HELSINKI UNIVERSITY OF TECHNOLOGY
Telecommunications Software and Multimedia Laboratory

Joni Bask

Heterogeneous Handoff Algorithm for Mobile IPv4 without Foreign Agents

Master's Thesis: May 5, 2003

Supervisor — Professor Antti Ylä-Jääski

Instructor — Panu Pietikäinen, MSc

HELSINKI UNIVERSITY OF TECHNOLOGY

ABSTRACT OF MASTER'S THESIS

Author and title:
Joni Bask
Heterogeneous Handoff Algorithm for Mobile IPv4 without Foreign Agents

Date: May 5, 2003 **Number of pages:** 9 + 96

Department: Department of Computer Science and Engineering

Chair: Tik-110

Supervisor:
Professor Antti Ylä-Jääski

Instructor:
Panu Pietikäinen, MSc

Mobile IP is a protocol that allows a host in the Internet to change its point of attachment to the fixed network without losing connection to its peers. Handoff is a mechanism by which the point of attachment is changed from one access network to another. Handoffs cause interrupts in the data flow until the mobile node has established its presence in the new network.

The number of mobile nodes is increasing, and there is a trend towards more time-sensitive applications. This calls for efficient handoff management to minimize the undesirable effects of handoffs. At the same time, new wireless transfer media have become available.

This thesis presents a heterogeneous handoff algorithm that achieves maintenance of undisturbed connectivity and makes the best use of available network resources. The solution uses link layer information and statistical methods to improve handoff target decisions, to make the handoffs faster, and to prioritize available connections to reach the improved overall quality of service.

As part of this thesis, I identified several problems involved in simultaneous Mobile IP registrations with timestamp- and nonce-based replay protection. Based on the findings, I developed a registration scheme that allows efficient registering through multiple interfaces having dynamic priority order.

Keywords: Mobile IP, heterogeneous handoff, fast handoff, seamless handoff, QoS, velocity adaptation, ping-pong detection, IEEE-802.11, IEEE-802.3, GPRS, registration management, implementational feasibility

TEKNILLINEN KORKEAKOULU DIPLOMITYÖN TIIVISTELMÄ

Tekijä ja työn nimi:
Joni Bask
Heterogeeninen yhteydensiirtoalgoritmi Mobile IPv4:ään ilman vierasagentteja

Päivämäärä: 5. toukokuuta, 2003 **Sivumäärä:** 9 + 96

Osasto:
Tietotekniikan osasto

Professuuri:
Tik-110

Työn valvoja:
Professori Antti Ylä-Jääski

Työn ohjaaja:
Diplomi-insinööri Panu Pietikäinen

Mobile IP on protokolla, joka mahdollistaa päätelaitteiden liikkumisen Internetissä istuntojen katkeamatta. Yhteydensiirto on mekanismi, joka mahdollistaa siirtymisen kahden eri aliverkon välillä. Samalla kuitenkin aiheutuu katkos tiedonsiirtoon.

Liikkuvien IP-päätelaitteiden määrä on lisääntymässä. Samalla on huomattavissa trendi kohti reaaliaikaisia ja aikakriittisiä sovelluksia. Tämä vaatii tehokasta liikkuvuuden hallintaa, yhteydensiirron aiheuttamien haittojen minimoimiseksi.

Tämä työ esittelee heterogeenisen yhteydensiirtoalgoritmin, joka pyrkii turvaamaan häiriöttömät yhteydet, painottaen eri tiedonsiirtoteknologioiden optimaalista käyttöä. Ratkaisussa käytetään linkkikerroksen informaatiota ja tilastollisia menetelmiä parantamaan kohdeyhteyden valintaa, nopeuttamaan yhteydensiirtoprosessia ja priorisoimaan tarjolla olevia yhteyksiä.

Osana tutkimustani yksilöin useita Mobile IP rekisteröintiin liittyviä ongelmia, kun rekisteröinti on mahdollista tehdä usean eri aliverkon kautta. Ongelmat aiheutuvat paljolti aikaleima- ja nonce-pohjaisen uudelleenlähetyssuojauksen rajoitteista. Kehitin havaintojeni pohjalta menetelmän, joka nopeuttaa rekisteröintiä sen tapahtuessa usean eri priorisoidun yhteyden kautta.

Avainsanat: Mobile IP, heterogeeninen yhteydensiirto, nopea yhteydensiirto, saumaton yhteydensiirto, palvelun laatu, nopeuteen mukautuminen, ping-pong:n havaitseminen, IEEE-802.11, IEEE-802.3, GPRS, rekisteröinnin hallinta, toteutettavuus

Preface

I want to thank my instructor Panu Pietikäinen for an opportunity to carry out this work at Netseal and for his valuable insights about wireless technologies. Thanks to Professor Antti Ylä-Jääski for his academic guidance, and for taking the time and the effort to supervise this thesis.

Special thanks belong to Sami Vaarala who helped me in numerous ways during the work. I want to express my appreciation for enlightening discussions and for his contributions in formulating the research problem. I thank him for excellent advises in many areas involved in IP mobility management. His comments on the work have definitely been valuable.

I want to thank my friends and colleagues at Netseal mobility team, Nuutti Kotivuori, Ilkka Pietikäinen, Kristian Slavov, and Mikko Varso for pleasant and inspiring working atmosphere. I thank Ilkka Pietikäinen for his consultation and throughout knowledge of Mobile IP.

Antti Nuopponen provided me with many useful LaTeX tips that eased the writing process. Elizabeth Heap-Talvela helped me with matters of the English language.

Finally, I want to thank my family and friends, particularly Anne Sinisalo, for support while doing this work.

Otaniemi, May 2003

Joni Bask

Contents

List of Figures

List of Tables

Glossary

access point (AP) Any entity that has station functionality and provides wireless access to the fixed network.

access router (AR) An IP router residing in an access network that is connected to one or more APs. An AR offers IP connectivity for mobile nodes.

agent discovery A process through which a mobile node learns the mobility agents residing in the visited network.

association An IEEE 802.11 service that enables the mapping of a wireless station to the fixed network via an access point.

base station Telecommunications radio equipment located at one fixed location, and that is used for serving one or several cells.

basic service area The conceptual area within which members of a basic service set (BSS) may communicate.

basic service set A set of IEEE-802.11-compliant stations that operate as a fully connected wireless network.

care-of address (CoA) The termination point of a IP-IP tunnel toward a mobile node.

co-located care-of address (CCoA) A care-of address that is bound to one of the network interfaces of the mobile node. Compare to foreign agent care of address.

correspondent node (CN) A peer with which a mobile node is communicating. A correspondent node may be either mobile or stationary host.

dynamic host configuration protocol (DHCP) A protocol that is used for dynamic provision of IP addresses for Internet hosts.

fast handoff A handoff that is completed faster than a simple Mobile IP handoff. Fast handoff is not loss-free. Packets may be dropped during the process.

foreign agent (FA) A router on a mobile node's visited network which provides routing services to the mobile node while registered. The foreign agent detunnels and delivers datagrams to the mobile node that were tunneled by the mobile node's home agent. For datagrams sent by a mobile node, the foreign agent may serve as a default router for the registered mobile nodes.

foreign agent care-of address An IP address of a foreign agent that serves tunnel end-point for mobile nodes.

foreign network Any network other than the mobile node's home network.

gratuitous ARP An ARP packet sent by a host in order to spontaneously cause other hosts to update the entry in their ARP cache.

handoff A process by which an Internet host changes its point of attachment from one subnet to another.

handoff latency The duration of interruption to data flow from and to the mobile node caused by a handoff.

handoff management A process that initiates, executes, and controls handoffs of a mobile node to ensure the best possible connectivity.

handoff policy A set of rules that identify the next point of attachment to be used if a handoff is needed.

handoff target The destination point of attachment, to which connection is transferred during the handoff.

heterogeneous handoff See inter-media handoff.

home address An IP address that is assigned for an extended period of time to a mobile node. The home address remains unchanged regardless of where the mobile node is attached to the Internet.

home agent (HA) A router on a mobile node's home network which tunnels datagrams for delivery to the mobile node when it is away from home, and maintains current location information of the mobile node.

home network A network, possibly virtual, having a network prefix matching that of a mobile node's home address.

ingress filtering Routers exercise ingress filtering to prevent source address spoofing of IP packets. This is done by matching the source address to network prefixes of known networks behind the link. If the match is not found, the packet is dropped.

inter-media handoff A handoff, during which connection is transferred over one type of link layer to another.

link A facility or medium over which nodes can communicate at the link layer. A link underlies the network layer.

link layer (L2) Handoff Change of MN's link layer connection from one location to another. No change in off-subnet routing reachability information is required if the L2 handoff occurs within a subnet.

link layer address The address used to identify an endpoint of some communication over a physical link. Typically, the Link- Layer address is an interface's Media Access Control (MAC) address.

mobile node (MN) A host or router that changes its point of attachment from one network or subnetwork to another. A mobile node may change its location without changing its IP address; it may continue to communicate with other Internet nodes at any location using its home address.

mobility agent Either a home agent or a foreign agent.

mobile controlled handoff A handoff, whereby the mobile node is responsible for determining its new point of attachment and carries out the necessary protocols for making the determination as well as establishing the presence at the point of attachment.

mobility binding The association of a home address with a care-of address, along with the remaining lifetime of that association.

mobility security association A collection of security contexts, between a pair of nodes, which may be applied to Mobile IP protocol messages exchanged between them. Each context indicates an authentication algorithm and mode, a secret (a shared key, or appropriate public/private key pair), and a style of replay protection in use.

move detection A method that mobile node uses to notice the change of interfaced subnet.

network controlled handoff A handoff, whereby some entity in the serving network directs the establishment of a new link between the mobile node at some point of attachment determined by the network elements.

network layer (L3) Handoff Change of a mobile node's routable address from one subnet to another. An L3 handoff results in a change in the mobile node's routing reachability, that will require renewal of the mobility binding.

ping-pong A situation in which the connection of the mobile node rapidly oscillates between two point of attachments, causing excessive number of unnecessary handoffs.

point of attachment Location of network side termination point of the mobile node's physical layer connection to the network.

registration management A procedure that is followed when forming and sending Mobile IP Registration Requests via multiple interfaces to fast establish mobility binding.

reverse tunneling Mobile IP routing method that tunnels packets from a mobile node to the corresponding node via the home agent. As source IP addresses of packets are topologically correct, ingress filtering is not a problem as with triangular routing.

seamless handoff A handoff that causes no packet loss nor handoff latency. A seamless handoff can not be noticed simply by examining the data stream.

Security Parameter Index (SPI) An index that identifies security context between pair of nodes. The SPI is chosen among the contexts available in the Mobility Security Association.

smooth handoff A handoff that causes interruption to the flow of packets, but zero packets are lost during the handoff. Usually smooth handoff includes buffering of packets during the interruption.

station Any device that contains an IEEE 802.11 medium access control (MAC) and physical layer (PHY) interface to the wireless medium (WM).

target network The network that is targeted in handoff. Network layer handoff means always change of visited subnet, but link layer handoff may occur within a subnet.

tunnel The path followed by a packet while it is encapsulated. Encapsulated packet is routed to a knowledgeable decapsulating agent, which decapsulates the datagram and then correctly delivers it to its ultimate destination.

velocity adaptation A method by which a mobile node adjusts handoff target decisions and schedules handoffs in accordance to its current velocity.

visited network A network other than a mobile node's home network, to which the mobile node is currently connected.

1 Introduction

The concept of mobility has been the subject of a lot of interest within the Internet community during the past decade. The telecommunications industry has revealed the benefits of being able to communicate anywhere and anytime. With true IP level mobility, nomadic Internet users would be able to access network resources and exchange data with each other in a similar manner. This need becomes every day more evident. The number of mobile hosts is increasing with the emergence of lightweight digital personal assistants and laptop computers capable of wireless communication.

Mobile IP is the most prominent way to introduce mobility to the Internet Protocol based communications. The idea is to make communicating hosts unaware of the underlying movement between the subnets. This allows applications to operate as if the hosts were stationary. A *mobile node* is a host that can change its point of attachment to the Internet from one subnet to another. To maintain ongoing communications with other hosts, the IP address of the mobile node must remain immutable. Otherwise, the higher layer protocols lose the end-to-end connection they need for exchanging packets. A permanent IP address called a *home address* is allocated from the home network of the mobile node. Packets destined for the mobile node are routed with the home address to the home network. The packets are intercepted by a *home agent*, a network entity residing at the home network. The home agent tunnels packets to the current location of the mobile node. Whenever the mobile node moves to a new subnet, it registers its new location to the home agent.

Moving from one subnet to another causes a handoff. During the handoff, a mobile node's point of attachment to the fixed network is transferred to another. This causes a disruption in the ongoing data flow. The overall delay can be separated into four components. First, the mobile node has to set up a link layer connection to

a new point of attachment. Then, the movement has to be detected in the network layer. This can be done by employing network layer mechanisms or by sharing information between the link and the network layer. After that, a temporary IP address called a *care-of address* must be acquired from the new subnet. Finally, the care-of address must be registered to the home agent in order for packets to be redirected to the mobile node. The connection is broken until a new tunnel is formed from the home agent to the care-of address.

If handoff is not handled properly, it causes user perceivable degradation of quality of service. Especially real-time and delay-sensitive applications suffer from the interrupted data flow. The handoff algorithm seeks to execute handoff as efficiently as possible to minimize packet drop and latency. When a mobile node possesses multiple network interfaces, the handoff algorithm must also choose the best interface to be used at a time. The right timing of the handoff is essential in wireless networks. Delayed handoff causes increased packet drop as signal quality decreases when the distance between the transmitters grows. In the transport layer, TCP protocol interprets the break caused by the handoff as network congestion. The congestion control mechanism delays packet transmission even more. Increased use of multi-media services calls for efficient handoff management.

