival time and arrival order. Even when a live video feed is being streamed to the mobile device, a small buffer of a few seconds will not affect the perception that the user is watching a live video feed. Live television broadcasts are often delayed by a few seconds for editing and commercials. For these reasons, video streaming was not examined during each test scenario. However, when a real-time voice conversation is taking place on the mobile device, delays of even a few hundred milliseconds will severely disrupt a smooth conversation. For this reason, it was decided that the performance requirements of Voice over IP would be used to determine the success and quality of each test scenario. Voice over IP has strict requirements for packet latency and packet jitter [12]. Packet latency is the time between a packet leaving the sending node and its arrival at the receiving node. This 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 is the second important parameter for Voice over IP. Packet jitter is the delay variation between each packet arriving at the receiving node. If packet jitter is low, then packets arrive at regular intervals with very little variation. If packet jitter is high, then packets arrive at seemingly random times. Voice over IP codecs generally package the speaker's voice into small packets, each containing a fraction of a seconds worth of voice. If packet jitter is high, then large breaks will occur between these small samples of voice, breaking up the conversation. For quality Voice over IP communication, 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 maximums outlined here will be used to determine if the scenario would support quality Voice over IP transmission over multiple interfaces.

3.5.2 Packet Sequencing Requirements

Low latency and low jitter are important for quality communication, but packet sequencing does also play 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. If it is found that packets are arriving out of order, then these out of order packets could be discarded by the application running on the mobile device. While this is the worst case, many applications and protocols can tolerate packets arriving out of order by using a buffer. Any out-of-order packets are held in the buffer until the next packet in the sequence arrives. The packet flow is then released from the buffer once the correct ordering has been established. But these buffers add further latency and jitter to the packet stream. For this reason, the evaluation platform specifically records out-of-order packets. These buffers need to take into account that Voice over IP can only tolerate 100ms of latency and 75ms of jitter end-to-end. If packets are to travel over different interfaces with different delay and jitter, then when these packets are recombined at the mobile node it is expected that under certain conditions these packets will arrive out-of-order. It is clear that packet order is thus an effective measure of the success of a test scenario.

3.6 Chapter Discussion

The expected evolution of 3G networks into 4G networks have been briefly discussed with emphasis on how a multiple interfaced node will complicate the implementation of the next generation of networks. With such complex implementation issues to solve, it is very important to investigate the possible functionality that a multiple interfaced device will bring to the user. For this reason, the evaluation platform was developed. To properly determine the success of the simulated multiple interface scheme, it was decided that Voice over IP parameters would be used as benchmark figures to compare the test results against. With these performance figures, the complexity of a multiple interface management scheme can be justified if the test results show that the benefits of such a scheme justify the disadvantages.

Chapter 4

Design and Implementation of Evaluation Platform

The previous chapter presented an overview of the issues that need to be considered when designing a multiple interface enabled network with emphasis placed on the type of functionality desirable for an effective and useful scheme. When designing the evaluation platform, it was decided that the most useful data would show whether the potential functionality of a multiple interface scheme would in fact be useful. A complete end-to-end system is not required so it was decided that the evaluation platform would instead implement several scenarios that would be possible in a multiple interface environment. The testbed would be hard coded to perform each test scenario in isolated test runs and all intelligence would be pre-built into the configuration script. This means that there was no need to implement any decision policies from user preferences or exchange any preference information between the mobile node and router.

This chapter begins by describing the testbed hardware to emulate a basic multiple interface evaluation platform, consisting of a Corresponding Node, Router and Mobile Node. The operating system and software router are then described and their suitability for the tasks is discussed. The rest of the chapter describes the setup and configuration of each test scenario, including the scripts used to configure the routing software and the program used to generate the packet flow that would be used to perform the tests. The data collection process is also described and the result gathering method is outlined.

4.1 Multiple Interface testbed Components

Based on the literature studied in Chapter 2, a clear trend emerges in terms of the hardware topology that a multiple interface scheme would require. Each component is necessary to perform the routing that allows packet flows to travel from the sending party (Corresponding Node) through the network and over the correct interface to the user's mobile device. The typical components of a multiple interface scheme are shown in Figure 4.1:




Figure 4.1: Critical Components of a Multiple Interface Scheme

- The mobile node. This device is equipped with multiple interfaces and is carried around by the user. The mobile node needs to connect to various access networks, so each interface built into the mobile node needs coverage by an appropriate access network.
- 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. These access networks are provided by different service providers and network operators. Two wireless access networks are shown in Figure 4.1.

- Router. The router acts as the redirection point for packet flows that are destined for the mobile node. The router's function is to carry out packet routing decisions based on the preferences set by the user for each type of traffic flow and interface. The router will be the point at which all packet flows destined for the mobile node converge, in a simple scheme with only one router. For the evaluation process, only one router is required, but in larger schemes, the number of routers present could scale with the size of the network. This could provide advantages such as distributed work load, redundancy and geographic placement to optimise packet routing, but not without added complexity and implementation issues arising from the need to co-ordinate the routers.
- 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. Traffic in the reverse direction will be sent from the Mobile Node's interfaces to the Router and will leave the Router on this interface.
- Corresponding Nodes. Any node sending traffic to the Mobile Node or receiving traffic from the Mobile Node is classified as a Corresponding Node. These nodes are agnostic to the fact that the Mobile Node is equipped with more than one interface. It is the Router's responsibility to receive multiple traffic flows from Corresponding Nodes and route this traffic to individual interfaces on the Mobile Node. The Corresponding Nodes communicate with a single IP address that is associated with the Mobile Node and this IP Address is reachable on the Router's External Interface, acting like a Home Address in the case of Mobile IP. Corresponding Node's will appear to be communicating with this single Home Address, where in fact the Mobile Node has a different IP address associated with each interface and only the Router is aware of these addresses.

These components are the minimum requirement for the Multiple Interface evaluation platform. These components will be realised in the form of suitable testbed hardware and software in the rest of this chapter.

4.2 Testbed Hardware

Each of the components in the previous section are emulated with hardware. The choice of hardware components was based on achieving maximum results from common off the shelf hardware available, since these components are readily available, well documented and cost effective. A break down of each components follows:

4.2.1 Mobile Node

The evaluation platform requires a Mobile Node equipped with more than one interface. Since these devices do not yet exist in a form that would allow the correct functionality required for the evaluations, a suitable substitution needed to be assembled. The Mobile Node therefore is emulated by a standard PC equipped with multiple interfaces. The PC itself was not physically mobile, but the evaluations did not require mobility as this was not being tested. For the purposes of evaluation, it was decided that two interfaces would be suitable and provide sufficient useful data in determining the success of the evaluation scheme.

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 and further test the interaction of different access hardware under multiple interface test scenarios. It is also important to include the wired Ethernet connection to represent a hypothetical user connecting to a wired access network, since these access networks will be included in 4G Networks. The wireless LAN connection is representative of general wireless access networks having increased packet delay, packet loss and signal strength variation compared to wired access networks. These varying conditions influence the test results and provide insight into the interactions of wired and wireless connections.

4.2.2 Router and Access Networks

A router needs to be both configurable and programmable to perform the packet routing tasks required in a multiple interface scheme. A router also needs to be equipped with different access technologies to simulate the different types of access networks available

to the user. For these reasons, it is simplest to create a router from an ordinary PC. 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) depicted in Figure 4.1 provided by a wired 100BaseT Ethernet connection directly to the Corresponding Node described next. This Ethernet connection carrying IP packets is chosen for the External Interface since in a real world implementation connections to the Internet and thus Corresponding Node's will be over the IP-based Internet. It therefore follows that Interface One and Two also use the IP protocol and in fact the entire testbed is IP-based. In a real world implementation, there could be non-IP networks at certain points in the multiple interface enabled network core and access network, but it is sufficient to use IP exclusively in this evaluation framework, since this will provide a simple but powerful method of performing packet routing without the added complexities of dealing with non-IP entities.

4.2.3 Corresponding Node

The Corresponding Node is the entity from which traffic will originate in the evaluation platform. To perform simple packet flow generation, a PC is again sufficient and is representative of PC's that would be connected to the Internet in a real world scheme. 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.

4.2.4 Ethernet and IP Address Configuration

The testbed uses two Ethernet-based interface types, 802.11b Wireless LAN and 100BaseT wired Ethernet. Each Ethernet interface by default is manufactured with a unique Ethernet Address. The Ethernet Address is used by the Ethernet protocol to establish which hardware interface is referred to by an IP address in the Ethernet Frame Header. This allows the Ethernet frame to be delivered to the correct hardware interface on a common Ethernet network. The IP address of each interface is however determined by configuration. Each interface in the testbed was assigned a unique IP Address based on the subnet associated with each pair of communicating interfaces. In Figure 4.2 the entire testbed is shown, with the Linux 'eth' identifier, Ethernet Address and IP Address assigned to

each interface. It was decided that a pair of communicating interfaces, for example the two Wireless LAN Interfaces, would exist on their own subnet, to symbolise the real world scenario of each access network having its own IP Address subnet to allow connected users to each have their own unique IP Address in the subnet.