Recent years have seen the introduction of several new technologies for wireless data transfer. Currently the assortment includes for instance GSM data, GPRS, UMTS, wireless LAN, and Bluetooth. Telecommunications networks have wide coverage areas, but modest bandwidth compared to the technologies of smaller cell sizes like IEEE-802.11. Inter-media, or heterogeneous handoffs are the way to make the best use of the currently available network resources. Wired connection can be used at the office, whereas wireless LAN is the best option for a mobile user. Cellular technologies fill the gaps in the wireless LAN coverage. A fast-moving mobile node benefits from the large radio cell sizes of the cellular technologies. Using wireless LAN, a high velocity mobile node may not occupy one cell long enough for the

handoff to complete at all. This results in frequent care-of address registration attempts that only increase network traffic. At the same time, zero user data is transferred. A large cell size diminishes the handoff frequency of a moving mobile node. Heterogeneous handoff algorithm collects data from each available transfer medium to make optimal inter- and intra-media handoffs.

Mobile IP specification assumes that in general, each access network comprises a network element called a *foreign agent*. Foreign agents detunnels and delivers packets from the home agent to the mobile node. They provide care-of addresses for the mobile nodes. Additionally, foreign agents can assist mobile nodes in handoff execution. In the absence of a foreign agent, a mobile node can perform detunneling by itself and acquire a care-of address with other means. At the time of writing this thesis, only an insignificant portion of networks comprise foreign agents. To operate widely, a mobile node has to cope without them. The lack of foreign agents has its implications for handoff management. Most of the current research to improve handoff management involves foreign agents. It can be questioned if foreign agents will ever be deployed on a larger scale. Benefits of the foreign agents may not justify the investments needed. It is more likely that foreign agents will be deployed only for isolated regions like hotspots. Finally, the commonly available hotspots must also serve clients that do not execute Mobile IP. Those clients will acquire care-of addresses without communicating with the foreign agents.

This thesis studies Mobile IP protocol to identify its implications for multiple network interface handoff management in networks that lack foreign agents. Based on the findings, a new heterogeneous handoff algorithm is presented. The algorithm does not assume any changes to intermediate network components or Mobile IP protocol. Handoffs are strictly mobile controlled and do not rely on foreign agent support of any kind.

2 Problem setting

2.1 Handoffs in Mobile IP

Mobile IP was designed to provide *macro mobility* for mobile nodes. A macro handoff occurs whenever there is a need to register a new care-of address (CoA) to the home agent (HA). Similarly, the term *micro mobility* is referred when the mobile node (MN) moves within a subnet or an administrate domain. Registration Requests are either unnecessary or terminated inside the administrative domain to update the location of the MN locally. Mobile IP is ill suited for frequent handoffs. The main source of the problem is the latency and packet loss that is induced by lengthy registration processes: Registration messages must traverse all the way to the home agent and back. Besides, the Mobile IP network layer movement detection mechanism is slow. This delays the initiation of the registration process even more. Nowadays, real-time requirements are rapidly infiltrating the IP world and the Mobile IP network layer handoff procedure is seen insufficient.

Network layer (L3) handoff is a procedure, during which the MN moves to a new subnet and establishes the corresponding mobility binding with the home agent. Efficient execution of the handoff requires two things: The time that the MN is unable to exchange packets with the correspondent nodes should be minimal. The number of packets dropped because of the handoff should as well be minimal. The handoff process requires several tasks from the MN:

1. At the link layer an association to the new subnet is formed. This is simple with wired media, but with wireless medium like IEEE-802.11 a link layer handoff algorithm is needed. The link layer association may be possible with multiple access points at a given time, and the most suitable access point should be chosen. The link layer handoff algorithm is responsible for selecting the target access point.

2. The network layer may or may not be aware of the link layer handoff. If not, the network layer handoff algorithm has to detect the movement by other means. Mobile IP suggest that this is done by examining Agent Advertisements sent by foreign agents (FA) in the visited network. Alternatively, the link layer may provide triggers to inform the network layer about the movement.

3. Once the MN has found out that the movement has occurred, it has to start acquiring a new CoA.

4. The MN sends a Registration Request to the HA to update the mobility binding with the newly acquired CoA. The Registration Request is sent via a FA or directly to the HA in case of co-located CoA. (See Section 4.1)

5. When the HA receives the Registration Request, it answers with a Registration Reply and starts to tunnel arriving packets to the new care-of address of the MN.

In this thesis a handoff is considered to start when the link layer association to the old subnet is lost. The handoff is completed when a Registration Reply message arrives from the HA to the MN. A network layer handoff starts when the L2 association has been changed. In this sense, the time it takes to notice the movement to a new subnet is considered part of the network layer handoff. The term handoff latency denotes the time it takes to commit the whole handoff process. During the handoff the MN is unable to exchange packets with the CNs. Packets destined to the MN are lost until the HA has updated the mobility binding of the MN.

2.2 Complexities of handoff management

Handoff management achieves to maintain the communications of the mobile node at an acceptable quality of service (QoS) level despite of the mobility. This involves efficient handoff execution to minimize the handoff latency and packet drop, and

intelligent handoff target selection to avoid unnecessary handoffs and to ensure that the best possible connection is used at a time.

Possible handoff targets can be prioritized for example by bandwidth, latency, packet error rate, cost of the bandwidth, and the expected time of availability of a certain connection. The handoff algorithm makes trade-offs between these measures when selecting the handoff target. To make optimal trade-offs, the QoS needs of the currently used applications had to be known. For example, voice-over-IP needs low latency above all, whereas file download benefits mainly from high bandwidth. Unfortunately, current platforms do not commonly provide this information for the use of the handoff algorithm. Even the application itself is often unaware of its future QoS needs.

Handoffs themselves affect the user perceived QoS. Advantage of a handoff to a point of attachment of a better connection has to be balanced against the interrupt caused by the handoff itself. The gain increases concurrently with the duration that the new point of attachment can be used. The expected time at the current location may be estimated for instance by examining the past movements and the current velocity of the mobile node. Of course, this assumes that some kind of regularity can be found in the movement pattern of the particular mobile node.

Many difficulties in handoff management arise from the unpredictability of the networking environment. The network topology changes during the time as new network elements are deployed and old ones are removed. Routing is fluctuating too, routes from host to host are under, usually slow, but continuous change.

Wireless technologies make the environment even more unpredictable. The cell coverage is dependent on surrounding terrain. Buildings and other barriers introduce sudden shadowing effects, which cause rapid decaying of the signal even within few meters of distance. Components of the signal may be received out of phase, as they take different paths and reflect from surfaces before arriving to the receiver [55].

This multi-path propagation causes part of the signal to encounter delay because travelling longer paths to the receiver. Represented information symbols overlap, which impedes the demodulation of the signal. Physical barriers are often temporary, like bypassing vehicles, and therefore difficult to be prepared for. Radio links suffer also from the interference inflicted by other transmitters in the same frequency channel.

2.2.1 Divided protocol layers

Handoff management is divided between link layer handoff algorithms of IEEE-802.11 wireless network adapters and a network layer handoff algorithm. The link layer algorithm decides the access point to associate with and the timing of the L2 handoff. The network layer algorithm chooses the interface to use, acquires a CoA from the visited subnet, and handles the Mobile IP registration process. If OSI (Open Systems Interconnection) [12] reference model is complied with, the two layers are separated and they do not communicate with each other. Yet, the layers share a common goal. An independent link layer handoff algorithm is unable to choose the best target access point without the knowledge of the quality of the network layer connection that would be provided through the link. For example, while the radio signal of the access point is good, the interfaced subnet may not provide route to the home agent at all. In practise, the situation is just as described. Protocol stacks seldom offer coupling of link and network layers. Link layer algorithms are usually proprietary and executed at the device driver or the firmware of the wireless network interface card. When the behaviour of the two layers are not coordinated, the resulting behaviour is inevitable suboptimal.

2.2.2 Link layer handoff decisions

Link layer handoff decisions of an IEEE-802.11 interface are commonly based on the signal qualities of the available access points as other information can be hard to acquire. The received signal strength (RSS), signal-to-interference ratio (SIR), bit error rate (BER) or some combination of these can be used. A strong signal level implies a low packet drop, but also the likelihood of the proximity of the access point. The goal is to associate with the access point that provides the best link layer connection at time and that can be used for an extended period of time. Again, there is a tradeoff between the signal quality level and the number of reassociations needed over time. Every reassociation decreases the overall quality of service. The lack of predictability of the movement of the the mobile node and the unknown coverage areas and placing of access points makes optimal decisions impossible. The nature of the radio network causes signal strengths to alternate frequently and unexpectedly. At the physical layer, only a few centimeters movement of the receiving antenna can affect drastically to the signal quality. For the use of the link layer, the changes are averaged over time, but the problem remains at larger scale. Rapid, but temporary variations in the signal quality disturb handoff decisions. One approach is to delay the handoff until it is more evident that the weakening of the signal is definitive, and the handoff is truly needed. This remedies a "ping-pong" situation, in which the mobile node rapidly oscillates between multiple access points. Naturally, this handoff policy has its weaknesses. When the signal degradation indeed turns out to be definitive, the handoff to the new access point is unnecessarily delayed. L2 handoffs between the access points of one subnet are tolerable in comparison to those between separate subnets. Between subnets, a L3 handoff is needed and the overall handoff latency is increased usually by an order of magnitude. The L2 handoff algorithm running in the network interface card does not have the access to the information about the subnet boundaries. Due to the dynamic propagation environment, any link layer handoff algorithm that uses constant parameters to determinate the necessity of the handoff from the signal quality is suboptimal in

certain situations [58].

2.2.3 Network layer difficulties

A network layer heterogeneous handoff algorithm forms mobility bindings for care-of addresses. The decision of which network interface to choose at a time is a complex problem. The criteria includes for example the quality of the network layer connection to the corresponding nodes and to the home agent, the type of user application, the current handoff policy, the velocity of the mobile node, the cost of bandwidth, the duration that the care-of address can be used, and the services available at the interfaced network. The characteristics of the route to the home agent are often unknown at the decision making time. Without this information, the interface of the fastest available transfer media is usually chosen. Likewise, it is unknown how long a specific link layer connection is up. A fast moving mobile node may associate to an access point at L2 layer, but does not have enough time to complete the registration of the CoA (L3 handoff) until the link is lost again. The knowledge of the exact moments of the link layer handoffs would be beneficial. The L3 handoffs could be scheduled accordingly. In this case, the next network interface could be prepared in advance. This means for example, acquiring a CoA and sending the Registration Request in advance to schedule L3 handoff to the same moment that the link layer connection to the old point of attachment is lost.

2.2.4 The network aspect

From the network aspect, handoff management has other requirements. Frequent handoffs generate control traffic that burden network and should be avoided. Minimal number of handoffs allows more networking resources to be used for actual user traffic. One way to diminish the number of link layer handoffs is to delay a handoff until it is absolutely sure that one is needed. On the other hand, the delayed handoff

results in a degraded signal level and in situation, where the QoS requirements of the client may not be met. A delayed handoff in IEEE-802.11 means often that the mobile node is moving across cell boundaries and more and more into the neighbouring cell. This may interfere with the signals of other mobile nodes communicating inside the new cell. Changing traffic distribution adds in its own difficulties, especially in the wireless networks that serve mobile nodes. Mobile nodes should be evenly allocated to available radio cells. A common problem is that the traffic distribution gets unbalanced. Common parameters in the handoff algorithms may cause situations, in which the majority of the mobile nodes decide to associate to the same access point at the same time. This leads to one burdened cell, while the neighboring cells are only slightly utilized.

3 Problem statement

This thesis studies handoff management of a multiple network interface Mobile IP client in networks that lack foreing agents. The main objective is to design a new heterogeneous handoff algorithm suitable to the situation.

I will emphasize straightforward deployment of mobile nodes and implementational feasibility in the design of the algorithm. The handoff management must conform completely to the Mobile IP specification. No new network entities nor modifications to the protocol are acceptable. The design should not postulate unrealistic implementational efforts. The capabilities of the standard network interface cards and the operating system[1] limitations are taken into account. The design of the algorithm will be detailed to the degree that no ambiguities are left in the handoff management functionality of the mobile node.

Two major drivers call for this research.

- A mobile node has to cope in networks, where foreign agents are not present. In the current networking environment, foreign agents are extremely rare. It is uncertain if foreign agents will ever be deployed widely. The main benefit of a foreign agent is that the use of it spares already burdened IPv4 address space. There is a work in process at the Internet Engineering Task Force (IETF) to introduce Network Address Translation (NAT) traversal to the Mobile IP [29], which would solve the problem. Other benefits of the foreign agent like packet tunneling on behalf of mobile nodes and handoff assistance may not be found important enough to upgrade the current networks.

- A mobile node with multiple network interfaces, each of different type of media, benefits from the extended coverage area and the ability to use higher bandwidth communications whenever they are available.

[1] Windows XP and Linux are referred when the feasibility of the implementation is considered. The algorithm itself will be platform independent

The Mobile IP specification acknowledges that the mobile node may have multiple interfaces in use at the same time. However, the protocol imposes restrictions to the multi-interface handoff management. These implications are not explicitly stated in the Mobile IP specification. I will include this Mobile IP analysis to my work. Functioning without foreign agents will introduce other difficulties, especially in the movement detection. The current network interface hardware may add further limitations to the possible solutions of the algorithm.

3.1 Criteria

The following criteria should be fulfilled by the resulting heterogeneous handoff algorithm.

Minimal interference to data flow The algorithm must seek to minimize handoff originated interruptions to the data flow. This includes as efficient as possible handoff execution to reduce the handoff latency and packet loss. Unnecessary handoffs must be avoided. Reasonable trade-offs must be found between the handoff frequency and the overall quality of service of the connection.

Detailed algorithm design The handoff management must be presented to the detail that leaves no ambiguities in the functionality of the mobile node in any situation.

Implementational feasibility The implementation of the algorithm must be possible with a reasonable effort. The design must not pose unrealistic requirements to the platform, protocol stack or assumed Mobile IP implementation. The capabilities of standard network interface cards and the operating system limitations must be taken into account. The algorithm itself will be platform independent, but when the implementational feasibility is considered, Windows XP and the latest Linux distributions are treated as reference operating

systems. These operating systems are chosen for their popularity among laptop computers.

Conformance to Mobile IP Full conformance to the Mobile IP specification is required from the algorithm. No new protocol messages, message extensions or changes to the functionality of the network entities are allowed.