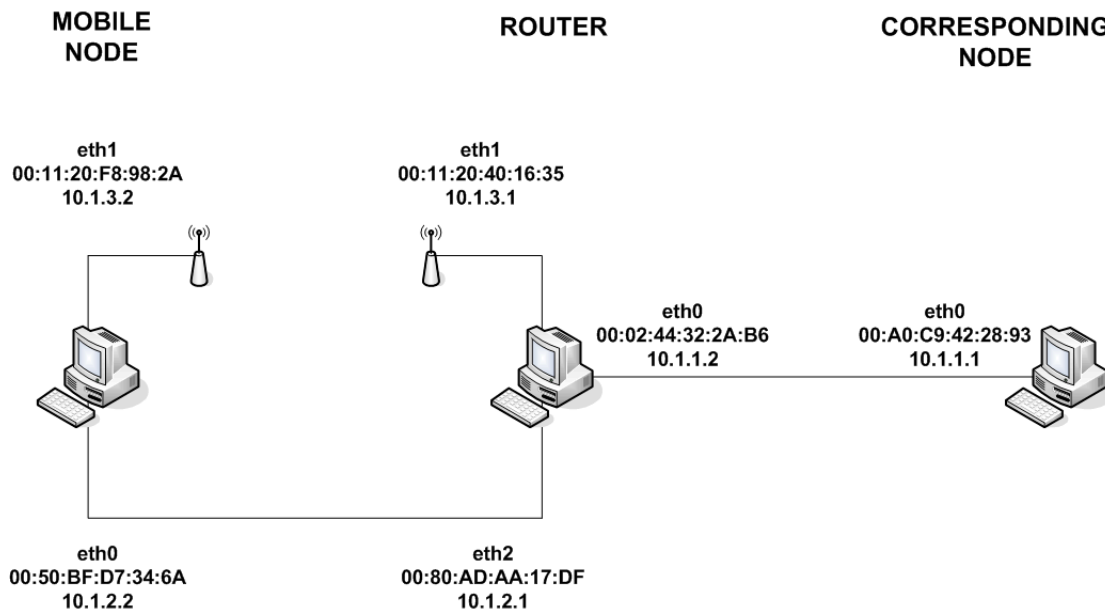


Figure 4.2: testbed Hardware and Ethernet/IP Address Assignments

4.3 Operating System

The three hardware components are basic PC's and therefore need an operating system. When choosing an operating system, it is important that it allows the user to configure every aspect of the system in as much detail as is necessary. For this reason, Microsoft Windows would not be suitable, as this operating system tends to hide too much functionality from the user. The various versions of Linux are very appealing in this regard, since every part of the operating system is completely accessible and configurable due to the open source nature of the product. Linux was thus chosen, more specifically, Debian Linux version 2.4.32. A further motivator in the choice of Debian Linux specifically, was because the routing software chosen for the evaluation platform recommends Debian Linux. Moreover, when configuring a network and various Ethernet interfaces, Linux allows the user to change and configure every aspect to their requirements. For this reason, Debian

Linux proved an ideal platform on which to build the testbed. Each of the three PC's had Debian Linux installed and the Ethernet interfaces were all setup according to Figure 4.2. The 'eth' identifier shown for each interface is the Ethernet identifier that Linux assigns to each network card connected to the PC. This 'eth' identifier is used when referring to a specific network card during configuration and testing.

4.4 Router Software

The router at this stage consists of a PC with three network card interfaces and Debian Linux installed as the operating system. In this form, it is not yet able to route packets between the three interfaces according to the requirements of the evaluation platform. Suitable routing software was needed that would allow packets to be intercepted individually from each network card and then be processed and routed according to the test scenario underway. A suitable software router is the *Click Modular Router* developed by MIT [13] and documented in Eddie Kohler's PhD Thesis [14]. This software router allows the user to build complex router configurations and packet processing entities out of simpler modules called *elements* that are connected to one another resembling a flow diagram. Even on a modest 700MHz Pentium III computer, up to 435 000 64-byte packets can be routed every second, well beyond the requirements of the evaluation platform. Due to the open source nature of Click, other research groups have developed additional *elements* and configurations for free distribution. The details of the configuration scripts for each test scenario are presented later in this chapter.

4.4.1 Packet Modification

Since packets will be routed over different interfaces according to the test scenario currently scripted into the Click Router software, they need to be modified to suit the interfaces they will be transmitted and received on. Refer to Figure 4.3. A packet is generated on the Corresponding Node (CN) with a Source IP address of 10.1.1.1 (the IP address of the CN) and Destination IP address of 10.1.1.2 (the IP address of the Router's External Interface or Incoming Interface). The packet is intercepted by the Click Modular Router software directly from the network interface card. If the packet is destined for eth1 on the Mobile Node (MN), the Source IP Address needs to be modified to the Router's eth1 IP Address (10.1.3.1) and the Destination IP Address needs to be changed to the MN's eth1 IP Address

(10.1.3.2). This ensures that the newly modified IP packet is accepted by the MN. A similar process is shown for eth0 and eth2 between the Router and MN. The IP Addresses have been modified, but the UDP Header for each packet also needs to be modified as the UDP Header also contains Source and Destination IP Addresses as well as Port numbers. The Router modifies the UDP Headers to match the IP Address changes already shown. Finally, the IP Checksum and UDP Checksum of each packet needs to correspond to the newly modified addresses. These two checksums are recalculated by the Router before packets leave the respective interface. Finally, the correct Ethernet header is added to each packet to correspond with the sending and receiving Ethernet addresses of the two interfaces that are communicating. All these processes ensure that a packet leaving the Router is topologically correct and will be accepted by the MN as if the packet originated from the Router itself, not the CN.

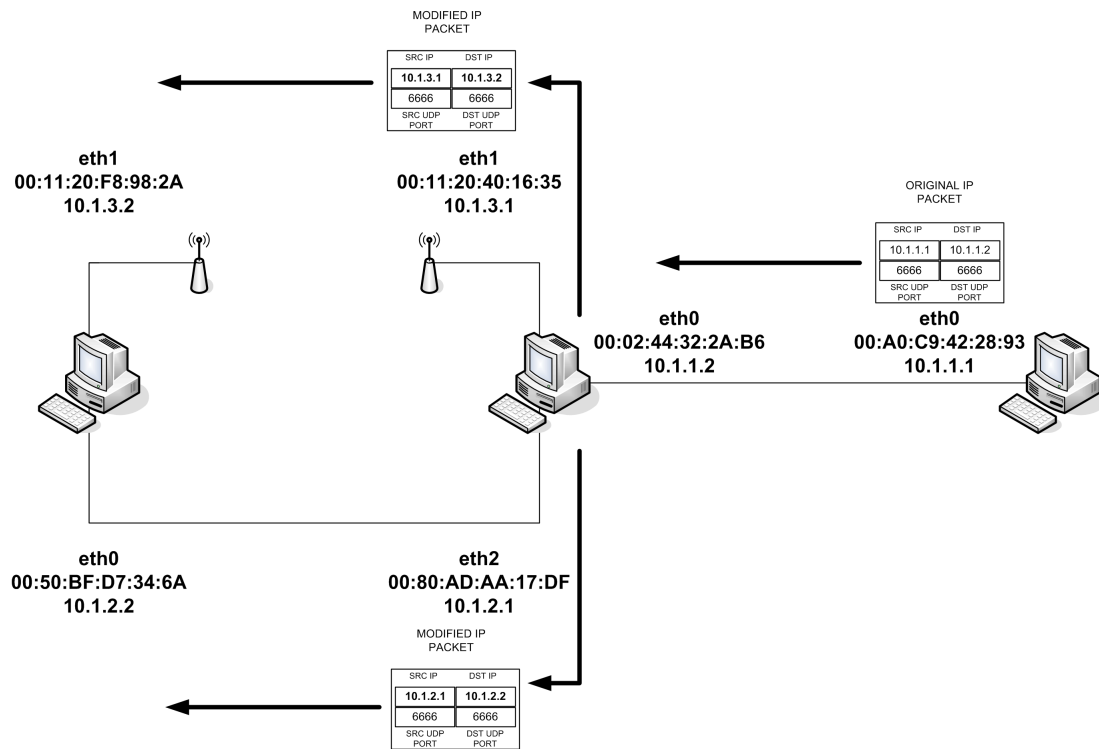


Figure 4.3: Packet modification diagram for testbed.

4.5 Packet Generation

For each test scenario, a steady flow of packets needs to be sent from the Corresponding Node to the Router. This packet flow would need to have certain parameters for each test. It was decided that all tests would consist of a stream of 172 byte UDP Packets as this is a common packet size and transport protocol for streaming voice and video traffic over IP networks. For the purposes of this evaluation platform, it was decided that the focus would be on measuring the performance impact of each test scenario on real time traffic only, since this type of traffic has strict delay requirements with little tolerance for packet loss or high latency and packet jitter. The UDP protocol will not guarantee delivery of each and every packet. For this reason, it would be clear in each test's results whether all packets were delivered timeously and reliably and whether the routing performed during each test scenario had influenced the packet flow.

To generate a stream of UDP packets with certain parameters for each test, a simple C based program was used, developed by Albert Hasson [15], and modified for each test scenario. Parameters such as number of packets sent, inter-packet transmission delay and UDP port numbers were set before each test scenario was conducted. To keep track of packet order at the receiving Mobile Node, each packet was numbered and this value was stored in it's data field. The packet flow was sent on the Corresponding Node's outgoing Ethernet interface to the Router's incoming Ethernet interface. Once the Router had received each packet, it could be modified according to the test requirements of each scenario.

4.6 Data Collection

In order to analyse the data produced during each test scenario, it was important to collect information about each packet received by the Mobile Node. Linux is equipped with this functionality in the form of the *tcpdump* command. This command has many options for data display and collection and was run in waiting mode before each test was started. As each packet was received by the Mobile Node, *tcpdump* would display important parameters like inter-packet arrival time, packet headers and packet size, as well as the total number of packets received. This information was used to check that the test scenario was performing as expected. Once it was established that the scenario was performing correctly, the packet data was recorded.

To record each packet received by the Mobile Node, a C based program developed by Albert Hasson called *dgram-rec.c* was modified to produce columns of information about each packet that would be compatible with spreadsheet programs. Each packet had its sequence number embedded in its data field at the Corresponding Node. This number was extracted and displayed in the first column. Any out of order packets were labelled with a '1' in the second column. The inter-packet arrival time was recorded in the third column. The total number of out of order packets was recorded as a running total in column four and the interface the packet was received on was recorded in the last column. See Table 4.1 for an example of *dgram-rec.c*'s output. All five columns of information were written to a text file to be imported into Microsoft Excel. The data was used to generate graphs and figures for Chapter 5.

Table 4.1: Example output of *dgram-rec.c*

Packet #	Out-of-Order	μsec	Total Out-of-Order	Interface #
1	0	200	0	1
2	0	205	0	2
4	1	198	1	1
3	0	202	1	2
5	0	210	1	1

4.7 Test Scenarios and Configuration Scripts

As outlined in Chapter 3, it was decided that four test scenarios would be implemented on the testbed to emulate possible functionality with a multiple interface scheme. The results of each test scenario will be used as an indicator of the practicality and usefulness of each example of multiple interface functionality. The four test scenarios were transformed into a configuration script for the Click Modular Router software and were executed individually for each test. The details of each test scenario and the corresponding configuration scripts follow.

4.7.1 Mid-Flow Interface Transfer

When considering the possible functionality of a multiple interface scheme, the first possible benefit that needs to be investigated is the case of mid-flow interface transfer. Since the mobile device is equipped with two interfaces, a handover of packet flow from one interface to the other could largely mitigate the problem of vertical handover in the access network - traditionally, with a single interfaced mobile device, the connection from the one interface to the previous access point must be broken before a new connection is made to the new access point. Various schemes aim to reduce this delay as much as possible, but there will still be a period of time when the single interface is not receiving packets (when it is disconnected from the previous access point and associating with a new access point). Now considering a mobile device with two interfaces, the potential immediately exists to have both interfaces associated with their respective access points simultaneously or if only one interface is active, the second can be activated and associated with an access point before the first interface disconnects - assuming that both access points are available simultaneously. Once both interfaces are ready to receive packets, a packet flow handover can take place from the first interface to the second. The performance of such a handover needs to be investigated.

For this test scenario, the Corresponding Node was configured to transmit a continuous stream of UDP packets of 172 bytes each to the Router with a fixed inter-packet transmit delay. This delay time was varied in separate test runs to see how packet sending rate affected the interface transfer handover point. The router was programmed with the configuration script shown in the flow diagram in Figure 4.4. For all four scenarios, all interfaces have IP addresses and are ready to transmit or receive packets. In order to reproduce a mid-flow interface transfer, a Click *element* called SplitFirst() was used. This *element* has a single input and two outputs. The input receives IP packets from output port 2 of the Classifier *element* and increments a counter until a threshold value is reached. Before this threshold is reached, all IP packets received are emitted on output 0. The threshold was set at 100 packets - an arbitrary value which produced consistent results but kept the size of the results data manageable. Once this threshold is reached, all further IP packets are emitted on output 1.

The SplitFirst() *element* therefore creates two separate packet flow paths and allows each path to be routed to a different outgoing interface of the Router - this will achieve the goal of forcing a packet flow handover from one interface to the other. IP Packets flowing down

either path are stripped of their IP and UDP headers, leaving just the packet data. A new IP and UDP header is appended with the correct IP source and destination addresses and UDP port number necessary to route each packet over the respective outgoing interface from the Router to the Mobile Node. Referring to Figure 4.4, it is clear to see that the first 100 packets are sent to 'eth1' while all further packets are sent to 'eth2', using the final element in both paths. 'eth1' refers to the wired Ethernet interface of the Router and 'eth2' refers to its Wireless LAN interface.

Since a total of 500 UDP packets are sent to the Router by the Corresponding Node, the first 100 packets are routed over the 'eth1' (wired Ethernet) of the Router and received on 'eth0' (wired Ethernet) of the Mobile Node. Then all further packets are routed over 'eth2' (Wireless LAN) of the Router and received on 'eth1' (Wireless LAN) of the Mobile Node. All 500 UDP packets were detected and recorded by *dgram-rec.c* running on the Mobile Node. This entire process simulates a mid-flow interface transfer from the wired Ethernet link to the Wireless LAN link. But a handover in the opposite direction (Wireless LAN link to wired Ethernet link) also needed evaluating, so the Router script was simply modified to send the first 100 packets to the Wireless LAN link, then all remaining packets to the wired Ethernet link. The different delay characteristics of the wired and wireless links affect the handover point differently. Chapter 5 has detailed results of this scenario.

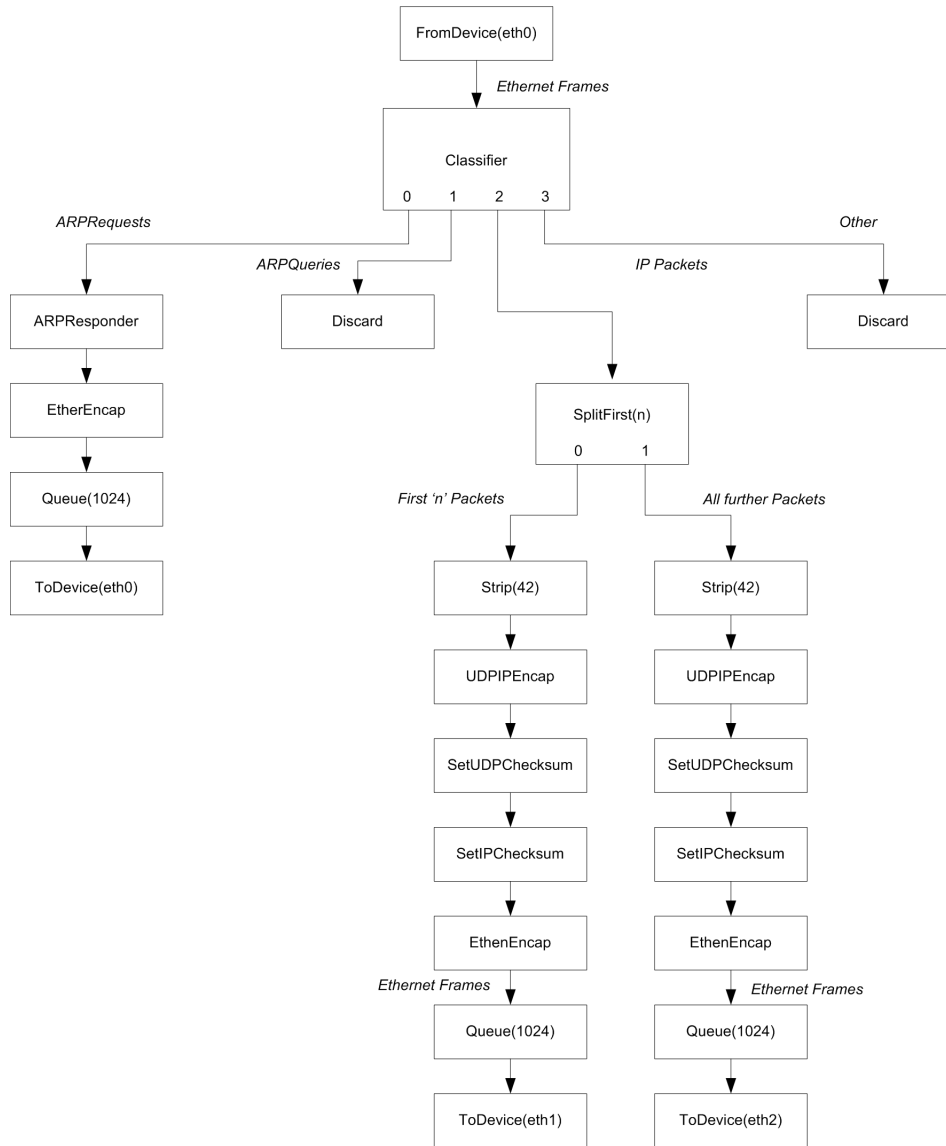


Figure 4.4: Click Element Flowchart for Mid-Flow Interface Transfer Scenario.

4.7.2 Dual Interface Aggregation

When considering a mobile device with two interfaces, it is desirable to have the ability to split a single packet flow and stream it over both interfaces - it is expected that this would increase the available bandwidth to somewhere in the region of the combined bandwidth of both interfaces. This process is referred to as *dual interface aggregation*. A logical method of splitting the single packet flow would be to send one packet to interface one, the next

packet to interface two, the next to interface one and so on. This scheme is called Round Robin switching where each alternate packet is sent to each interface in turn. A more efficient method of packet switching is to transmit the next packet to whichever interface is ready to send at that moment in time. The efficiency arises from the fact that each interface will have different transmission delays, available bandwidth, inter-packet delays and physical layer and link layer delays characteristic of the access technology. In the evaluation platform, two types of interfaces were chosen specifically to investigate this phenomenon - wired Ethernet provides a faster, more reliable connection with a lower ping time than 802.11b Wireless LAN does, in laboratory conditions. Evidence of this can be found in the results analysis in Chapter 5.

The Click Modular Router includes an *element* called RoundRobinSwitch() which will be suitable for the task of switching a single incoming packet flow over two interfaces. RoundRobinSwitch() is equipped with one input port and n output ports - in this scenario, $n = 2$. This element works by allocating the next incoming packet to the output port that is ready to receive a packet for transmission as detailed above - this method is more efficient than simply alternating packets between the two interfaces regardless of whether either is ready to send a packet or not. Figure 4.5 shows the configuration script for this scenario. As in Section 4.7.1, two packet flow paths are created, each stripping off the packet header, re-encapsulating the packet in a new IP and UDP header, and then transmitting it out of either 'eth1' or 'eth2' of the Router.