Efficient use of bandwidth There is a close relation between the bandwidth usage and the handoff performance. A fast movement detection may require extra network traffic. Verification of the bi-directional route to the home agent may require extra registration messages. The trade-off exists also between the frequency of the handoffs and the network usage. Naturally, each network layer handoff requires a new registration process. The active scanning of nearby access points consume the bandwidth of the air interface. The algorithm must use bandwidth sparingly.

Minimal user interaction A considerable goal is to make the handoff management transparent to the user. This is perhaps even more important than the seamlessness of the data flow. Handoffs between different transfer media must not involve user interactions. Furthermore, the user should be liberated from choosing different handoff management profiles e.g. when the velocity of the mobile node changes. This would be a notable step towards better usability of a Mobile IP client using heterogeneous transfer media.

Security Design choices made on behalf of the handoff management must not weaken the security of Mobile IP.

3.2 Scope

This thesis studies handoff behaviour only in the networks that do not comprise foreign agents. This does not prevent the implementation of the mobile node from having extra functionality for foreign agent supported handoffs. Even in case that

the network includes a foreign agent, it is likely that other methods to acquire care-of address are also available and that the algorithm can be used. Not all mobile hosts execute Mobile IP.

Media types that the heterogeneous algorithm will support are ethernet (IEEE-802.3), wireless LAN (IEEE-802.11b), and GPRS (General Packet Radio Service). These are the most likely media types used by a mobile node. Moreover, these media types differ enough to have an advantageous complementary effect on each other. The algorithm itself is not media dependent, but some exotic transfer media might impose extra requirements to the handoff management. The power consumption of different transfer media is not taken into account. Nevertheless, it might be considered, because handoffs increase network traffic and thus drain batteries.

The aggregate behaviour of the mobile nodes executing the handoff algorithm is not studied within this thesis. When multiple mobile nodes execute the same handoff algorithm with the same set of parameters at the same location, the aggregate behaviour may differ substantially from the behaviour of a sole mobile node [59]. There is a risk of oscillating congestion between access points and other unwanted phenomena. In general, this is mostly the problem of the link layer handoff algorithm, and it is acceptable to exclude it from this thesis.

The design of the handoff algorithm is dependent on the Mobile IP implementation. In this thesis, I will assume that Mobile IP registrations can be performed simultaneously through multiple interfaces. The mobile node has to be able to receive and transmit data packets simultaneously via multiple interfaces as well. These are realistic assumptions. The Mobile IP protocol requires such a functionality if simultaneous bindings are to be used. No platform limitations exists if Mobile IP is implemented under the operating system protocol stack.

This thesis is a theoretical study, and no software components are implemented, excluding some prototypes to test the feasibility of parts of the system.

3.3 Measurement of results

Most of the criteria related to the design of the handoff algorithm and the Mobile IP analysis are qualitative. Consequently, the best that can be done is to argue why a certain goal has been achieved, or to what degree it has been achieved.

Each feature, parameter or design choice of the algorithm is evaluated to identify its impact to the overall handoff management. The goal is to state the improvements, disadvantages, and limitations in a clear way to allow the reader to understand what kind of an effect the changing of each feature would bring.

A formal prove of the optimality of the algorithm is out of scope. There are simply too many variables involved in the system comprising the whole network layer handoff management. However, there are many design choices, which can be justified. To some degree, the handoff management can be separated to individual parts, whose optimality can very well be presented.

3.4 Organization of the thesis

The rest of this thesis is organized in a following manner. In Chapter 4 I will go through the Mobile IP protocol related RFCs. Chapter 5: Related work, discusses about previously presented handoff improvement methods. Mobile IP is analysed in Chapter 6 for the protocol implications to the handoff management of a multi-interfaced mobile node functioning without the foreign agent support. Chapter 7 presents a handoff algorithm that should fulfill the given criteria. In Chapters 8 and 9, I will evaluate the algorithm, analyse results, and make conclusions about the work.

3.5 Assumptions about the reader

The reader of this thesis is assumed to be familiar with the TCP/IP protocol suite and Internet routing. Comer's *Internetworking with TCP/IP* [11] gives excellent introduction to the issue. A basic knowledge of the wireless LAN medium access control and physical layer is beneficial [60].

3.6 Referencing Internet-Drafts

The Internet Engineering Task Force (IETF) requests that Internet-Drafts submitted to the IETF must not be referenced except as "Work In Progress" because they are not officially published documents. This is specified in the Internet Standards Process [5]. Contrary to this request, I have referenced Internet-Drafts by their file names, which are stated in the beginning of each draft. Following reasons advocate my decision.

1. It aids to locate drafts that are possibly older than six months. The IETF does not guarantee that drafts older than six months can be found. Internet-Drafts are meant to be only temporary working documents in the IETF standards process.

2. Internet-Drafts may contain newer information than older standards track RFCs and therefore should be taken to account.

3. Many Internet-Drafts (e.g. individually submitted ones) will never become RFCs even though they contain valuable information.

4 Mobile IP

When Internet Protocol (IP) was designed, mobility was not an important objective. Afterwards the need to move workstations to different locations was recognized. This would require that host specific routes had to be updated through the whole Internet routing fabric. The solution is obviously not feasible because of the involved scaling problems. Alternatively, a mobile node could use a new IP address in each visited network. This would prevent other hosts from starting communications with the mobile node, unless the IP address was communicated to all possible peers. Additionally, the ongoing communications would be interrupted every time the IP address changes. Mobile IP solves the problem with a cost of triangular routing, extra control traffic, and interrupts to the data traffic when the point of attachment to the Internet is changed. Mobile IP protocol was published year 1996. Despite of that, mobile IP implementations have been scarce. There are several reasons for this kind of development. Suitable wireless hardware hasn't been available for more than few years. Moreover, the standard Mobile IP includes new network entities called foreign agents. Deployment of foreign agents introduces costs and administrative efforts. The investment might have been hard to justify as the event of appearance of Mobile IP clients has been unpredictable.

Three RFCs considering the Mobile IP protocol have been published by IETF Mobile IP Working Group since 1996. The first version, RFC 2002 [42] carries name "IP mobility support". It described the Mobile IP protocol and the architectural entities involved in detail. The specification remained valid for six years until RFC 3220: "IP Mobility Support for IPv4" [45] was published. "IPv4" was added to distinguish the protocol from the mobility protocol of the Internet Protocol version 6. RFC 3344 [46] was published soon after. It merely corrected typographical and other minor errors in RFC 3220. Despite newer versions, the Mobile IP protocol has remained the same by its most essential parts.

4.1 Overview

Mobile IP allows transparent routing of IP packets to mobile nodes in the Internet. Each *mobile node* (MN) is identified by its *home address*, regardless of its current point of attachment to the Internet. A host that communicates with the mobile node is called a *correspondent node* (CN). The corresponding node does not have to be aware of the mobility of the mobile node. It can function as if the mobile node would be a normal stationary host in the Internet. The Mobile IP protocol identifies also two type of *mobility agents* (MAs), the *home agent* (HA) and the *foreign agent* (FA).

The home agent resides at the home network of the mobile node, the same network, in which the home address of the mobile node is allocated. Packets from the CN to the MN are routed using the home address. The HA intercepts the packets and tunnels them to the current location of the MN. To be able to tunnel packets to the MN, the HA must be aware of the current location of the MN. For this purpose, the MN acquires a *care-of address* (CoA) from the network it is visiting. Every time the MN moves, the current CoA is registered to the HA. A foreign agent resides in the visited network. It offers routing services for the registered mobile nodes. the FA provides the MN with a CoA and detunnels and delivers datagrams to the MN that were tunneled by the HA. The FA can also serve as a default router for the registered MNs.

Datagrams sent by the CN to the home address of the MN are tunneled by the HA to the CoA of the MN. IP-IP encapsulation [43] is the default tunneling method in Mobile IP. The IP-IP encapsulation attaches an extra IP header to the IP datagram. The routing is done according to the outer IP header, whose destination address is the CoA and the source address is the IP address of the HA. Thus, the resulting datagram is topologically sound. Inside the tunnel, the original packet remains unchanged. At the CoA, the inner IP packet is extracted and delivered to the MN. It appears to originate from the CN. The protocol stack of the MN receives the inner

packet, whose destination IP address is the home address of the MN and source IP address is the address of the CN. The upper protocol stack of the MN does not have to be aware of the mobility. From the application point of view, the home address can be used as a static IP address of the host. As the MN sends packets to the CN, it uses the home address as a source address. Packets can be routed directly to the CN or tunneled via the HA. The later method is called *reverse tunnelling* [36]. In this case, the HA is the tunnel end-point for encapsulated datagrams from the MN to the CN. The HA removes the outer header and redirects the packets to the CN. Reverse tunneling must be used in case of *ingress filtering* [17] employed by many routers. An ingress filtering router drops packets, whose source IP address does not appear to belong to any of the known networks behind the link. Either way, the CN observes that the packets are coming from a stationary host, identified by the home address of the MN. The Mobile IP specification allows three tunneling methods in total. Besides the IP-within-IP encapsulation, the General Routing Encapsulation [39] and the Minimal Encapsulation [44] methods can be negotiated. Figure 1 demonstrates the Mobile IP triangular routing and the tunneling of the packets.

The care-of address determines the tunnel endpoint of datagrams from the HA to the MN. There are two types of CoAs. The *foreign agent care-of address* is the IP address of the foreign agent, that serves as the tunnel end-point on behalf of the MN. The FA decapsulates arriving packets and delivers the inner packets to the MN. The MN obtains the foreign agent CoA by listening the Agent Advertisements sent by the foreign agents in the network. A FA can allow multiple mobile nodes to use its IP address as a CoA. The FA maintains visitor list about the link layer addresses and the home addresses of the registered MNs to deliver packets to their right recipients. This saves the Internet address space as every MN does not need its own care-of address. A *co-located care-of address* is a CoA acquired by the MN itself. The co-located CoA is associated with one of the network interfaces of the MN. In this scenario the MN acts as a tunnel endpoint and performs the encapsulation and

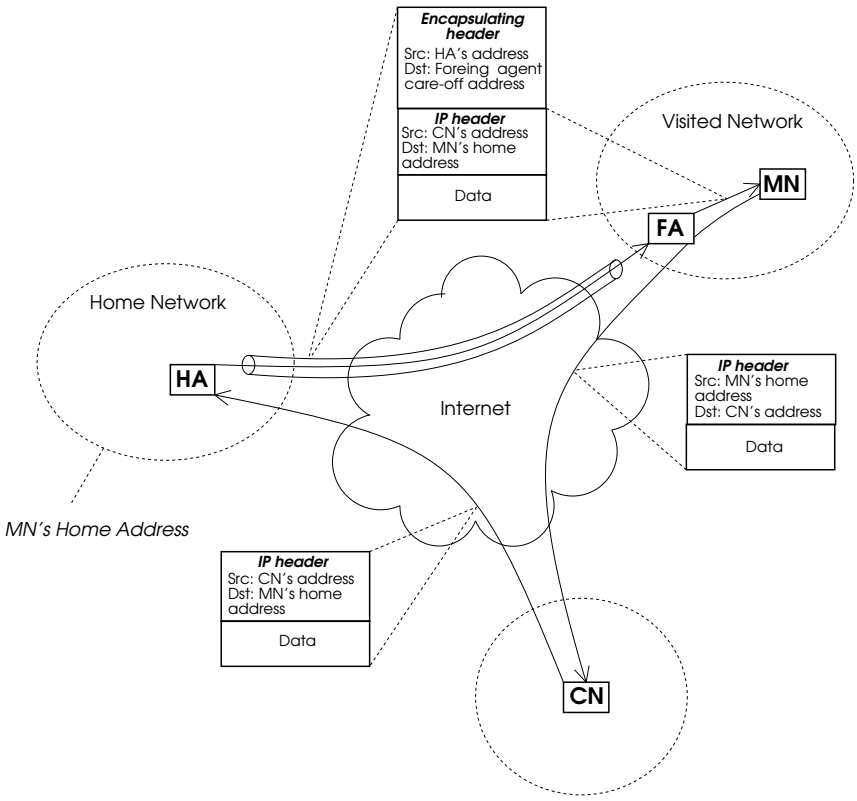

Figure 1: Mobile IP routing

decapsulation of the packets itself. A co-located CoA can be acquired dynamically using e.g. Dynamic Host Configuration Protocol (DHCP) [16]. Alternatively, the address may be owned by the MN as a long-term address for its use only, while visiting a specific foreign network. The mode of using a co-located CoA has the advantage that it allows the mobile node to function without foreign agents.

4.2 Protocol messages

Mobile IP defines two new message types: a *Registration Request* (RRQ) and a *Registration Reply* (RRP). The messages are called *Mobile IP control messages* and they are sent with User Datagram Protocol (UDP) [50] using port number 434. Control messages are used for the registration of the CoA to the HA and for committing the registration to the MN. Mobile IP defines also an *agent discovery* process. It allows mobility agents to inform mobile nodes about their existence and the services on

the link with periodic *Agent Advertisements*. Mobile nodes can solicit agent advertisements with *Agent Solicitations* to speed up the agent discovery. Both messages make use of the existing *Router Advertisement* and *Router Solicitation* messages defined for ICMP Router Discovery [13].

4.2.1 Agent discovery

A mobile node uses agent discovery process to determine whether it is connected to a foreign network or to its home network. Consequently, Agent Advertisements reveal movement between two foreign networks. Agent Advertisements are ICMP Router Advertisements with a *Mobility Agent Advertisement Extension*. The extension defines the type of the mobility agent (home or foreign agent), several capabilities of the mobility agent, and available CoAs (in case of FA). Two optional extensions are defined: A *Prefix-Lengths Extension* for network prefixes of the router addresses listed in the ICMP Router Advertisement portion of the Agent Advertisement and a *One-Byte Padding Extension* to make the length of the ICMP message even number of bytes, required by some IP protocol implementations. Other extensions may be defined in the future. When the mobile node receives an Agent Advertisement from its home agent, it should deregister and start to act as a stationary Internet host. The Agent Solicitation is identical to the ICMP Router Solicitation. No authentication is required for the Agent Advertisement and Agent Solicitation messages.