The RoundRobinSwitch() *element* therefore schedules packets in such a way that each is sent to the interface that is ready to transmit, eliminating the problem of having different delay characteristics for each interface.

For this test scenario, the Corresponding Node sends a stream of 500 packets of 172 byte length to the Router, with various fixed inter-packet transmission delays for each test run. The Router performs the Round Robin switching and packets leave the Router via either of its two Outgoing Interfaces. These two split packet streams are received by the Mobile Node and recorded using the *dgram-rec.c* program. The two split streams would need to be recombined on the Mobile Node and delivered to an application as the original single packet flow from the Corresponding Node in order for the application to function correctly as it is not expecting two split packet flows from two separate interfaces. For the purposes of this evaluation, this step was not necessary as it did not affect the results obtained in Chapter 5.



Figure 4.5: Click Element Flowchart for Dual Interface Aggregation Scenario.

4.7.3 Lossy Link Performance

Wireless access technologies are prone to interference, low signal strength and congested frequency spectrum together with the network congestion and packet loss of all networks in general. In the case of a single interfaced device, the prevailing access network conditions can affect the quality of connection that the user is experiencing and there is no alternate connection available. With a multiple interfaced device (two interfaces in this case), each access network will experience different network conditions. While at a certain time a wireless connection might be lossy, congested and have a high packet loss, a wired con-

nection might be operating more efficiently. The opportunity exists to use both interfaces to achieve a higher combined bandwidth, even though one interface may be more lossy than the other. More importantly, it is possible to send the same packet flow over both interfaces simultaneously. Since each interface will likely drop different packets, there will still be packet loss, but at a significantly reduced level when the two lossy packet flows are combined again at the receiver. Statistically, this can be proved using probability theory. If both links experience 5% packet loss, the probability of the same packet being lost on both interfaces is $0.05 \times 0.05 = 0.0025$. So the total packet loss of both links combined is only 0.25%. A link with 5% packet loss would be unusable for a critical real time voice or video application, since the quality of a Voice over IP stream is severely affected with even 1% packet loss. Yet statistically it is clear that combining the two interfaces with 5% packet loss each will produce a combined packet loss of only 0.25% - an acceptably low value for critical real time communication. Here two unusable interfaces have been combined into one path for reliable data transmission. While it might cost the user more to send the same packet flow over two interfaces, for critical applications such as an important real time voice or video feed, this cost may be justified.

For this scenario, both interfaces were configured to have the same packet loss. Five different packet loss rates were tested, 5%, 10%, 15%, 20% and 25%. To simulate a lossy link, both interfaces would be instructed to randomly drop packets. Click provides the `RandomSample()` *element* that will randomly drop packets at the rate specified. In order to duplicate the single incoming packet flow over both interfaces, the `Tee()` element was used. Packets are duplicated and sent to port 0 and 1 of the `Tee()` *element*, dropped at the specified % using the `RandomSample()` *element*, stripped of headers, allocated new IP and UDP headers as in the two previous scenarios, and emitted on 'eth1' and 'eth2'. See Figure 4.6. With this configuration, it is expected that of the 1000 packets that are sent to the Router by the Corresponding Node, a much lower packet loss will be experienced once the duplicate packet flows reach the Mobile Node over both interfaces.

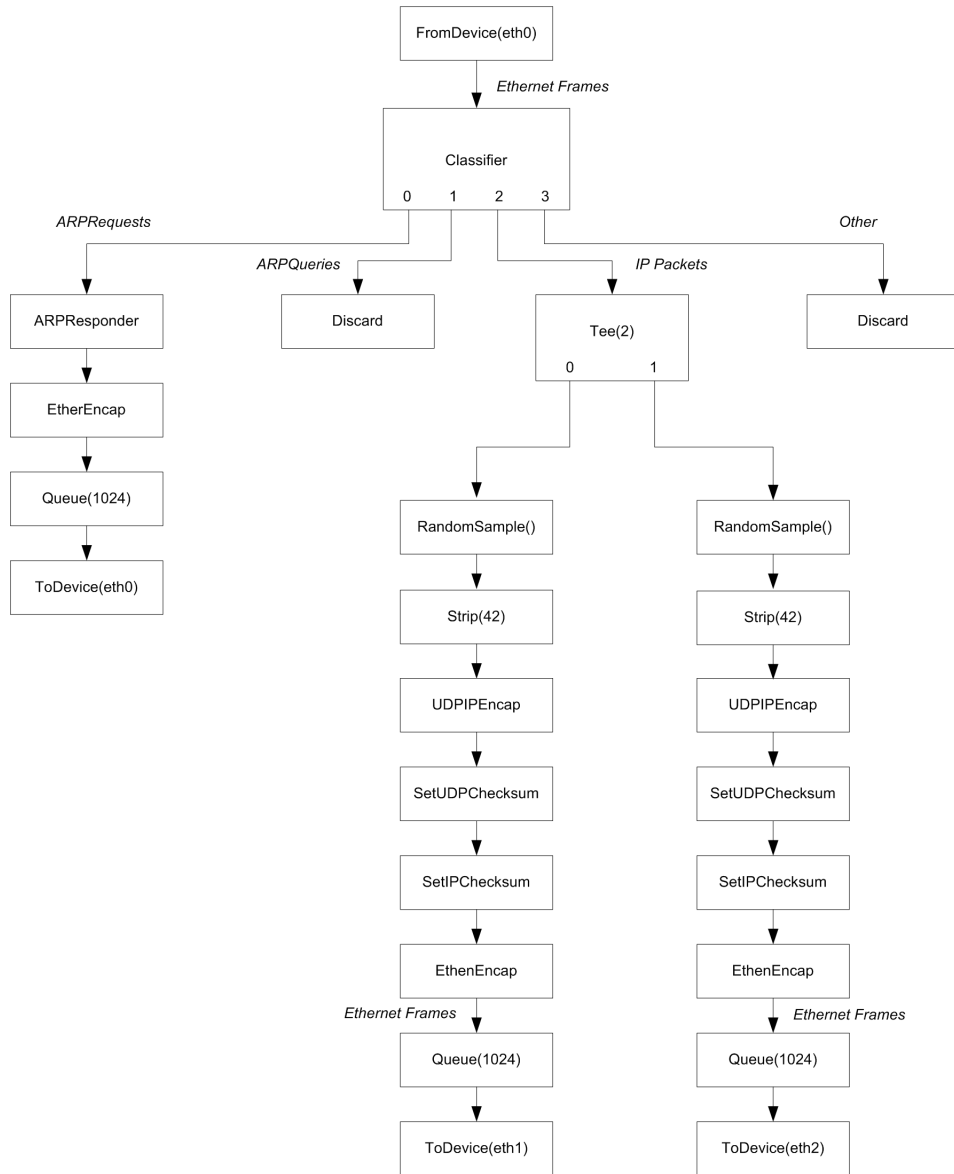


Figure 4.6: Click Element Flowchart for Lossy Link Scenario

4.7.4 Bandwidth On Demand

The final scenario under investigation is the case of a multiple interfaced device allowing the option of some form of Bandwidth-On-Demand scheme. For example, the user could be streaming a voice over IP call over one interface with enough bandwidth to support the call session. During the call, the user might add a video feed, but this video feed and voice

stream together require more bandwidth than the single interface is able to provide. So the device's second interface is activated and the two streams are carried by both interfaces in collaboration.

To emulate this type of scenario, an element was needed that would emit all IP packets out of one output port until a bandwidth threshold was reached. Above this threshold, all additional IP packets would be emitted out of the second output port. Click provides an ideal *element* in the form of `BandwidthRatedSplitter()`. It was decided that the threshold level for activating the second interface would be set at 384kbps. This is a common data rate for 3G networks and a voice over IP call commonly requires around 64kbps to support good sound quality - so the first interface can easily carry the voice over IP call. Referring to Figure 4.7, the two output ports of `BWRatedSplitter()` lead to *elements* that cause packets to have their headers re-written and emitted out of either 'eth1' or 'eth2' as was the case in the three previous scenarios.



Figure 4.7: Click element Flowchart for Bandwidth-on-Demand scenario.

4.8 Chapter Discussion

In this chapter, the details of each required component of the evaluation platform were discussed together with the implementation process. The three hardware components, Corresponding Node, Router and Mobile Node, were created from off the shelf hardware and an operating system- Pentium III desktop computers running Debian Linux. A suitable configurable software router, Click Modular Router, was found that will enable the evaluation platform to run four test scenarios emulating four important real world functions

possible in a multiple interface enabled network. In order to analyse the results generated by each test scenario, the *tcpdump* tool in Linux will be used and the resulting text files will be processed in Microsoft Excel.

In the next Chapter, the results of each of the four test scenarios are detailed, analysed and discussed. Each test scenario will provide evidence as to the advantages and disadvantages shown by the functionality under test - this data will be used to discuss whether each test scenario shows enough benefit to recommend that such functionality should be included in a multiple interfaced network.

Chapter 5

Performance Evaluation

The design of the multiple interface evaluation testbed was presented in Chapter 4. Four test scenarios were implemented in order to determine if their functionality would prove useful in a multiple interface network. The results of each scenario are discussed in this chapter and analysed for practicality, performance and suitability. These test scenarios will show that a multiple interfaced device does provide performance benefits over a single interface device to justify the added complexity required to support such functionality.

5.1 Results

For each of the four test scenarios, the three testbed computers were configured according to the parameters described in the beginning of each section below. The test scenario was then executed and the results were captured using the *dgram-rec.c* tool in Linux as described in Section 4.6. The output of this tool was then imported in Microsoft Excel to generate the graphs and figures below. The first scenario tested on the evaluation platform was Mid-Flow Interface Transfer.

5.1.1 Mid-Flow Interface Transfer

The Mid-flow Interface Transfer scenario detailed in Section 4.7.1 was executed as follows: Two separate mid-flow handovers were investigated. The first handover is from the Ethernet interface to the Wireless LAN interface. In order to test the handover process under

different loads, five separate tests were conducted with different packet sending rates from the Corresponding Node. The Corresponding Node was configured to generate a single stream of UDP packets for each run, with the parameters outlined in Table 5.1 and send them to the Router on its incoming interface. Table 5.1 shows the inter-packet transmit time for each packet sent by the corresponding node, the effective data rate of the packet stream and the number of Voice over IP (VoIP) streams represented by the packet flow for comparison purposes.

Table 5.1: UDP Packet Parameters from Corresponding Node.

Test Run	Inter-packet Transmit Time	Packet size (bytes)	Effective Data Rate	Effective # of VoIP Streams
1	20ms	172	8.6 KB/sec	1
2	2ms	172	86 KB/sec	10
3	1ms	172	172 KB/sec	20
4	500 μ s	172	344 KB/sec	40
5	200 μ s	172	860 KB/sec	100

The Click Modular Router was configured with the script shown in Figure 4.4 to perform the packet flow transfer from the wired Ethernet interface of the Router to the wireless LAN interface after 100 packets had passed through the Router. This would result in the Mobile Node receiving 100 packets on its wired Ethernet interface and thereafter 100 packets on its wireless LAN interface. See Figure 5.1.




Figure 5.1: Packet flow - Mid-flow Interface Transfer Scenario Test A

The scenario was run five times, once for each different set of parameters of the UDP packet flow from the Corresponding Node. The packet recording program on the Mobile Node `dgram-rec.c` recorded the arrival of each packet. In all five tests, all packets were received in order by the Mobile Node. This means that these results could be used to generate the graphs below showing inter-packet-arrival time vs packet number. If out-of-order packets were detected, then these graphs would not be an accurate representation of the packet arrival times and a different representation would be used - Elapsed Time vs. Packet Number. See Test B below where this different representation shows the out-of-order packets vs. time.

Results of Test A

The mid-flow handover point is not clear in Test 1 as the 20ms inter-packet transmit time is high enough to disguise the difference in average packet latency of the Ethernet and Wireless LAN interfaces. From the router script, the handover point has been set at packet 100, so this is where the handover takes place in all five graphs. A noticeable trend in Test 1 is that beyond the handover point at packet 100, there is an increasing amount of inter-arrival time spikes in the data, characteristic of a Wireless LAN. This is due to the fact that Wireless LAN's experience greater variance in the time it takes to transmit each packet, due to their physical characteristics. Ethernet has a more constant and much lower value of variance in packet transmit time and this will be more evident in the rest of the graphs.

Test 2 shows similar trends to Test 1, with no obvious handover point, although we know it occurs at packet 100. It is only in Test 3 that the handover point becomes obvious. Packets 0 to 100 show a near constant inter-arrival time of 1ms and at the handover point at packet 100, the inter-arrival time becomes erratic with a high variance, characteristic of the Wireless LAN interface. However, Voice over IP can tolerate 40ms of jitter, so while these erratic inter-arrival times may affect the amount of buffering needed to keep the stream constant, the jitter is acceptable at less than 1ms.

Tests 4 and 5 show distinct jumps in inter-arrival time at packet 100 and beyond. The handover point and the difference in Ethernet and Wireless LAN in terms of jitter is clear. It should be noted that even though the inter-packet sending rate is decreasing in each test, all packets are received in order at the Mobile Node. The handover from Ethernet to Wireless LAN is therefore a seamless process, with packet 101 received directly after




Figure 5.2: Test 1 - 20ms (left) Test 2 - 2ms (right)




Figure 5.3: Test 3 - 1ms (left) Test 4 - 500; μ s (right)

packet 100 with only an increase in inter-arrival time evident. This would result in an adjustment to the jitter buffer of the application running, but the Voice over IP stream will continue unaffected since jitter is still well below the 40ms limit. From these results, it is clear that a mid-flow handover from the Ethernet to the Wireless LAN interface is seamless, with no out-of-order packets and would not affect the quality of the Voice over IP stream. This is partly due to the fact that the Ethernet interface is capable over very low packet transmission time and very low variance in jitter, so that packet 100 is sent and received before packet 101 can be transmitted over the Wireless LAN interface. It is now important to investigate whether a handover from the Wireless LAN interface to the Ethernet interface will cause problems for the packet flow, since Wireless LAN show

significantly higher packet transmission time and jitter.




Figure 5.4: Test 5 - $200\mu s$

Results of Test B

The router script was modified to transmit the first 100 packets of the Wireless LAN interface and then all further 100 packets over the Ethernet interface as shown in Figure 5.5. The same five test runs were executed according to the parameters outline in Table 5.1.

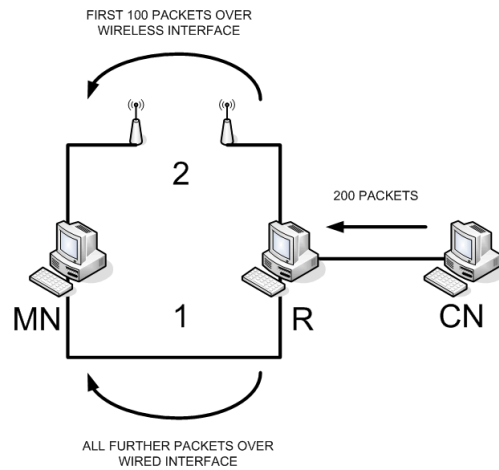


Figure 5.5: Packet flow - Mid-flow Interface Transfer Scenario Test B

From the test results, it was immediately clear the the Wireless LAN to Ethernet handover process was not seamless. Out-of-order packets were observed in a number of the test runs.

For this reason, it is no longer appropriate to plot packet inter-arrival time on a graph, since packets are arriving out-of-order. Consecutive packets received by the mobile node were not the same as the order of packets transmitted by the Corresponding Node, so inter-arrival time becomes largely meaningless. To show out-of-order packets, it is necessary to determine the exact time at which each packet is received by the Mobile Node. This time stamp can then be plotted for each packet. Out-of-order packets will arrive later than they should, so the graph will deviate from a straight line. A straight line graph will mean that all packets are received in order, so as packet number increases, time stamp increases. The graphs below show these characteristics more clearly.




Figure 5.6: Test 1 - 20ms (left) and Test 2 - 2ms (right)

In Figure 5.6, Test 1 and Test 2 both show no out-of-order packets, with only slight deviation from a straight line evident in Test 2. This shows that the packet sending rate from the Corresponding Node is not overwhelming the Wireless LAN interface, so all packets are received in order at the Mobile Node. The handover point for both Test 1 and 2 is not obvious, but occurs at packet 100. These two tests show that when a handover from a Wireless LAN to an Ethernet interface occurs, packet sending rates of 2ms inter-transmit time or more result in perfect transmission and no out-of-order packets.

In Figure 5.7, the first signs of out-of-order packets appear in Test 3. The handover point indicated by the box on the left has been magnified in the graph on the right. Here it is shown that packets 101 to 103 are received during the same period as packets 97 to 100. This overlap is a period of 3.2ms during which a buffer will be needed to store packets 101 to 103. However, a constant sending rate is maintained indicated by the parallel nature of




Figure 5.7: Test 3 - 1ms (left) and close-up of handover point (right)

the two separate lines. The Voice over IP transmission would need a buffer of the order of 4ms, but quality will be maintained as jitter is not affected by these out-of-order packets.




Figure 5.8: Test 4 - $500\mu s$ (left) and Test 5 - $200\mu s$ (right)

In Figure 5.8, Test 4 shows a more severe case of out-of-order packets and a different sending rate for the packets. The first 100 packets are transmitted at a lower rate over the Wireless LAN than is required to support the steady stream from the Corresponding Node. This shows that the Wireless LAN interface is not capable of supporting this high packet flow rate. The sending rate is one packet per $500\mu s$, but the average receiving rate

at the Mobile Node is $1071 \mu s$. This would result in delays and breakup in the Voice over IP stream. Therefore the Wireless LAN interface is not suitable for this sending rate. Once the handover occurs at packet 100, the higher data rate of the Ethernet interface means that packet 101 is received after packet 59. This results in an overlap of 40.5ms where Ethernet packets are being received while the Wireless LAN is still transmitting packets. Once the handover has occurred, the Ethernet interface is capable of supporting an average sending rate of $628 \mu s$, still above the $500 \mu s$ sending rate from the Corresponding Node, but significantly better than the Wireless LAN interface.

Test 5 shows an even greater difference in the Wireless LAN and Ethernet interfaces. Neither interface can support the sending rate of $200 \mu s$ from the Corresponding Node, although the Ethernet interface is capable of an average sending rate of $265 \mu s$ while the Wireless LAN only manages $1057 \mu s$. Out-of-order packets occur from packet 28 onwards and the overlap extends beyond the limits of the 200 total packets, meaning packets are being received from the Wireless LAN after all 100 packets have been received by the Ethernet interface. Test 5 together with Test 4 therefore illustrate that the Wireless LAN interface should not be used for a sending rate of 500 or $200 \mu s$, but the Ethernet interface is capable of supporting these sending rates. The importance of having two different interfaces in a mobile device is again confirmed, since in these two cases, only the Ethernet interface can support the test parameters.