4.2.2 Move detection

Fast network layer movement detection is important in efficient handoffs. The procedures to obtain and register a new CoA can begin only after the movement to a new network has been noticed. Mobile IP suggests two methods for the network layer movement detection. The first method is based on the registration lifetime field of the Mobility Agent Advertisement Extension. If currently negotiated life-

time expires while the mobile node fails to receive further advertisements[2] from the foreign agent, the connection is assumed to be lost. The second method is based on the network prefixes. Once the mobile node receives an Agent Advertisement with a different network prefix than the network prefix of the visited network (or of the previous advertisements), the change of the network has occurred. Without foreign agents, the movement can be detected with the network prefix method if router advertisements are received, and the mobile node knows the network prefix associated with its co-located care-of address. However, not all routers send router advertisements.

4.2.3 Registration messages

The mobile node exchanges registration messages with the home agent, and optionally with foreign agents, to form mobility bindings i.e. to register its CoAs. The process associates the mobile node's home address with the CoA for a certain lifetime. If the registration is done via a foreign agent, it processes and relays the Registration Request to the home agent. Similarly, when the foreign agent receives an authenticated Registration Reply, it forwards the Reply to the mobile node after updating it's visitor list. To renew the mobility binding, the mobile node repeats the registration process. The packet formats of the Registration Request and the Registration Reply will be described in more detail in Chapter 6. The Mobile IP registration messages and the agent discovery process are illustrated in Figure 2.

Each mobile node and home agent maintain a *Mobility Security Association* for authentication of the registration messages. They are indexed by *Security Parameter Indexes* (SPIs) and the home addresses of the mobile nodes. Messages between the mobile node and the home agent are authenticated with a *Mobile-Home Authentication Extension*. Optionally, security associations may also exist between the MN and the FA, or between the FA and the HA. A *Mobile-Foreign Authentication Extension*

[2]At least three advertisements must be sent during the advertisement lifetime.

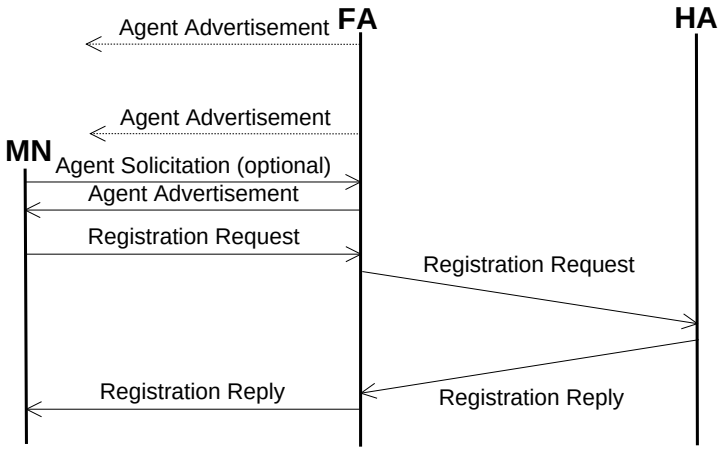

Figure 2: Agent discovery and registration messages

and A *Foreign-Home Authentication Extension* are defined for these purposes. The authentication extensions contain A *Message Authentication Code* (MAC) to ensure that the messages have passed unmodified over the Internet. The MAC is a 128-bit message digest calculated over the packet beginning from the end of the UDP-header to the beginning of the extension. The default authentication algorithm is HMAC-MD5 [28]. In all registration messages, The Mobile-Home Authentication Extension must be the first extension.

Mobile IP does not provide any data encryption, only authentication for the control messages. To achieve privacy for user traffic *IP Security* (IPsec) [25] can be employed on top of Mobile IP.

4.2.4 Extensions

Several message extensions have been defined in the Mobile IP related RFCs. RFC 3012 [47] defines a *Mobile IP Agent Advertisement Challenge Extension* for Agent Advertisements and a *MN-FA Challenge Extension* for the registration messages. They can be used to enable replay protection of registrations between the foreign agent and the mobile node. A *Mobile Node NAI Extension* [8] provides the mobile node with a mechanism to identify itself with a Network Access Identifier (NAI) [1].

AAA (Authentication, Authorization and Accounting) servers at foreign domains use the NAI to identify their clients. RFC 2356 [37] defines a *Traversal Extension* for the registration messages. The Traversal Extension identifies IP address of a SKIP firewall, which access policies can then be dynamically configured in order to allow traffic pass to and from the mobile node. RFC 3025 [14] defines vendor specific extensions to Mobile IP messages. The extensions allow Mobile IP vendors to include information of their own interest to the messages. There are two type of vendor specific extensions: A *Critical* (CVSE) and a *Normal* (NVSE) *Vendor/Organization Specific Extensions*. The main difference is that messages that contain an unknown CVSE must be discarded, while unknown NSEVs can simply be ignored.

5 Related work

5.1 Overview of handoff management

Handoff management is involved in every mobility supporting networking architecture. Handoffs are managed for example in cellular networks, Internet Protocol (IP) networks and even between earth orbiting satellites and their ground stations. While technologies and protocols vary, there are general concepts by which different types of handoff management schemes can be classified.

K. Pahlavan et al. [40] classified handoff management by architectural and decision time algorithm issues. The handoff architecture defines the framework of the system in which handoffs occur. The handoff architecture can be divided into three concepts: *Handoff control*, *handoff methodology* and *handoff procedures*.

The handoff procedures involve a set of protocols to notify all the related entities of a particular connection that the handoff has been or will be executed. In Mobile IP, these entities are mobile node, home agent, and involved foreign agents. At the IEEE-802.11 link layer, handoff procedures would include Inter Access Point Protocol (IAPP) [26], which access points use to share information in reassociation situations. IAPP enables old access point e.g. to free resources and to transfer the security context to the new access point.

The handoff methodology classifies handoffs by their qualities. A soft handoff occurs if the MN can communicate simultaneously with two or more access points during the handoff. A hard handoff occurs if only one connection is available at time. A fast handoff introduces only a short interruption to the communication. A smooth handoff causes an interruption, but no packets are lost. A seamless handoff occurs if no interruption or packet loss can be perceived.

Handoff control denotes the party that makes the handoff decision. In a *network*

controlled handoff (NCHO) the network side entities decide when the handoff should occur. The network controls the execution of the handoff. A *Mobile controlled handoff* (MCHO) stands for exactly the opposite approach. The mobile terminal makes the handoff decision and handles the registration at the new base station or access point. A *Mobile assisted handoff* (MAHO) is between the NCHO and MCHO. Usually, the mobile terminal sends measurements about the received signal to a network entity, and the network makes the handoff decision based on the data.

Decision time algorithms are classified by the metrics that are used and by the type of the algorithm. The algorithm can be anything between a linear function measuring signal strengths and a neural network operating with various types of inputs. Handoff management classification to architectural issues and decision time algorithms by K. Pahlavan et al. [40] omits the heterogeneous network layer handoff management altogether. They identify Mobile IP as one solution for IP mobility, but do not discuss about controlling the registration processes between multiple interfaces of different transfer media.

Another good classification for the handoff management frameworks has been given by G. Dommety et al. [15]. They state that the handoff procedure is determined by the following factors.

The element that makes the handoff decision Either the network or the mobile node can initiate and/or execute the handoff.

The amount of information available before the handoff The handoff procedure is accelerated if the configuration parameters needed at the new point of attachment can be obtained before the handoff occurs. Knowing the gateway IP and MAC addresses, or the new care-of address to be used in the visited network allows a faster network layer handoff.

The extent of link layer support Link layer triggers and statistics can be used to improve handoffs. The amount of coupling between the link layer and the

network layer is an essential factor in speeding up handoffs.

The amount of prediction that is possible If the movement of the mobile node can be predicted, it is possible to achieve lower latency handoffs and choose more wisely the next point of attachment.

The type of link layer The type of link layer has its implications to handoffs. A connection oriented link layer has longer setup times to a new point of connection. In contrast, it may support quality of service features. If the link layer supports connecting to two point of attachments at the same time, it is possible to achieve seamless handoffs.

5.2 Mobile IP handoff enhancements

A lot of research has been conducted to improve the network layer handoff of Mobile IP. The vast majority of the improvements are based on foreign agent co-operation and assistance in handoff situations. The improvements can be classified to four categories.

Localized registrations Foreign agents can be arranged in hierarchies to localize registrations within the hierarchy. The round-trip time of the registration messages is almost completely eliminated. The registration is terminated at some higher hierarchy level foreign agent much nearer to the mobile node than the home agent. This foreign agent remains at the location where routes to the old and the new visited subnet would diverge.

Multicasting Multicasting is an efficient handoff improvement method inside a foreign agent hierarchy. Seamless handoffs can be acquired with less bandwidth consumption compared to the Mobile IP simultaneous binding. The simultaneous binding duplicates the traffic between the home agent and the mobile node, while the multicasting can be limited to the hierarchy.

Buffering Packet buffering is a method that eliminates the packet loss, but not the handoff latency, thus resulting to soft handoffs. In this method, foreign agents buffer packets addressed to the mobile node. In a handoff situation, the packets are not lost at the target network, but delivered when the registration process has completed.

Forwarding Foreign agents may execute protocols to form tunnels between each other to redirect the mobile node intended traffic. When a mobile node is handed over to a new foreign agent, the old foreign agent starts to tunnel packets to the new location. The packets may traverse via multiple foreign agents until the mobile node is reached. The actual registration process can be postponed to a time that is suitable to the mobile node.

Most of the handoff architecture proposals use some combination of above-mentioned methods.

Localized registrations have been used in many micro mobility architectures. Cellular IP [9] developed by A. Campbell et al. introduces mobility management ideas from telecommunications networks like GPRS to IP mobility. Mobile IP handles the macro mobility i.e. moving between Cellular IP domains, while Cellular IP manages the micro mobility. The location management is integrated to routing, and schemes like paging are used to track the mobile node. HAWAII [52] is another micro mobility extension to Mobile IP. It introduces Quality of Service support to handoffs within the HAWAII hierarchy.

S. Seshan [53] developed a handoff scheme for cellular data networks that relies on mobility hints i.e. in-cell location information and user movement patterns to predict handoffs. Packets are multicasted to several neighbouring base stations to achieve fast handoffs. Buffering at the base stations ensure that no packets are lost. R. Cáceres and V. N. Padmanabhan [7] presented another buffering-based solution. Their solution employ packet buffering and interaction between access points to

forward traffic. The new access point requests buffered packets from the old access point and then delivers them to the mobile node via the new link. Multicast-based architectures have also been studied by C. Tan and S. Pink [54]. Two layer foreign agent hierarchy is utilized. Domain FA at the top of hierarchy multicasts packets to the lower FAs. The FA that has a link layer association with the mobile node transmits them over the wireless link.

5.2.1 Mobile IP low latency handoffs

K. E. Malki et al. [30] present two techniques to support delay-sensitive and real-time applications in Mobile IPv4: *Pre-Registration* and *Post-Registration*. The Pre-Registration is based on link layer triggers, and proxy Agent Advertisements and Agent Solicitations. In advance of a handoff, the neighbouring foreign agents exchange Agent Advertisement messages with each other. A link layer trigger at the mobile node or, either the old or the new foreign agent triggers the handoff. Note that this happens before the actual L2 handoff. The mobile node receives a *Proxy Agent Advertisement* from the new foreign agent relayed by the old foreign agent. Now, the mobile node can issue a Registration Request to the new FA via the old FA. This allows the mobile node to use the old FA until the registration completes. If the L2 handoff is scheduled to the same moment that the registration completes, the overall handoff latency will be reduced only to a L2 handoff latency. The Post-Registration method forms tunnels between foreign agents to forward arriving packets to the current location of the mobile node. Packets from the home agent are first received by an anchor FA. The anchor FA is the FA that relayed the last Registration Request/Reply pair to the home agent. The anchor FA forwards packets to the subnet of the foreign agent that currently serves the mobile node. There is no need for further registrations until the anchor FA has to be changed. The mobile node can postpone the registration to a time that it sees the most appropriate. Tunnels between foreign agents are established by a *Handoff Request* and

a *Handoff Reply*, the messages that are exchanged upon the L2 trigger that indicates the handoff. The Post- and Pre-Registration methods can be combined: The Pre-Registration method is started before the L2 handoff. If the registration does not resolve at a time, the old FA starts to forward packets to the new FA as in the Post-Registration method.

J. Kempf and J. Wood [24] have analysed Mobile IPv4 low latency handoff algorithms. The Pre-Registration, Post-Registration and standard Mobile IPv4 (SMIPv4) handoff algorithms were compared. Their conclusions were that The Pre-Registration method is difficult to implement. Furthermore, it behaves badly if the time interval between the L2 trigger and the beginning of the L2 handoff is either too short or too long. When the interval is too short, the Pre-Registration does not have enough time to complete, and the handoff procedure falls back to the standard Mobile IP handoff. If the interval is too long, packets are lost because they are routed to the new FA before the link layer association has been formed. The Post-Registration algorithm is straightforward to implement, but required good operating system support for the tunnel management. Both the algorithms achieved about the same level of optimization, provided that the Pre-Registration method can be tuned. In reality this is not possible, which favors the Post-Registration method.

5.2.2 Mobile IP regional registration

Approach of E. Gustafsson, A. Jonsson and C. Perkins [20] to improve handoffs utilizes foreign agent hierarchies. Their work defines Mobile IPv4 Regional Registrations, which allow mobile nodes to register locally within a visited domain. This diminishes the number of signaling messages to the home agent and reduces the handoff latency. The specification introduces a new network entity: a *gateway foreign agent* (GFA) and new messages: a *Regional Registration Request* and a *Regional Registration Reply*. When a mobile node enters a domain supporting regional registrations, it notifies the home agent about the new care-of address, the address

of the gateway foreign agent. This care-of address will not change if the mobile node changes a foreign agent under the same GFA. Location updates are handled within the hierarchy using regional registrations. The hierarchy can have multiple levels and cover wide geographical areas.