Discussion

In both Test A and B, a seamless handover was achieved, with only Test B showing out-of-order packets under packet transmit times of less than 1ms. It is important to observe that the Wireless LAN interface is not capable of supporting packet transmit rates of less than 1ms due to the physical characteristics of the technology. This will result in the Voice over IP streams being affected, since packets will be delayed through the Wireless LAN interface. Due to out-of-order packets, buffering is needed store later packets arriving early. However, the quality of the Voice over IP stream has already been affected by the limits of the Wireless LAN interface, so buffering does not help in this regard.

When the inter-packet transmit time was high enough, in Test 1 and 2, the Wireless LAN and Ethernet interfaces could both support the Voice over IP stream, with no out-of-order packets observed in either Test A or B. This shows that under the right conditions, these do different access technology can work seamlessly together. Once the limits of one interface

are reached, the overall quality of their combined performance is affected, as shown in these tests. This again emphasises the need to be careful in selecting which interfaces should be used for different types of traffic in a multiple interfaced device.

5.1.2 Dual Interface Aggregation

Dual interface aggregation is the practice of channelling the single IP packet flow received by the Router over both outgoing interfaces connected to the mobile node. The Corresponding Node was configured to generate a stream of UDP packets with the parameters outlined in Table 5.2 and transmit these packets to the Router's incoming interface.

Table 5.2: UDP Packet Parameters - Dual Interface Aggregation

Test Run	Inter-packet Transmit Time	Packet size (bytes)	Effective Data Rate	Effective # of VoIP Streams
1	20ms	172	8.6 KB/sec	1
2	2ms	172	86 KB/sec	10
3	1ms	172	172 KB/sec	20
4	500 μ s	172	344 KB/sec	40
5	200 μ s	172	860 KB/sec	100

The Click Modular Router was configured with the script detailed in Section 5.1.2 and Figure 4.5. This script will perform the routing necessary to transmit packets over both outgoing interfaces of the Router. The incoming stream of packets is held in a queue. When either the wireless LAN interface or wired Ethernet interface is ready to transmit a packet, the Click *element* RoundRobinSwitch() will send a packet out of that interface to the Mobile Node.




Figure 5.9: Packet flow diagram - Dual Interface Aggregation Scenario

Since the two outgoing interfaces have different characteristics such as transmission and physical layer delays, packets will not necessarily be transmitted one at a time to each interface alternately in a round-robin fashion, although in general this is the case. The five test runs were recorded by the Mobile Node and produced the following figures showing elapsed time vs. packet number. This representation is useful, since it will show out-of-order packets and even more importantly, any synchronisation problems between the two different interfaces.




Figure 5.10: Test 1 - 20ms (left) Test 2 - 2ms (right)




Figure 5.11: Test 3 - 1ms (left) Test 4 - 500 μ s (right)

Figure 5.10 shows Test 1 and Test 2. In Test 1, no out-of-order packets were observed. The Ethernet and Wireless LAN interfaces are synchronised and both able to maintain the sending rate from the Corresponding Node. In Test 2, only a few out-of-order packets are




Figure 5.12: Test 5 - 200 μs

received but a buffer of one packet would be sufficient to correct this.

In Figure 5.11 both Test 3 and 4 show more disruption to the packet flow. Test 3 resulted in 21 out-of-order packets out of the 500 transmitted, but again only a one packet buffer would be needed to correct this. Average jitter was only 0.7 ms with a maximum of 4.8ms, well below the 40ms limit for Voice over IP. Test 4 clearly shows that the Wireless LAN interface is starting to have significant difficulty in keeping synchronised to the Ethernet interface, as the inter-packet sending time from the Corresponding Node is reaching the Wireless LAN's limits. Out of 500 packets, 130 were out-of-order and now a buffer of 8.5ms (17 packets) will be needed to support a maximum jitter value of 8.1ms and average jitter of 1.4ms. However, the Wireless LAN is able to keep up with transmission rate of the Ethernet interface, so the two interfaces still work together with only an increase in jitter.

Test 5 in Figure 5.12 shows the synchronisation problem clearly. The two divergent lines illustrate that the Wireless LAN is unable to maintain a 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 and so under these test parameters, the Wireless LAN cannot be used for dual interface aggregation with the 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 has also been proven that using two interfaces together will not necessarily result in increased transmission rate. In theory, more packets are being sent per second. However, since these packets are not synchronised, the quality of the stream is decreasing. Therefore the possible advantage of the

increased bandwidth of the two interfaces working together is negated by a deteriorating stream quality at the Mobile Node.

Dual Interface Aggregation is shown to work when both interfaces are operating within their physical limits. Tests 1, 2 and 3 show near constant sending rates which indicates that the packets are arriving in time and maintaining the sending rate of the Corresponding Node. Tests 4 shows slight disruption to the constant sending rate as the limits of the Wireless LAN are being reached. Test 5 is a good illustration of potential for dual interface aggregation to fail to achieve the goal of increasing bandwidth and maintaining a quality stream of packets. Despite this, dual interface aggregation shows good performance when interface limits are not exceeded.

5.1.3 Lossy Link

The lossy link scenario aims to test whether having two interfaces to stream duplicated packets would reduce overall packet loss once the two stream are recombined at the Mobile Node. For this scenario, both interfaces were scripted to have the same percentage packet loss for each of the four test runs, with test parameters shown in Table 5.3. Figure 5.13 shows the test scenario. The Corresponding Node will stream 1000 packets to the Router. The router will then duplicate this packet flow and transmit each copy over each of the two interfaces.

Table 5.3: UDP Packet parameters - Lossy Link Scenario

Test Run	Packet Loss	Packet Size	Inter-packet transmit time
1	5%	172	5ms
2	10%	172	5ms
3	15%	172	5ms
4	20%	172	5ms

It is expected that packet loss will be greatly reduced once these two separate streams are recombined. Statistically it is possible to work out the theoretical packet loss of the combined streams at the Mobile Node. If each interface has 5% packet loss, then the combined packet loss will be $0.05 \times 0.05 = 0.0025$ which is 0.25%. The theoretical packet losses for each of the test runs is shown in Table 5.4.

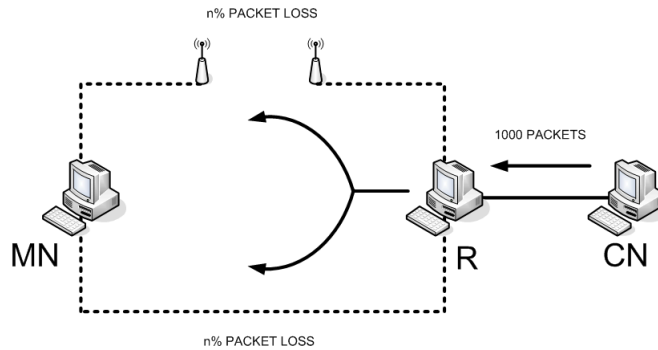


Figure 5.13: Packet flow diagram - Lossy Link Scenario

Table 5.4: Theoretical Combined Packet Loss

Test Run	Link Packet Loss	Combined Packet Loss
1	5%	0.25%
2	10%	1.00%
3	15%	2.25%
4	20%	4.00%

Since Voice over IP can tolerate up to 1% packet loss and still achieve good voice quality, it is theoretically possible to combine two link with very poor packet losses of 10% and result in a combined packet loss of 1%. This means that two unusable interfaces which cannot carry quality voice traffic are combined to produce a link that meets Voice over IP specifications. To confirm these theoretical statistics, the four test runs were executed on the testbed and the results for each test run are shown in the following figures.

In Figure 5.14, Test 1 shows that of the 2000 packets transmitted (1000 over each interface), 93 were dropped. But when these two duplicate streams were combined, only 1 packet of the 1000 transmitted by the Corresponding Node was missing. This gives an overall packet loss of 0.1%. This confirms the theoretical combined packet loss of 0.25% shown in Table 5.4. For Test 2, the results show 217 packets dropped with only 25 packets missing from the combined stream. This is a combined packet loss of 2.5%, which is much higher than the 1% theoretical packet loss. This shows that a real-world scenario can deviate from statistics over small sample sets like 1000 packets. It is expected that with a figure of 10 000 packets, the combined packet loss would tend towards 1%. However, it is important to see that theoretical figures can be exceeded on a real testbed, since a real world scheme




Figure 5.14: Test 1 - 5% packet loss (left) Test 2 - 10% packet loss (right)

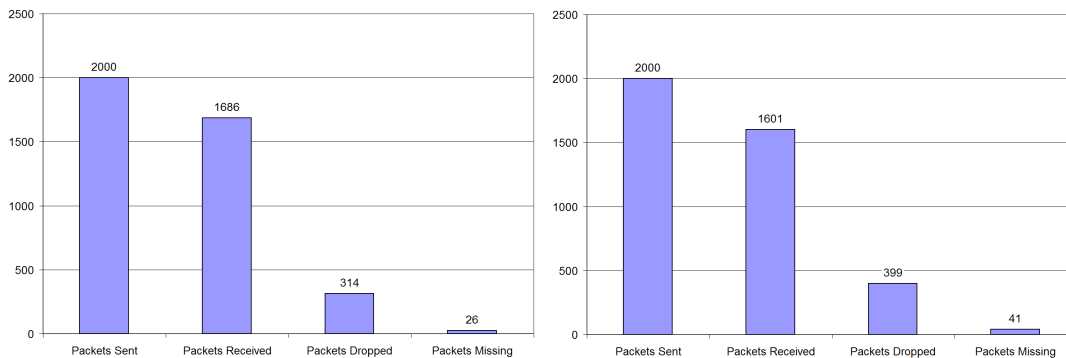


Figure 5.15: Test 3 - 15% packet loss (left) Test 4 - 20% packet loss (right)

would need to account for the worst case packet loss. Since 2.5% packet loss is well above the Voice over IP limit of 1%, these two interfaces with 10% packet loss each would be unsuitable for a Voice over IP stream.