5.2.3 Conclusions

There is pros and cons in all the presented handoff improvement methods involving foreign agents. Multicasting wastes bandwidth. Packet buffering slows down routing and requires resources at foreign agents. Forwarding between foreign agents adds latency and may introduce routing loops. Some foreign agents hierarchies execute protocols that are incompatible with standard mobile nodes. Nevertheless, benefits of seamless handoffs would overcome involved drawbacks. Additional messaging at the fixed side of the network is tolerable. Bottlenecks are anyway at the radio side of the network. Achieving universal seamless handoffs is mostly a matter of choosing one micro mobility solution, or several compatible solutions, and deploying foreign agents everywhere. Unfortunately, none of the micro mobility solutions has been widely accepted, and the worldwide foreign agent deployment may be unrealistic.

5.3 Transport layer enhancements

When handoff latency and packet loss cannot be avoided, it is reasonable to minimize the harm. Transmission Control Protocol (TCP) [51] retransmission and congestion control mechanisms are designed for fixed networks. There is no equivalent to handoff latency in traditional networks, hence the interruption is interpreted as network congestion. The action is to slow down the transmission, which in case of a handoff further increases the experienced latency. M. Miyoshi et al. [35] have studied the TCP performance over wireless links and came to the conclusion that some kind of a error correction scheme is required in addition to regular TCP. Similar work

has been carried out by A. Fladenmuller and R. De Silva [18]. They came to the conclusion that either TCP or Mobile IP has to be changed to accomplish satisfactory behaviour. They make also a clever note that most of the delay-sensitive applications are build over UDP, which reduce the criticality of TCP performance.

H. Balakrishnan and S. Seshan [3] solve the TCP performance problem by buffering packets at base stations and examining the TCP data flow. Should the packet be lost or corrupted in the wireless media, it is retransmitted across the link. Handoff situations are handled by a low-latency multicast-based handoff algorithm. Their solution offers even twenty-fold throughput speedups over regular TCP in environments that are extremely error-prone. Other TCP improvement work has been done in [10, 62, 41, 22].

K. Brown and S. Singh [6] modified the UDP protocol [50] for wireless links: M-UDP. Like UDP, the M-UDP does not provide reliable transfer, but focuses to keep the number of erroneous datagrams small. M-UDP is based on datagram buffering at the base stations. The mobile node may be unable to inform the base station about the start of the handoff (i.e. the link is suddenly lost). Instead, the mobile node tracks the time during the interruption and communicates it to the base station when it is able to transmit again. If the buffer was large enough, the lost packets can be determined and retransmitted.

5.4 Link layer handoff algorithms of IEEE-802.11

Prior to the network layer handoff, a new link layer connection must be established. Wireless transfer media use link layer handoff algorithms to control handoffs between base stations of telecommunications networks like GPRS or access points of IEEE-802.11. Handoffs in telecommunications networks have been researched extensively, and it can be stated that satisfactory solutions have been found. The characteristics of IEEE-802.11 limit the downright use of those algorithms in the IEEE-802.11 envi-

ronment. The cell size is modest compared to the telecommunications technologies, which makes handoffs (reassociations) more frequent. Different radio frequencies and symbol modulation schemes affect the obstacle penetration and the rapidness of the increase of errors caused by a weakening signal. The protocol does not offer sophisticated methods for access points to assists mobile nodes in handoffs. Placing of access points can be quite accidental, and there is no guarantee of a solid coverage. However, the concepts of the signal-quality-based handoff management of telecommunications networks can be applied also to IEEE-802.11. K. Pahlavan et al. [40] give an overview of IEEE-802.11 and GPRS handoffs.

Some handoff architectures treat handoffs of two layers as separate problems. However, information from the link layer (L2 triggers) can be used to improve the overall handoff performance. Better yet, if the two layers work in conjunction, network layer criteria can be considered while deciding target access points in IEEE-802.11.

Link layer handoff algorithms base their handoff decisions on various metrics. The most used metric is received signal strength (RSS) of the serving, and adjacent access points. G. Pollini et al. [49] identify other metrics like a signal-to-interference ratio (SIR), bit error rate (BER), block error rate (BLER), the path loss, carrier-to-interference ratio (CIR), symbol error rate (SER), power budgets, and cell ranking. There are many types of L2 handoff algorithms. Along with traditional algorithms, fuzzy logic [27], neural networks [38], pattern recognition [31] and signal-prediction-based [23] algorithms have been developed. Neural network and fuzzy logic algorithms usually provide good behaviour in test environments, but the proof of uniformly superior behaviour is hard to acquire. The signal prediction algorithms estimate future signal levels from the past values [23]. The prediction is possible because of shadow fading variations are correlated [19]. The correlation exists in cellular systems, but it may not be as strong in the IEEE-802.11b environment because of the modest obstacle penetration of the 2,4 GHz frequency and the smaller cell size. L2 handoff algorithms behave better if the measured signal values could

be mapped to the velocity of the mobile terminal. J. M. Holtzman [21] estimated the velocity by examining Doppler Effect in the received signal. In IEEE-802.11 this would require access to the physical layer, which is not supported by current hardware. N. D. Tripathi gives a great survey of L2 algorithms in his Phd. thesis [57].

G. Pollini [49] demonstrated features of a traditional L2 handoff algorithm that uses RSS or signal quality metrics. Traditional handoff algorithms compute moving average of the chosen metric to smooth rapid variations in the signal strength. Hysteresis margins, thresholds, and dwell timers can be used to enhance the algorithm. For example, simple RSS-based algorithm measures signal strengths of various base stations. The base station with the strongest signal is selected. Because signal strengths variate considerably during the time, this leads to many unnecessary handoffs. This could be remedied by widening the averaging window, but it would delay the handoff. The use of threshold cuts down the number of unnecessary handoffs. The base station or the access point is switched only if the signal level of the current base station falls below the threshold. The threshold value is usually set to the minimum signal strength that provides an acceptable quality of connection. The hysteresis margin allows the handoff only if the signal strength of the new base station is greater than the old one added with the hysteresis margin. This prevents a "ping-pong" situation, rapid handoffs between two access points, when the signal strengths are under the threshold. Sometimes the dwell timer is used in conjunction with the threshold and the hysteresis margin. The dwell timer is started when the handoff condition is true. The handoff is committed only if the condition is also true when the timer expires. Figure 3 illustrates the time of handoff initiation using several L2 algorithms.

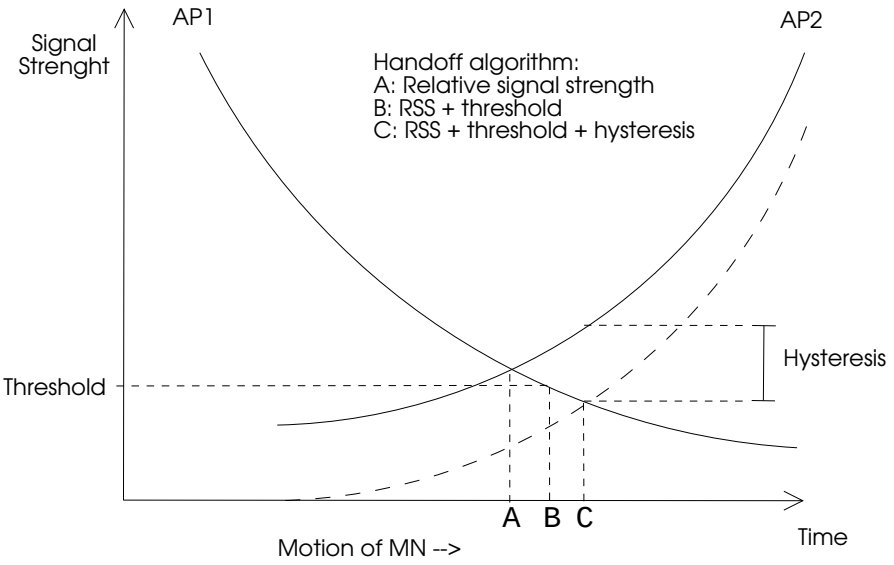

Figure 3: RSS-, threshold- and hysteris-based L2 algorithms

5.5 Handoff policies and heterogeneous algorithms

Handoff policy denotes the criteria, which is used in selection of the target point of attachment. Metrics like bandwidth, latency, cost of bandwidth, services at the network, power consumption, connection setup time, traffic load at the access point or characteristics of the network layer connection to the home agent or the corresponding nodes can be used. The handoff policy determines what criteria is weighted over another. H. J. Wang [59] did the ground work in handoff policies by presenting a *cost function*. The cost function allows comparison of the metrics of different connections and produce a prioritization between the available network connections. Being able to select the best possible network is not enough by itself. For a moving mobile node, the best network may not be available after awhile, and a new, subsequent handoff is needed. To provide good overall QoS, the number of handoffs must be minimized. V. Bharghavan and M. Jayanth [4] approached the best cell selection problem by trusting that people tend to move along routine paths. The technique collects mobility profiles from multiple users to a profile server, where all the profiles are aggregated. The movement of a mobile node can be predicted to reasonable accuracy in familiar environments.

Heterogeneous handoff algorithms have been studied by M. Ylianttila et al. [61] and N. An et al. [2]. Heterogeneous handoff algorithms control handoffs between transfer media, and in case of mobile IP, handle also the registration process. K. Pahlavan et. al [40] describe how Mobile IP and the L2 handoff algorithms of GPRS and IEEE-802.11 are related. In addition, they present a rather simplistic heterogeneous handoff algorithm for GPRS and IEEE-802.11. The current handoff research is mostly divided to the L2 handoffs of each transfer medium and media independent network layer algorithms handling the registrations. Difficulties in heterogeneous handoff algorithms and the scheduling of Mobile IP registrations are seldom addressed in handoff research.

6 Mobile IP and multiple network interfaces

In this chapter I will examine Mobile IP to identify the protocol's implications to the handoff management of a multiple interfaced mobile node in the environment that lacks foreign agents. RFC 3344 [46] requires that a mobile node has to be able to handle multiple registration processes simultaneously i.e. to send multiple Registration Requests and getting several, possible re-ordered Registration Replies back. This is not straightforward in case the mobile node has multiple network interfaces. Registration Requests can not be issued arbitrarily. Replay protection mechanisms and rate limitation of Registration Requests pose restrictions to the sending of multiple registrations. If simultaneous Registration Requests are issued without care, they interfere with each other. Replay protection at the home agent may reject some of the Registration Requests, depending on the receive order. Mobile IP does not identify any scheme to be followed when registering via multiple interfaces. I will present two registration schemes that optimize simultaneous registration processes. These schemes are called the *initial registration scheme* and the *basic registration scheme*.

A network layer handoff algorithm perceives interfaces either to have or not to have L2 connection. Only the interfaces with working L2 connection are usable from the L3 point of view. This list of usable interfaces changes dynamically as the L2 handoff algorithms of network adapters form and lose L2 connections. I will assume that the interfaces have a priority order i.e. it is more preferred to register via one interface than another. The priority of an interface can be related to transfer media type or other criteria discussed in Chapter 7.

Reachability of the home agent via an interface can not be taken as granted. Route to the home agent may be only uni-directional, because of congestion or misconfigured routers. Some access networks may not provide connection to the Internet at all. Many organizations separate their wireless LANs from the Internet, allowing

only users with strong encryption methods like VPN to go past the access network, possibly to the company's intra network. WEP (Wireless Equivalent Privacy) may not be employed at all in WLANs, as it is seen encryption wise inadequate anyways. This allows all mobile nodes to form useless L2 connections to the network. Furthermore, nowadays lot of wireless networks are used for testing purposes only. Thus, no access is provided beyond the network. If the registration process is started in a network with a broken route to the home agent, handoff latency is doubled at minimum. The mobile node has to wait and retransmit RRQs until the failure is concluded. Depending on the retransmission period, this usually lasts longer than a successful registration via another interfaced network. Although these situations may be rare in general, the impact to the handoff latency is severe. More importantly, if a user tends to move routinely in the vicinity of broken networks, inefficient handoff management is disastrous for that specific mobile node.

Mobile IP does not provide means to verify home agent reachability via an inactive interface. Registration Requests can not be used for advance verification, as the mobility binding changes along each accepted Registration Request. In handoff situations, the mobile node therefore possesses no advance information about the reachability of the home agent via any interface. The L3 handoff (registration) decisions have to be done blindly.

Following sections discuss how Registration Requests should be send in handoff situations to quickly obtain a mobility binding via the highest priority interface. I will assume that the mobile node is willing to allocate co-located care-of addresses and send Registration Requests simultaneously through all its interfaces. In practice, the mobile node may want to limit the number of simultaneous registrations to save network bandwidth. This study is applicable if the registration is done using one or more interface at a time. Along with scheduling Registration Requests, I will also cover more general issues about registration management. Finally, I will consider alternative methods to test the home agent reachability via lower priority interfaces

without committing to the new mobility binding.

6.1 Registration messages and replay protection

Mobile IP uses timestamp- and nonce-based replay protection to protect against replay attacks of registration messages. Both the Registration Request and the Registration Reply have a 64-bit identification field for this purpose. The identifications of consequent requests and replies are matched to prevent the attacker from copying messages of the previous registration and then issuing them later on. A replayed Registration Request would change the mobility binding of the mobile node, and a replayed Registration Reply would fool the mobile node to believe to a false mobility binding. The support for timestamp-based replay protection is mandatory for all mobile nodes and home agents. Nonce-based replay protection is optional, mainly because of IPR (Intellectual Property Rights) issues involved. Figure 4 shows the message format of the Registration Request (RRQ).