Figure 5.15 shows Test 3 and Test 4. Test 3 showed 26 missing packets, with 314 packets dropped. The combined packet loss is therefore 2.6%, which is close to the theoretical 2.25% from Table 5.4. Again, 2.6% is above the Voice over IP limit of 1%, but still shows that two completely unusable links with 15% packet loss each can be combined to reduce packet loss to only 2.6% which might suit a less demanding application than Voice over IP. For Test 4, 41 packets were missing from the combined streams. This amount to 4.1% packet loss compared to the theoretical amount of 4.0%.

Figure 5.16 compares the total number of packets that were dropped over both links to the number of packets that were missing once the two streams were recombined at the Mobile

Node. In all four cases, significantly less packets were missing.




Figure 5.16: Packet losses for the four test runs

In order to illustrate the percentage packet loss for each test run compared to the Voice over IP limit of 1% packet loss, Figure 5.17 was created. The Voice over IP limit of 1% is indicated by the dotted line. Of the four test runs, only Test 1 showed packet loss below 1%, but it was expected from statistics that Test run 2 would also meet the 1% requirement. This was not the case in during the test, but is due to a small sample of only 1000 packets. Nevertheless, it is evident that in all four test runs, packet loss was reduced by at least 400% when the duplicated streams were combined at the Mobile Node.




Figure 5.17: Link and Combined Packet Loss

This scenario shows that lossy links can be significantly improved if two are combined. Two links that have high packet loss of 5% each can be combined to support a Voice over IP stream with less than 1% packet loss. This proves that multiple interfaces can be used

to take two unusable links and combine them into one link which meets the requirements of the application. For this reason, a multiple interfaced device would be extremely useful in a situation where link quality is poor or coverage and network congestion are unsuitable for a sensitive application like Voice over IP. A single interfaced device would not be suitable.

5.1.4 Bandwidth on Demand

In order to simulate a Bandwidth on Demand scenario, the Router was scripted according to Section 4.7.4 and Figure 4.7. This scenario was split into two tests.

Test 1 - One interface active

For this test, the Corresponding Node will stream packets to the Router with an inter-transmit time of 20ms. This represents a single Voice over IP stream of 64 kilobits/sec. The Router is scripted to limit both the Ethernet and Wireless LAN links to 384 kilobits/sec. To simulate a bandwidth-on-demand system, any incoming packet flow from the Corresponding Node that is below 384 kilobits/sec will be routed over the Ethernet interface only. If the incoming packet flow should exceed 384 kilobits/sec, then the Wireless LAN interface will carry the excess packets. Since the incoming packet rate is only 64 kilobits/sec, this stream is easily accommodated by the Ethernet interface. Refer to Figure 5.18. The wireless LAN interface is not activated and is greyed out since the 384 kilobit/sec Ethernet link is capable of carrying the entire 64 kilobits/sec packet stream. The wireless LAN interface is therefore still active and ready to transmit packets should the incoming packet stream at the Router exceed 384 kilobits/sec.




Figure 5.18: Packet flow diagram - Bandwidth on Demand Scenario Test 1

Figure 5.19 shows the packets received at the Mobile Node. Out of 200 packets transmitted, none were out of order. The 64 kilobit/sec stream is easily accommodated by the Ethernet link of 384 kilobits/sec, so packet jitter is very low. This is evident from the straight line nature of the graph.




Figure 5.19: Ethernet link with 64 kilobit/sec stream

The test scenario was modified to reverse the roles of the two interfaces. The 64 kilobit/sec stream was first routed over the Wireless LAN interface, with the Ethernet interface acting as the on-demand link for packets in excess of 384 kilobits/sec. Figure 5.20 shows packets received when the test scenario was modified so that the 64kbps stream would first be routed over the wireless LAN link. Again, the packet stream is well within the transfer limits of the Wireless LAN, so packet jitter is very low. The graph resembles a straight line characteristic of low packet jitter and regular packet arrival.

Test 2 - Both interfaces active

In Test 2, the Corresponding Node increases its packet sending rate ten times, as inter-packet transmit time is now 2ms. The data rate is now 640 kilobits/sec. The Router is scripted to send up to 384 kilobits/sec of packets over the Ethernet interface. Since the packet stream exceeds this limit, the wireless LAN interface (also limited to 384 kilobits/sec) will carry the excess packets amounting to 256 kilobits/sec. Refer to Figure 5.21.




Figure 5.20: Wireless LAN link with 64 kilobit/sec stream




Figure 5.21: Packet flow Diagram - Bandwidth on Demand Test 2

The triggering of the wireless LAN interface is automatic, as is desirable in a Bandwidth-on-Demand system. Figure 5.22 shows the two separate packet flows of 384 and 256 kilobits/sec combined at the Mobile Node. It is clear that packet jitter is low and packets are arriving regularly at the Mobile Node. Only 3 packets out of 500 were out-of-order and a single packet buffer would be sufficient to correct this. It was observed that packets were transmitted alternately between the Wireless LAN and Ethernet interface, with only a few instances of two or more packets being transmitted over the same interface in succession. This shows that the Bandwidth-on-Demand configuration was efficiently scheduling packets over both interfaces.

As in Test 1, the test scenario was reconfigured to reverse the roles of the two interfaces. The Wireless LAN interface was made the primary interface, with the Ethernet interface




Figure 5.22: Ethernet and Wireless LAN link with 640 kilobit/sec stream

acting as the Bandwidth-on-Demand standby. Figure 5.23 shows the two separate packets flows combined at the Mobile Node. As expected, this graph is similar to the previous one, with low jitter and only 3 out-of-order packets out of 500. Since both interfaces are within their transfer limits, it is expected that whether the Wireless LAN or the Ethernet link is the standby should make little difference to the performance of the scenario.




Figure 5.23: Wireless LAN and Ethernet link with 640 kilobits/sec

5.2 Discussion of Scenarios

In this Chapter, four test scenarios were implemented on the evaluation testbed to investigate the functionality possible when a mobile device is equipped with multiple interfaces

(in this evaluation platform, two interfaces were used). The tests performed during each scenario show how a packet flow behaves when being altered and routed according to the script implemented on the Router. In all four scenarios, having two interfaces was beneficial.

In the first scenario, it was shown that redirecting an incoming packet flow to a new interface adds no significant jitter or packet loss to the packet flow if the transfer limits of either interface are not exceeded. It was interesting to observe the increased packet jitter caused by the wireless LAN interface due to its physical layer characteristics. The ability to redirect packet flows during mid-stream will benefit the Mobile Device user should a handover be necessary, specifically during an ongoing communication session. In this scenario, it was shown that this handover is seamless and caused no disruption to the packet flow if the Wireless LAN was not over-stressed, as it has higher packet jitter than the Ethernet interface. This is a highly desirable feature of a handover scheme. It should be remembered that both interfaces were activated and had IP Addresses assigned before the test took place, so it is incorrect to compare the results of the packet flow handover to those of handover schemes which aim to minimise the disruption to packet flows during a break-before-make vertical handover. It is evident that the ability to perform a Mid-Flow Interface Transfer while data is being sent to the mobile device will benefit the user, since their communication session is not interrupted at any stage of the handover process, as long as the new interface is capable of carrying the full data stream. It was also interesting to observe that combining two different interfaces can be problematic if the performance of one interface differs from the other. In this case, the Wireless LAN interface was the weaker link and did not support packet inter-transmit times of less than 1ms without a significant number of out-of-order packets.

In the second scenario, Dual Interface Aggregation was investigated. In order to split a single incoming packet flow over two outgoing interfaces, a suitable scheduling *element* was required to perform the task of correctly distributing each packet over one of the two interfaces according to which interface was ready to transmit. The RoundRobinSwitch() *element* proved ideal for this task. It was observed that most packets were routed in an alternating pattern between the Wireless LAN and Ethernet interface. This is an efficient scheme for routing packets when both interfaces are below their transfer limits. When the Wireless LAN interface could no longer support the packet sending rate from the Corresponding Node, out-of-order packets were evident. The number of out-of-order packets reached an unacceptable level below $500\mu\text{s}$ inter-packet transmit time. Jitter was

of the order of 10ms and a buffer of almost 10ms would be needed to store these out-of-order packets. This is approaching the Voice over IP limit of 40ms total jitter. However, these values are still within acceptable limits. It was in the final test run using $200\mu\text{s}$ inter-packet transmit time that the critical problem with dual interface aggregation was observed. The Wireless LAN was unable to maintain synchronisation with the Ethernet interface as packet were arriving too quickly. This resulted in an ever increasing jitter, which no buffering would be able to correct. This test proves that dual interface aggregation is not a trivial matter of routing packets over two separate interfaces. Unless both interfaces are well within their transfer limit and are evenly matched in capability, synchronisation will be a problem. Although all packets do eventually arrive at the Mobile Node, they are so out of sync that the strict requirements of Voice over IP are not met. For real time applications then, dual interface aggregation is successful when both interfaces are able to deliver packets timeously.