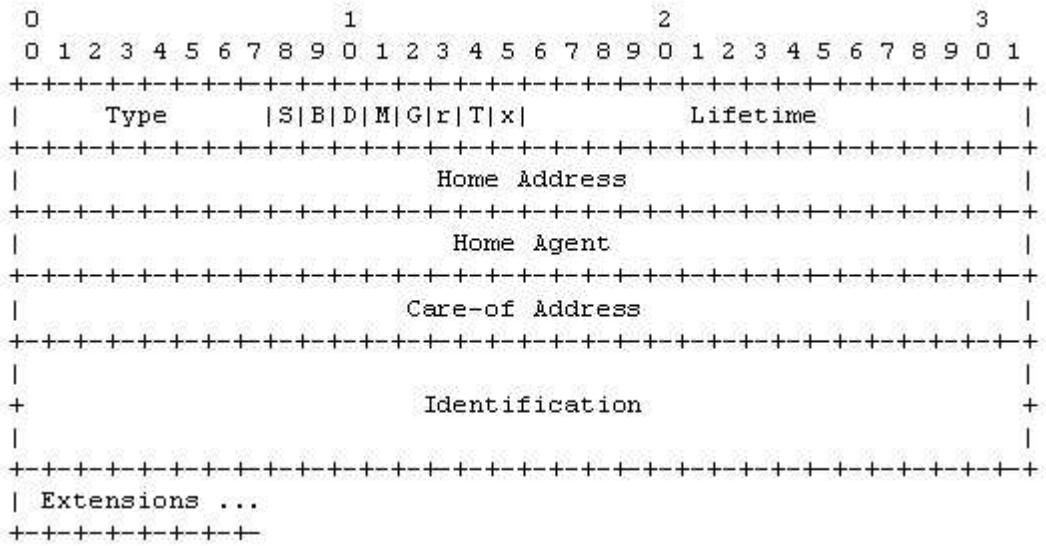

Figure 4: Registration Request

Table 1 explains the fields of the Registration Request message. The IP source address of the Registration Request is set to the co-located care-of address, when

the mobile node is registering directly, without a foreign agent. The IP destination address is the address of the home agent. The destination UDP port is set to 434 and the source UDP port is variable.

Table 1: Registration Request fields

Field	Description
Type	Value 1 indicates Registration Request.
S	If the "S" bit is set, the mobile node requests simultaneous binding.
B	Broadcast datagrams. A request to the home agent to tunnel also datagrams that are broadcasted in the home network.
D	Decapsulation by the mobile node i.e. the mobile node is using CCoA.
M	Minimal encapsulation requested.
G	GRE encapsulation requested.
r	Sent as zero; ignored on reception.
T	Reverse tunneling requested.
x	Sent as zero; ignored on reception.
Lifetime	The requested lifetime for the mobility binding of this CoA. Zero indicates deregistration.
Home Address	The IP address of the mobile node.
Home Agent	The IP address of the mobile node's home agent.
Care-of Address	The IP address of the end of the tunnel.
Identification	A 64-bit number, constructed by the mobile node, used for matching Registration Requests and replies. Used for replay protection.
Extensions	Message extensions like Mobile-Foreign Authentication Extension

The mobile node chooses the value of the Identification field in accordance with the replay protection mode. Replay protection is part of the mobility security association the mobile node shares with its home agent. Regardless of which replay protection mode is used, the lower 32-bits must remain unchanged between the Registration Request and the corresponding Registration Reply. Foreign agents use low-order 32 bits to match the Registration Request to the Reply. The Identification field must not be the same in two adjacent Registration Requests, and the same value should

not repeat while the same security context is being used between the mobile node and the home agent.

6.1.1 Timestamp-based replay protection

Timestamp-based replay protection fills the Identification field of the Registration Request with a NTP timestamp [33]. The highest 32 bits are seconds since "0 hours on 1 January 1900", and the lower 32 bits are fractional seconds. The home agent checks that the timestamp in the Identification field is sufficiently close to its own time of day. The default accepted deviation is 7 seconds. The timestamp acts also as a sequence number. The home agent accepts Registration Requests only with rising consequent Identification values. Without this, it would be possible for an attacker to repeat Registration Requests within the clock synchronization. Upon reception of an authenticated Registration Request, the home agent replies with a Registration Reply. If the Identification value satisfies the above-mentioned conditions, the entire Identification field is copied to the Registration Reply. If the Identification is not valid, only the lower 32 bits are copied to the Registration Reply. Additionally, the home agent fills the higher 32 bits with its own NTP timestamp. In this case, the home agent rejects the registration and supplies the Code 133 (identification mismatch) in the Registration Reply. As the mobile node receives the Registration Reply indicating rejection, it has to verify that the low-order 32 bits of the Identification field match with those of the Registration Request. After this, the high-order bits can be used for clock synchronization. Lost Registration Requests are retransmitted with Identification field containing a new timestamp. In effect, they are counted as new registrations. Technically, the home agent is not obliged to send a Registration Reply at all if there is an identification mismatch. In practise this is out of question. The mobile node would be no longer able to synchronize, and thus register, with the home agent.

Message format of the Registration Reply (RRP) is shown in Figure 5. The *Type*

field is set to 3 for a Registration Reply, and the *Code* field is a value indicating the result of the registration.

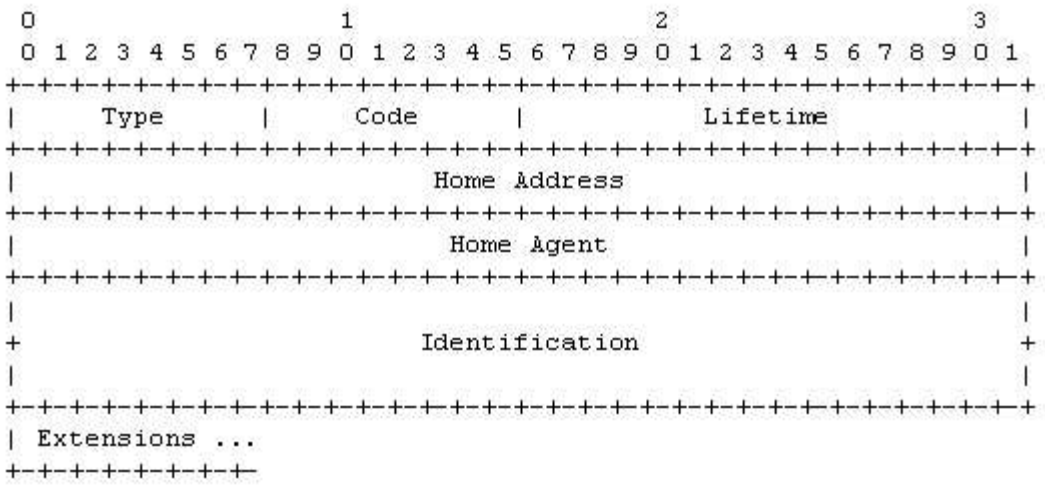

Figure 5: Registration Reply

6.1.2 Nonce-based replay protection

Nonce-based replay protection is utilizes one-time random numbers i.e. nonces that are exchanged between two communicating nodes. The basic principle is that each node supplies a random number in each authenticated message it sends. This nonce is expected to be returned in the replying message from the other node. When both the nodes employ this strategy, they can be sure that no messages are replayed by an adversary.

When the home agent receives a Registration Request, it copies the low-order 32 bits of the Identification field to the same location in the Registration Reply. The high-order 32 bits are filled with a new nonce. The mobile node receives the Registration Reply and checks that the low-order 32 bits are the same it sent in the Registration Request. It saves the high-order 32 bits of the Registration Request to be used later as high-order bits of the next Registration Request. For that Registration Request, the mobile node generates a new nonce and inserts it to the low-order 32 bits. This way, the nonces of the mobile node are carried in the low order 32

bits of the Identification field of registration messages, and the nonces of the home agent are carried in the high-order 32 bits. Nonce-based replay protection is self-synchronizing. Every time a Registration Request is rejected because of an invalid nonce, the Registration Reply provides the mobile node with a new nonce to be used in the next registration. Lost Registration Requests are retransmitted unchanged in contradistinction to timestamp-based replay protection. This evades the need for resynchronization, as the home agent always receives a valid nonce. Table 2 illustrates differences between timestamp- and nonce-based replay protection modes.

Table 2: Timestamp- and nonce-based replay protection

	Timestamp	**Nonce**
RRQ, Identification field:	Higher 32 bits: NTP timestamp. Lower 32 bits: fractional seconds.	Higher 32 bits: Copied from the last RRP. Lower 32 bits: MN's nonce.
Validity checks performed by the HA:	Higher 32 bits larger than in the previous RRQ and within the time synchronization.	Higher 32 bits must be the same that in the last RRP
RRP, (registration accepted) Identification field:	The whole field is copied from the last RRQ.	Higher 32 bits: HA's nonce. Lower 32 bits: copied from the last RRQ.
RRP, (identification mismatch) Identification field:	Higher 32 bits: HA's NTP timestamp. Lower 32 bits: copied from the last RRQ.	Higher 32 bits: HA's nonce. Lower 32 bits: copied from the last RRQ.

6.2 Simultaneous registration processes

The replay protection modes limit the sending of Registration Requests as seen in the previous section. In addition, RFC 3344 [46] states that the mobile node must not send Registration Requests more often than once per second on average. If there is a link layer move detection method that indicates a change of network, Registration Requests may be issued more often. And finally, The minimum time

between retransmitted Registration Requests must not be less than one second. The retransmissions should follow a binary exponential back off algorithm. These matters require that Registration Requests must be scheduled carefully between multiple interfaces.

As explained earlier, interfaces are assumed to have a priority order. The priority of an interface can be seen as a function of the type of the link layer media and other, interfaced network related characteristics. Consequently, the priority order may change upon the link layer handoff that changes the interfaced subnet. The prioritization of the interfaces is discussed later in Section 7.7 describing the design of the heterogeneous handoff algorithm. For now, it is just assumed that the prioritization exists. The goal of the mobile node is to fast achieve the best possible connection i.e. to register via the highest priority interface. Interfaces with no L2 connection are excluded from the list of usable interfaces.

Registration Requests are not needed if the L2 handoff is made within a subnet. The mobile node can merely issue a gratuitous ARP [48] to update possible bridges or switches in the subnet. These situations are not covered in this discussion

6.2.1 Managing multiple interfaces

I define three registration states for interfaces:

Active The mobile node has the current mobility binding via the interface.

Failed The mobile node has tried to register via this interface several times, but failed. The *failed* state should be changed to the *unknown* state after some reasonable time has passed.

Unknown This is a state that the mobile node has no information whatsoever about the probability of success if registration was attempted. Interface is in the *unknown* state after a L2 handoff to another subnet.

In addition to the state information, the mobile node should record the time of last registration attempt for each interface. When the interfaced network changes, the registration state and the time of the last registration are cleared.

The behaviour of the mobile node is dependent on the states of its interfaces. Generally, if there are higher priority interfaces available than the *active* interface, the registration should be tried via those interfaces. If no higher interface provides connection to the home agent, they should be put to the *failed* state. *Failed* interfaces should eventually be tried again as the reachability situation may change. Since Registration Requests change the mobility binding, it is out of question to try registering via lower priority interfaces than the *active* interface. It must be ensured that registration attempts do not interfere with the renewing of the mobility binding of the *active* interface.

The registration states of available interfaces constitute configurations that require different registration actions. The relevant *interface configurations* are:

- There is an *active* interface, but also higher priority interfaces in the *unknown* state. Registrations should be attempted via the *unknown* interfaces, starting from the highest priority interface.
- There is an *active* interface, and all the higher priority interfaces are in the *failed* state. No tasks are required.
- The highest priority interface is in the *active* state. No tasks are required.
- There is no *active* interface, but *unknown* interfaces exists. registrations should be attempted via the *unknown* interfaces according to the interface priorities.
- There is no *active* interface, and all other interfaces are in the *failed* state. No tasks are required until an interface is set to the *unknown* state after a sufficient fall back period.

The registration behaviour of the mobile node depends on the current interface configuration and the replay protection mode in use. Timestamp-based replay protection gives more freedom for registration than nonce-based replay protection.

6.2.2 Initial registration scheme

The initial registration scheme is only applicable with timestamp-based replay protection. It can be employed in two interface configurations. One is that the mobile node has interfaces of *unknown* state, but no *active* interface. This is a typical case after a L2 handoff in the formerly *active* interface. A mobility binding should be acquired fast via the highest priority interface possible to resume the data flow. The other configuration is that there is an *active* interface and multiple higher priority interfaces become *unknown*.

The initial registration scheme sends RRQs simultaneously via all network interfaces in the *unknown* state. If an *active* interface is present, only the higher priority interfaces are qualified. The Identification fields should be assigned in the following manner: The RRQ of the lowest priority interface gets the lowest Identification value. The RRQs of the other interfaces get increasing Identification values according to their priorities. Consequently, the RRQ of the highest priority interface gets the largest Identification value.

Latencies of different routes vary, and the RRQs are probably reordered when they reach the home agent. The assignment of Identification values ensures that the mobility binding is formed via the highest priority interface, whose RRQ reaches the home agent. If a RRQ with a lower Identification arrives first, it also changes the mobility binding. The mobility binding is overridden by subsequent RRQs with larger Identification values. After the RRQ with the largest Identification value arrives, all the later RRQs are rejected by the home agent. Multiple rapid changes in the mobility binding are not harmful. Data packets are possibly routed

to multiple care-of addresses, but the situation settles down fast, as more and more lower Identification RRQs become rejected. No data packets are lost. The mobile node is assumed to be able to receive packets simultaneously from multiple interfaces (see Section 3.2).

The increase in the Identification fields can be done in fractional seconds (lower 32 bits) or to the NTP timestamp (higher 32 bits). The fractional seconds is more appropriate, because it has smaller effect to the time synchronization. If the home agent should employ rate limitation of received RRQs against the specification, the higher 32 bits must be used. In this case, the time synchronization set limits to the amount of simultaneous RRQs.

Initial burst of simultaneous RRQs produce an accepting or a rejecting RRP for each RRQ that reach the home agent. The mobile node may use this reachability information in next registration attempts. The initial registration scheme can not be used two times in a row. For the second batch of RRQs, the Identification values had to be larger than in the first batch. Therefore, a delayed RRQ of the first batch would be overridden by a RRQ of a lower priority interface of the second batch. This does not satisfy the criteria of using the highest priority interface that provides route to the home agent. In all, the batch of RRQs can only be sent once, after the *active* interface is lost or if several higher priority interfaces are changed to *unknown* state at the same time. Subsequent RRQs are sent using the basic registration scheme. With nonce-based replay protection, the basic registration scheme must be used from the beginning.

The initial registration scheme is fast compared to the method of sending one RRQ at a time, and retransmitting if the RRQ is lost. If the highest priority interface fails, the mobility binding is formed via the next interface in priority without any delay. The standard retransmission method would more than double the handoff latency in failure situations. At minimum, one second would have to be waited before the RRQ could be retransmitted. Considering the handoff latency, there are

no disadvantages in the initial registration scheme. Even if all the RRQs are lost, the fall back to the basic registration scheme ensures that the highest priority RRQ is retransmitted after one second as usual.

For a RRQ, the mobile node may receive a RRP with the error code 133 (identification mismatch) indicating that the time offset to the home agent time must be adjusted. This is not a problem in the initial registration scheme. Those RRQs that do not fit in the time synchronization generate RRPs indicating rejection, and the mobile node can register again. The sole harm is that the mobility binding may not be formed first via the best possible interface.

RFC 3344 [46] states that the home agent should write a log entry upon each identification mismatch. The initial registration scheme would therefore generate lots of log entries. If this is considered undesirable, the number of interfaces participating to simultaneous registrations can be limited only to a couple highest priority interfaces. This provides almost the same probability of a successful mobility binding as with all the available interfaces. After all, broken routes in the Internet are fairly rare. If the highest Identification RRQ is sent slightly after the another RRQ, it is more probable that the identification mismatch will not occur at all. Error log entries are by no means intolerable. Retransmitted RRQs with the minimum of one second interval may be reordered and inflict log entry as well.

6.2.3 Basic registration scheme

Nonce-based replay protection requires that each RRQ is answered by a RRP before further RRQs can be issued. Otherwise, the nonce synchronization between the home agent and the mobile node is broken. I will present a basic registration scheme for nonce-based replay protection and for timestamp-based replay protection after the one time use of the initial registration scheme. The basic registration scheme can be used in any interface configuration requiring registration actions.

Simultaneous registrations can not be issued in nonce-based replay protection. Nonce-based replay protection would not guarantee that the simultaneous registrations would follow the interface priorities. Multiple RRQs would have to have the same Identification field and the first to arrive at the home agent would change the mobility binding despite of the priority of the sending interface. All subsequent RRQs would be rejected. Answering RRPs would give valuable reachability information to the mobile node, but there are more important drawbacks. An interface might be set to a low priority, because the network layer handoff algorithm is anticipating to lose the L2 connection. Should this happen, handoff latency would be experienced once again. Unnecessary use of a low priority interface is questionable, even if it continued only for time of one round-trip to the home agent (supposing, the mobility binding is changed to a better interface as soon as the rejecting RRP is received). What is gained in possibly faster mobility binding is lost due to decreased QoS.

Table 3 defines several parameters that are needed in the basic registration scheme. Suggested values should not be taken as granted. The desired behaviour of the mobile node determines the absolute values that should be used.

The basic registration scheme is employed whenever there are *unknown* interfaces available for registration attempts. If an *active* interface exists, only *unknown* interfaces of higher priority are considered. If no *unknown* interfaces higher than the *active* interface exists, the mobile node can concentrate only on renewing the mobility binding. Registering via an *unknown* interface is called probing hereafter.

RRQ retransmissions of each interface must be slotted in time. Otherwise, the RRQs would interface with each other. The duration of each time slot is determined by a retransmission_period parameter. This period must be short in order to fast traverse through all *unknown* interfaces to quickly find a mobility binding. To comply with RFC 3344 [46] There must be one second interval between each time slot.

A short retransmission_period calls for another parameter. A RRP_tolerance de-

Table 3: Parameters for the basic registration scheme

Parameter	Description	Suggested value
Retransmission_period	Determines the length of the time slot of RRQ retransmissions for each interface. The retransmission_period must be short in order to fast cycle through all the interfaces.	3 seconds allows enough time for three RRQs to be transmitted with a binary back off retransmission algorithm.
RRP_tolerance	RRP for an interface should be accepted within the duration of RRP_tolerance, counting from the last RRQ the interface sent.	2 * maximum tolerated latency (default 7 seconds).
Calm_interval	To prevent flooding of foreign networks, an interface can start a new retransmission_period only after the calm_interval has elapsed counting from the last RRQ it sent.	15 seconds
Fail_count	The number of failed registration attempts allowed before an *unknown* interface is put to a *failed* state.	3 ensures that failure is not depicted because of temporary error conditions.
Fail_penalty	The required period of inactivity, after which a *failed* interface can be changed to the *unknown* state.	A dynamic value that should be increased upon each consequent confirmed registration failure.

termines the time interval, during which a RRP can be accepted at the latest to change the interface's state to *active*. The RRP_tolerance starts from the last RRQ of the time slot of the interface. This means that the RRP can be accepted even if it does not arrive within the retransmission_period of the interface i.e. it is already another's interface turn to try registering. The RRP_tolerance is a precaution ensuring that the mobility binding can also be acquired via routes that suffer from high latency. Now the total tolerated round trip latency to the home agent is the retransmission_period plus the RRP_tolerance. The basic registration scheme is illustrated in Figure 6.

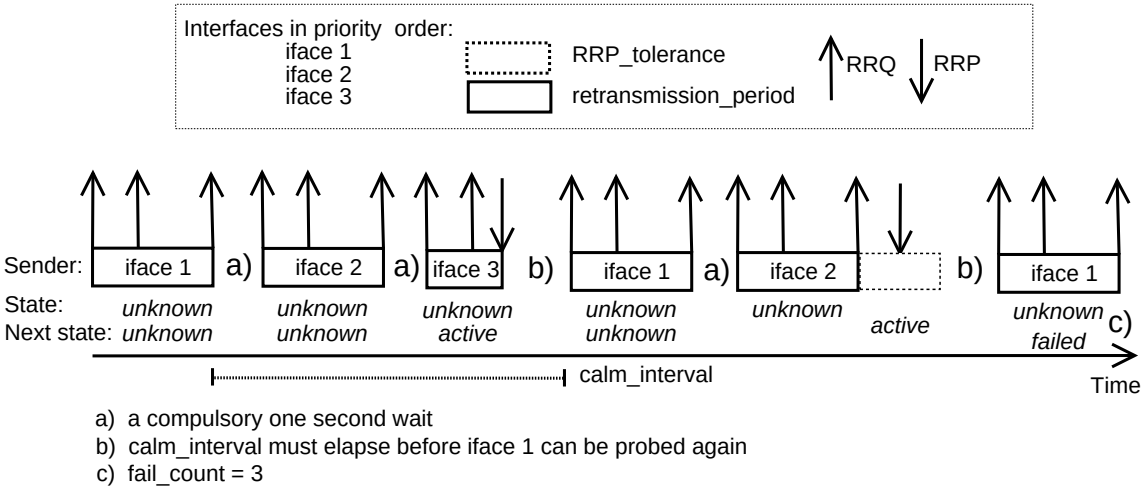

Figure 6: Basic registration scheme

The use of the RRP_tolerance has its consequences. A RRQ sent via a lower priority interface may change the mobility binding at the same time as a higher priority interface receives a RRP within its RRP_tolerance. Therefore, the RRPs to a lower priority interface than the *active* interface must be ignored. Moreover, the mobility binding must be renewed via the active interface immediately upon arrival of a lower priority RRP.

The probing order of the *unknown* interfaces is determined primarily by the recorded time of the last RRQ (last_RRQ_time) sent via each interface. The interface that has waited the longest time gets to register first. Each time an interface is added to the interface configuration, it gets "0 hours 1. January" as the last_RRQ_time.

If interfaces have the same last_RRQ_time, the probing order is determined by the interface priorities. The highest priority interface registers first. The resulting registration order prefers new interfaces, but maintains interface priorities otherwise. The idea behind preferring new interfaces is that once failed interfaces are more likely to fail again.

If an interface is probed unsuccessfully for fail_count times, it must be put to the *failed* state. A registration failure is accounted if no RRP arrives at the interface within the retransmission_period plus the RRP_tolerance. *Failed* interfaces must not be probed until they are put in the *unknown* state after a sufficient back-off period determined by the fail_penalty parameter. The fail_penalty should be increased upon each subsequent failure. The fail_penalty after the first registration failure should not be more than tens of seconds for IEEE-802.11 interfaces. The mobile node may have arrived at a hotspot. While the user enters credentials, the first registration attempt is likely to fail. If there is no mobility binding via any interface, the fail_penalty can be decreased to speed up the registration process.

The calm_interval determines the period of registration inactivity of an interface that is required before the interface is allowed to register again. That is, each time the fail_count is increased, the calm_interval is waited before the next registration. Without the calm_interval, the *unknown* interfaces would be fast cycled through multiple times and put to the *failed* state. The calm_interval prevents that the interfaces are not put to the *failed* state too hastily because of some temporary error condition. Additionally, the calm_interval suppress the flooding of the home agent with RRQs from multiple interfaces.

A visited subnet may contain multiple access routers. The router that is used for registering should be changed between each failed registration attempt i.e. each time the fail_count is increased. Likewise, if the registration latency via one router is increased drastically, forthcoming registrations should be done via another access router. See Section 7.2 for more information about recording registration latencies.

Time to time the *active* interface, if present, must register in order to renew the lifetime of the mobility binding. The *active* interface is privileged to register before any *unknown* interface. The time slot of the *active* interface is determined by the retransmission_period. After the mobility binding has been renewed, the basic registration scheme can continue normally. An UML activity diagram of the basic registration scheme is shown in Figure 7.

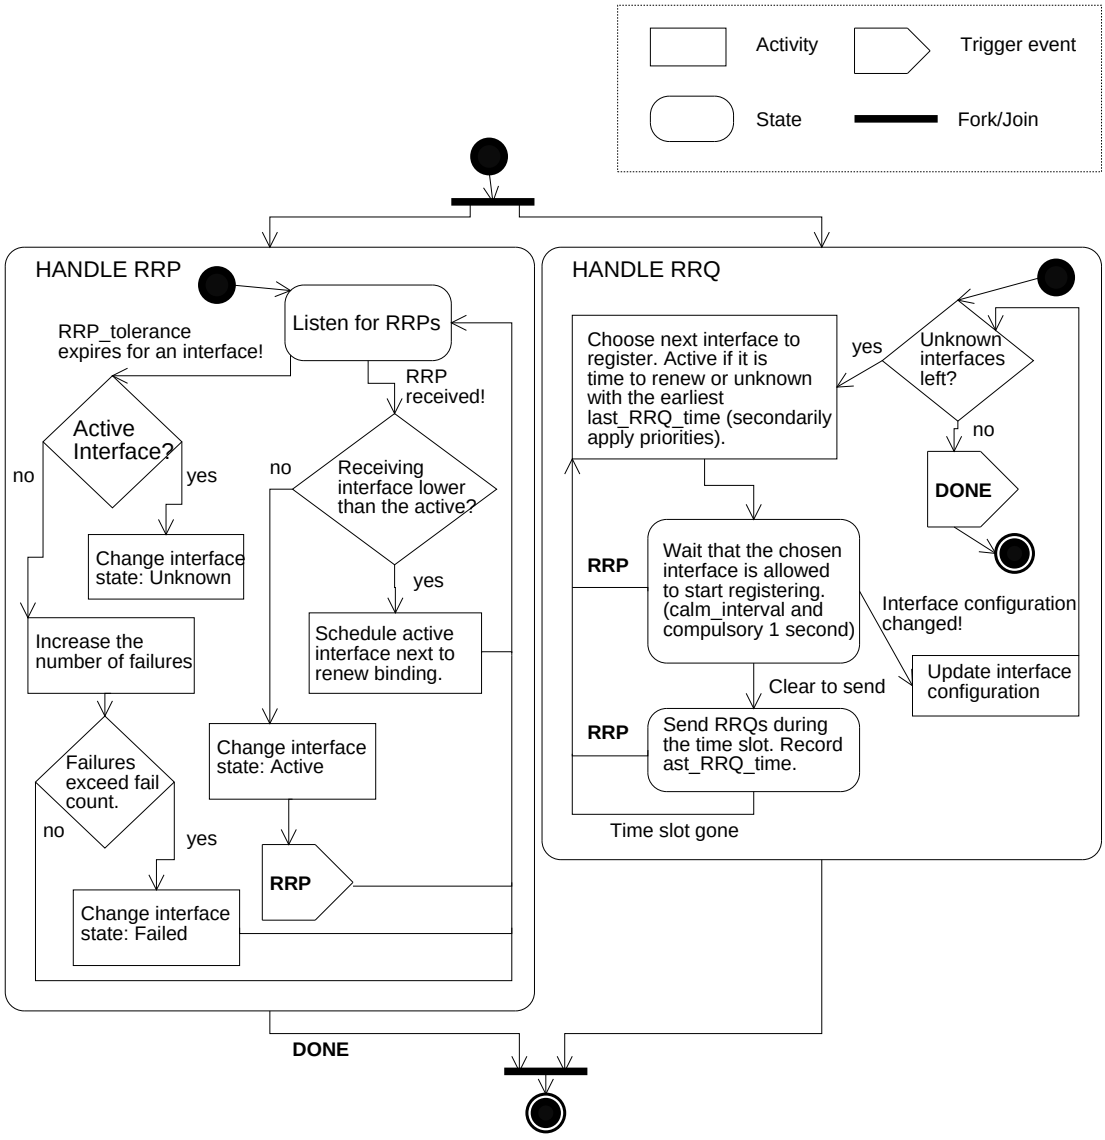

Figure 7: State machine of the basic registration scheme

The RRP_tolerance is applied also to the *active* interface. This makes it possible to probe other interfaces and gain mobility binding fast, in case the home agent is suddenly unreachable via the *active* interface. If *active* interface can not renew the

mobility binding within the RRP_tolerance, it must be put to the *unknown* state. The flow of data packets does not necessarily cease, even while the home agent is unreachable (uni-directional route). It is sensible to continue accepting data packets as long as it lasts.

The route to the home agent may be only uni-directional. A RRQ changes the mobility binding, but the RRP does not reach the mobile node. The same applies to the data traffic from corresponding nodes tunneled by the home agent. The unpleasant side effect is that the mobile node receives no indication about the change in the mobility binding. There is no remedy other than to register again via some working route in case a uni-directional route can be detected. One indication is a ceased data flow to the *active* interface after trying to register via some other interface. To prevent this harm, the *active* mobility binding should be renewed after all the *unknown* interfaces have been probed. Table 4 summarises the possible error scenarios with each type of registration scheme.

6.3 Alternatives to verify home agent reachability

The probing order of interfaces after a L2 handoff in the *active* interface would be optimized if the mobile node could verify the home agent reachability via lower priority interfaces without changing the mobility binding. The registration could be started with the interfaces that are known to provide a route to the home agent. There are several possibilities to investigate the home agent reachability. Unfortunately, all the methods have drawbacks and none of them can be used.