In the third scenario, it was shown that two equally lossy links with identical packet loss can be combined to form a single, channel with significantly reduced packet loss. The first test run showed involved two links with 5% packet loss. Neither of these interfaces can support a quality Voice over IP session, which requires 1% packet loss maximum. When these two interfaces both carried duplicate packet streams, the combined streams arriving at the Mobile Node experienced only 0.1% packet loss. This proves that two unusable links can be combined into one that will support a critical application like Voice over IP. Reduced packet loss was observed in all four test runs with values close to theoretical calculations. Particularly in wireless networks, packet loss is a common problem. Combining two lossy wireless links can provide a clear channel for the data stream that would not be supported on the individual links. Here it is clear that a multiple-interfaced device would prove to be beneficial in a situation where a single-interfaced device could not provide quality communication.

In the fourth and final scenario, it was shown that a Bandwidth-on-Demand scheme would operate as automatically - a link limited to 384Kbps would carry a 64Kbps data stream singlehandedly while the second interface remained dormant. When the incoming data stream exceeded 384Kbps, the second link was activated and packets were routed alternately over both interfaces at 384Kbps and 256Kbps respectively (for a total of 640Kbps for both links, matching the 640Kbps transfer rate of the incoming packet stream). The Wireless LAN and Ethernet links were both within their transfer limits during this scenario, so either interface proved to be an effective standby. No packets were lost during the transfer

and both interfaces carried the data stream reliably to the Mobile Node, with only a few out-of-order packets observed during the 640 kilobit/sec packet stream. A Bandwidth-on-Demand scheme such as this would be useful in a mobile device should the user activate a multimedia session that exceeded the current transfer rate of the single interface being used. Automatic activation of a second interface would be optimum and provide a seamless multimedia session for the user with no intervention required on their part besides setting this scheme as a preference on the mobile device beforehand. A single-interfaced device would not support this functionality and as soon as the application required a higher transfer rate than the single interface could support, all multimedia would be affected. With a multiple-interfaced device, the potential exists to automatically adjust the number of interfaces activated to support the total data rate of all running applications.

It is clear that a multiple-interfaced device provides useful functionality in all four scenarios as long as certain limits of either interface are not exceeded. The problem of out-of-order packets is the most critical to solve, since this results in adverse effects on the quality of the Voice over IP stream. The performance limits of each scenario were clearly shown and it is possible to have all four scenarios running on a practical multiple interface scheme, as long as the scheme is careful to account for jitter, out-of-order packets and synchronisation issues that can arise. Despite these problems, a multiple-interfaced device shows important and useful benefits over a single-interface device and therefore the complexity of developing a multiple interface management scheme can be justified.

Chapter 6

Conclusions

6.1 Conclusions

This study has investigated the functionality possible with a multiple-interfaced device in an IP network. Specifically, this work has emulated four possible scenarios that show the type of performance possible with a multiple-interfaced device.

In the literature, it was found that multiple interface management schemes tend to have a key component - the central packet redirection point. This central point is responsible for routing incoming packet flows to the correct interfaces of the mobile device. However, none of these scheme proved that the functionality they could provide would in fact work in a practical implementation. Thus the goal of this work is to emulate a generic scheme based on the components evident in the literature. The generic scheme consists of a Corresponding Node for packet generation, a central Router and an emulated Mobile Node equipped with two different interfaces.

An evaluation framework was built using these three components and emulated with three computers using Ethernet and 802.11 Wireless LAN interfaces. This test-bed performs packet flow redirection using a software router running on the central router, configured with scripts for each test scenario. Useful results were obtained that show good performance for each test scenario when two interfaces are used. However, problems were observed with each test scenario that prove a multiple interface scheme is not a trivial system to design. Based on the results and findings of previous chapters, the following conclusions are drawn:

- Single interfaced devices are no longer sufficient to enable ubiquitous access to emerg-

ing next generation networks. The networks will converge numerous different access technologies, so mobile device need more than one interface built-in to connect to these diverse access technologies.

- Mobile IPv6 can enable mobility of devices in fixed networks, but experiences difficulty in supporting a device with multiple interfaces. Since Mobile IPv6 is a network layer mobility technology, a control scheme is necessary to manage the various interfaces of the mobile device.
- Multiple interface management schemes tends to favour Mobile IPv6 for providing mobility, but add a control scheme for managing multiple interfaces on the mobile device. This is achieved by modifying or extending Mobile IPv6 with added functionality, or designing a new signalling protocol.
- A simple but effective emulation of a multiple interface scheme is possible using the Click Modular Router software. This software allows packets to be modified and routed efficiently to either of the outgoing interfaces of the Router. Click was able to route packets with negligible processing delay and set appropriate checksums and packet headers for correct packet modification.
- It was shown that UDP is a good choice of protocol for routing real-time traffic over multiple interfaces. Without sequence numbers to worry about, it is a simple process to modify UDP packet headers to travel over different interfaces.
- A mid-flow handover from one interface to the other is possible, with no packet loss, but both interfaces must be capable of supporting the full sending rate of the incoming packet stream, otherwise jitter will increase and data rate will decrease. In this regard, it is shown that the Wireless LAN interface had higher jitter and lower data rate capability than the Ethernet interface.
- Dual Interface Aggregation is a possible method of increasing the total data rate available to a mobile device. However, out-of-order packets cause delays that affect the quality of the packet stream. This will only occur when interface transfer limits are exceeded, so it is important to only use dual interface aggregation if both interfaces are capable of maintaining a high enough transfer rate to split the incoming stream in two.

- Two lossy links with unacceptably high packet loss for quality Voice over IP transmission can be combined into a single link that will meet the strict requirements of less than 1% total packet loss if both link have less than 5% packet loss each.
- A Bandwidth-on-Demand scheme will prove effective if both interfaces are carrying traffic within their transfer limits. A single interfaced device will not be capable of dynamically activating an additional interface as the application requires more bandwidth.

6.2 Recommendations and Future Work

While conducting this work, a number of issues surrounding multiple interfaced devices and their functionality were encountered. These issues are discussed here for future research on these topics:

- Multiple-interfaced devices are increasingly being released, but networks lack support for using these interfaces in collaboration. Further research is needed into designing schemes to manage these multiple interfaces. A suitable mobility scheme is needed, along with a protocol for controlling the changing connections to access networks. While Mobile IPv6 can provide certain mobility functionality, this protocol needs assistance in managing multiple interfaces. For this reason, research into a suitable control protocol is needed.
- Mobile devices will use applications that will use connection-oriented protocols such as TCP. The effects of multiple interface functionality on a TCP connection should be studied to determine whether this protocol can function in this environment. Although the focus of this study is on real-time applications using real-time protocols such as RTP and UDP, non-real time applications are also an important feature of mobile devices. For this reason, scenarios such as the four presented in this work should be developed for non-real time protocols travelling over multiple interfaces.
- Although Voice over IP traffic was simulated by a UDP stream of packets on the test-bed in this work, actual Voice over IP calls were not tested for qualitative evaluation. A real-time voice application running on the Corresponding Node and Mobile Node would provide further insight into the quality of the call during interface handovers and interface switching.

- A multiple-interfaced device with more than two interfaces should be studied, to see whether having more than two interfaces provides any significant benefits that justifies the added complexity of this configuration.
- In this study, only IP-based Ethernet interfaces were investigated on the test-bed. It would be worthwhile to test the interaction of multiple interfaces using technologies like a 3G interface or a WiMAX interface. The real-world delay and jitter conditions of these interfaces would provide additional insight into the challenges a mobile node will experience in managing flows over two or more access technologies of this type.
- The Router could be extended to support more than one multiple-interfaced mobile node. The performance of the Router under this additional load could affect the success of the multiple interface functionality under test. This configuration could give further insight into the performance of a multiple interface scheme supporting multiple mobile devices.

Bibliography

- [1] N. Fikouras and K. K. et al, “Multiple Access Interface Management and Flow Mobility,” *University of Bremen, Germany*, 2004.
- [2] D. Johnson, C. Perkins, and J. Arkko, “Mobility Support in IPv6,” Internet - <http://www.rfc-editor.org/rfc/rfc3775>, June 2004.
- [3] 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>, April 2006.
- [4] “Internet Protocol,” Internet - <http://www.ietf.org/rfc/rfc0791.txt>, September 1981.
- [5] Various, “OSI Model,” Internet - http://en.wikipedia.org/wiki/OSI_model, 2007.
- [6] A. Dutta, F. Vakil, J. Chen, and M. Tauil, “Application Layer Mobility Management Scheme,” *Computer Science Department, Columbia University, New York*.
- [7] I. Okajima, M. Inoue, K. Omae, H. Takahashi., and N. Umeda, “IP-based Mobility Management Technology,” *NTT Technical Review*.
- [8] 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.
- [9] J. Sun and J. Sauvola, “Mobile IP Applicability: When Do We Really Need It?” *International Conference on Parallel Processing Workshops*, 2004.
- [10] “Advanced Mobile Phone System,” Internet - <http://www.leapforum.org/published/internetnetworkMobility/split/node37.html>.

- [11] S. Hui and K. Yeung, "Challenges in the Migration to 4G Mobile Systems," *IEEE Communications Magazine*, December 2003.
- [12] Cisco, "A Clear Understanding: What makes Voice over IP Work?" Internet - http://www.cisco.com/warp/public/779/smbiz/community/voip_basics.html.
- [13] Various, "Click Modular Router Wiki," Internet - <http://www.read.cs.ucla.edu/click/>.
- [14] E. Kohler, "The Click Modular Router," Ph.D. dissertation, MIT, February 2001.
- [15] A. Hasson, "dgram-gen.c," 2004.