Simply pinging the home agent is possible, but as the ping is not authenticated, it should not be used.

A Registration Request with S bit (simultaneous binding) set would not change the *active* mobility binding. Unfortunately the data flow would be duplicated at least for the duration of one round-trip. This is not a good tradeoff, considering that

Table 4: Registration schemes and failure situations

Type of failure	**Initial registration scheme**	**Basic registration scheme**	
		with timestamps	**with nonces**
RRQ is lost	In case of the highest priority RRQ, the mobility binding is first formed via another interface. Remedied in the basic scheme.	RRQ is retransmitted.	RRQ is retransmitted.
RRP is lost	Data flow begins. The RRQ must be retransmitted only if the highest priority RRP was lost.	Data flow begins, but the registration must be repeated to get an accepting RRP.	Data flow begins, two more registrations are needed. The first to obtain the nonce synchronization, and the second to get an accepting RRP.
Identification mismatch	Success, if it was not the highest priority RRQ. Time offset adjusted.	Retransmission with adjusted time offset.	Retransmission with a new nonce.
Unidirectional route	No means to recover in case of the highest priority RRQ. Switch to the basic scheme.	Data flow to the *active* interface ceases. Remedied with a registration via the next interface in order.	Data flow to the *active* interface ceases. The nonce synchronization lost. Noticed and remedied by two RRQs sent via the next interface in order.

broken routes are rare in the Internet. For one optimized network layer handoff, data flow would be duplicated many times. More harmfully, the home agent may accept the RRQ without supporting the request for the simultaneous binding. This would change the mobility binding to the CoA of the lower priority interface.

Timestamp-based replay protection would allow the sending of two simultaneous RRQs. One from the *active* interface and one from a lower priority interface. Larger identification could be assigned to the RRQ of the *active* interface. Independent of

the receive order of the RRQs, the route to the home agent would be verified from the lower priority interface. The mobility binding would be changed only for the difference between the times of arrival of the RRQs, provided that the lower RRQ is received first. There is a problem if the highest Identification RRQ is lost or latency difference of the routes is considerable. A suboptimal mobility binding would be used for an extended duration.

Alternatively, the mobile node could sent only one RRQ from a lower interface. This RRQ should have an incorrect identification field. the home agent would reply with a RRP indicating identification mismatch. A possible error log entry at the home agent is the drawback in this method. Nonce synchronization would also be lost.

The care-of address of a RRQ could be "spoofed", meaning that a RRQ with the CoA (and source IP address) of the *active* mobility binding is issued via another interface. The answering RRP would arrive at the interface of the *active* binding. This would guarantee only a uni-directional route to the home agent from the interface that sent the RRQ. Spoofed RRQ would be topologically incorrect, and ingress filtering routers would drop it. This is a major drawback. Spoofed RRQs are also against the Mobile IP specification, and therefore this is a questionable method to use.

7 Heterogeneous handoff algorithm

In this chapter, I will present a heterogeneous handoff algorithm that comforms with the criteria in Section 3.1. Chapter 5 reviewed numerous handoff improvement methods. Unfortunately, their applicability is limited in the real environment. Operating systems and the current network adapters have their limitations. The absence of foreign agents in current networks drops many methods out of the consideration. The majority of the methods necessitated changes in the Mobile IP specification.

All the handoff improvement methods involving foreign agents had to be excluded. This includes foreign agent hierarchies for localized registrations, multicasting inside an administrative domain, and packet buffering and forwarding. RFC 3344 [46] specifies no micro mobility solutions that could be used for improving handoffs.

Transport layer enhancements to better cope in handoff situations are in no doubt beneficial. This study does not include them as they are out of scope. Enhancements to the TCP or UDP are not handoff management issues, even though they provide better overall quality of service.

Current operating systems do not implement means for applications to communicate their exact QoS needs to the handoff algorithm or to the protocol stack. Consequently, this information is not available for scheduling the handoffs or choosing between transfer media and target networks. Different user profiles could be created to determine the QoS needs. This would require the user to select a new profile whenever the networking surroundings or the application in use changes. I consider this is too cumbersome for the user.

Studies about handoff policies can be used to determine the best available network connection. There are difficulties in gathering the information to be compared, but the principle is applicable. People's tendency to move along routine paths [4] is a statistical method that leads to better handoff target decisions. I will use

policy-based method in the heterogeneous handoff algorithm to prioritize network connections.

The availability of the link layer information is very limited in practice and the information is unreliable. Implementing own link layer handoff algorithms for wireless LANs is not practical with the current IEEE-802.11 equipment. This makes it impossible to fully coordinate the actions of the heterogeneous handoff algorithm with the link layer handoff algorithms of each IEEE.802.11 network adapter. However, some link layer information can be used. The information provides the most powerful means to improve handoff management in the absence of foreign agents. The link layer issues are discussed in more detail in the following sections.

I do not consider using the Mobile IP simultaneous binding as an option to reach seamless handoffs. It is impossible to know when the application in use actually needs seamless connectivity. Each simultaneous binding doubles the data traffic. The bandwidth wasting is not justified considering that the home agent may reject the request for simultaneous binding. The new binding would override the old binding. This leads to unwanted consequences in the handoff management. Intermedia handoffs do not benefit from the simultaneous binding. They are already seamless.

7.1 Availability of link layer information

Link layer handoff algorithms for IEEE-802.11 are typically implemented in the firmware or device driver of the wireless network interface card. These algorithms are proprietary and they potentially differ substantially between devices. Currently there are no standards for L2 handoff algorithms. Most of the adapters make autonomous handoff decisions between access points and possibly inform upper the protocol stack about the decisions that were made. This leaves the heterogeneous network layer algorithm to choose only between available network interfaces without

ability to control the selection of access points to associate with. L2 connections have to be taken as granted. The availability of a certain association is subject to change at any moment. The resulting overall handoff behaviour is unavoidable suboptimal.

Operating systems provide interfaces to access and command IEEE-802.11 network interface cards (NICs). Windows systems have NDIS (Network Driver Interface Specification). Many Linux distributions have integrated *Wireless Tools* [56] developed by Jean Tourrilhes. Jean Tourrilhes has shown that adapters support mixed sets of L2 information. NDIS specification requires certain set of L2 support. Values for L2 parameters are set and queried using NDIS OIDs (Object Identifiers). Support for certain OIDs is obliged by Microsoft. According to my experiments, adapters sometimes meet these requirements by returning some constant, bogus value for the queried OID or, against the wishes of Microsoft, do not support the OID at all. In all, the L2 support implied by NDIS can not be trusted. In practice, implementations of the device driver and firmware are decided solely by the manufacturer and are subject to change between product releases.

Windows XP introduces an operating system specific problem. There is a service called *Wireless Zero Configuration* (WZC). The WZC has an important task in the management of IEEE-802.11 adapters. This service handles for example the setup of the adapter, when it is attached to the laptop for the first time. In addition, WZC supervises L2 handoffs through the NDIS interface. It is possible for the user to specify an access point to associate with. Furthermore, the WZC is exposed to problems arising from the incomplete NDIS interface support of IEEE-802.11 adapters. The WZC makes it impossible for a heterogeneous network layer handoff algorithm to implement its own NDIS-based L2 algorithms for each adapter. This would require the WZC to be turned off, which would waste important functionality. With the WZC operating, there would be two entities commanding the adapter without co-operation. It is doubtful that even the WZC works with all the NICs.

A NIC may execute its own L2 handoff algorithm despite the service. The exact functionality of the WZC is not published information.

The following set of L2 information is supported in the most of the current IEEE-802.11 adapters:

1. Basic Service Set Identifiers (BSSIDs) of the currently heard access points. BSSID is an access point equivalent to the MAC address of IEEE-802.3
2. Received Signal Strength Indication (RSSI) or signal quality (SQ) values of all the currently heard access points.
3. Media connect/disconnect event. This event is triggered when the network interface card has attained/lost the link layer connection to the access point. Applies also to IEEE-802.3 NICs

My heterogeneous handoff algorithm will use this set of L2 information to improve handoff management. It can not be trusted that the L2 information is valid with all the network adapters. There must be means to notice incorrect L2 information and to function without the L2 support.

7.1.1 Quality of link layer information

IEEE-802.11 adapters scan for new access points by listening beacon signals sent every 10 ms. This is called *passive* scanning. *Active* scanning includes sending of probe broadcast frames, which are answered by access points. A. Mitra et al. [34] have measured the AP discovery performance. If the NIC does not continuously update the list of passively heard APs prior to the reassociation, the active scan may introduce latency of several hundreds of milliseconds. Additionally, some IEEE-802.11 products are not able to transmit data frames during the process. For this reason, my algorithm will not request adapters to perform active scans. I will trust that NICs perform the scanning as efficiently as their implementation allows.

Because of this trade-off, it is not ensured that all the APs in the range can be noticed instantly.

Signal quality or RSSI values are meant to be used in a relative manner [60]. This makes it difficult to compare values between different NICs. Absolute values of RSSI do not reflect the real signal power level. Moreover, it is uncertain if the values are truly real time. NICs may update cached values using different intervals. The exact signal power oscillates with a considerable amplitude. NICs report moving averages of the real RSSI values. The averaging window is likely to vary between products. According to my experiences with Nokia C110, Orinoco WLAN PC card (5 volt), Cisco aironet 340, and D-link DWL-650 wireless LAN PCMCIA cards, the averaging window was smaller than one second. P. McCann [32] has studied the applicability of the L2 information of the IEEE-802.11 adapters for Mobile IPv6 handoffs and considered it unfit for real use where accuracy is necessary.

Section 5.4 presented prediction-based methods to estimate the time of losing the L2 connection by examining the past samples of RSSI values of the access point. This kind of prediction is useless with the current equipment. The sampling frequency is not adequate for just decisions, especially when the velocity of the mobile node is considerable.

There is another heavy reason for the heterogeneous algorithm not to control L2 handoffs. The user may have configured the NIC to associate with access points of certain Service Set Identifier (SSID). The SSID may identify one BSS or an Extended Service Set (ESS) containing multiple access points. If the ESS is a hotspot, the user will definitely want to maintain connections with the ESS after entering credentials and paying for the access.

All these matters limit the applicability of the L2 information. In the design of the algorithm, I will seek L2-information-based handoff improvement methods that are immune to the above-mentioned problems. The heterogeneous handoff algorithm

will be responsible for:

- Choosing the best interface/subnet to be used at time.
- Scheduling the switches between interfaces of different media to achieve best possible end user QoS.
- Detecting the movements (L2 handoffs) between subnets.
- Allocating the care-of address and detecting the access router in the foreign network.
- Registering possibly via multiple interfaces simultaneously in order to fast resume the data flow.

7.2 Architectural entities

Two architectural entities are needed to gather information for the use of the heterogeneous algorithm.

The *Scan log* periodically stores RSSI samples of the surrounding IEEE-802.11 access points. The purpose of the Scan log is to produce a rough estimate of the velocity of the mobile node. This helps the heterogeneous handoff algorithm to decide whether to use GPRS instead of WLAN. Each row of the Scan log contains one RSSI measurement. The fields are:

```
BSSID:RSSI:active:timestamp
```

RSSI values are recorded once per second. A row is written for each available access point or "BSSID". The "active" bit determines, whether the NIC has an association to the access point. By examining the Scan log, pathological ping-pong situations can be noticed.The entries older than five minutes should be purged. If the NIC continuously reports constant RSSI value, the Scan log should be disabled. A constant RSSI other than 0 indicates defective L2 information support. Similar conclusion

can be done if the NIC obtains a L2 connection, but it is not reflected to the RSSI values.

The *AR history* stores information about visited subnets, or to be more specific, used access routers. The reachability information and the latencies to the home agent are recorded. Using this information, the algorithm can prefer subnets and access routers that provide short latencies. In practice, it is difficult to obtain other network characteristics than the latency to the home agent. Throughput measurements are dependent on the amount of transferred data, but also on the capacity of the route from the corresponding node to the home agent (assuming reverse tunnelling). The change of a CN would affect to the throughput. The latency to the home agent has to be deduced from the registration messages. This is problematic. Variating congestion, load at the home agent , and possible dropping of UDP datagrams hamper the discovering of the true latency. Multiple registrations balance these variations.

By examining recent past latencies, the mobile node can estimate the registration latency when arriving at a new subnet. On a larger scale, the latency correlates with the geographic distance to the home agent. Subnets close to each other are likely to experience roughly similar latencies, since packets are likely to take routes that partially match with each other. The AR history contains multiple rows, which fields are:

`subnetID:AR_MAC:lat1:ts1:lat2:ts2:lat3:ts3:lat4:ts4:lat5:ts5`

The access router MAC (AR_MAC) address is followed by five the most recent latency samples and their corresponding timestamps. If the home agent was unreachable during the registration period, latency sample is set to infinite. The "subnetID" is a random id number for grouping access routers under the subnets they are reachable through. The registration schemes explained in Chapter 6.2.1 examine the AR history to register first via the lowest latency routers in the subnet. In case an AR is moved to a new location, the past latency samples become invalid. The error is

eliminated as new samples are generated. IEEE-802.11 introduces extra latency over IEEE-802.3. This latency is negligible if compared to the registration latency, and it is not taken into account. The entries older than one week should be expunged.

7.3 Move detection

Without foreign agents, the standard Mobile IP move detection method is not applicable. Instead, the mobile node uses L2 triggers to instantly notice changes in the L2 connectivity. Both IEEE-802.3 and IEEE-802.11 network adapters provide media connect and disconnect events. The IEEE-802.11 media connect event is triggered upon reassociation, possibly omitting the media disconnect event. In case the NIC does not support L2 triggers, there must be an alternative mechanism. The data flow itself is an adequate evidence of the L2 connection. In the absence of data flow, the mobile node should ARP [48] the access router to verify that it is reachable. I suggest that ARP requests should be sent every 600 ms. In case of no response, the subsequent ARP requests are sent every 200 ms. If the ARP response pends over 1000 milliseconds, a movement is deduced. The ARP-based move detection method is employed only for the interface that has the mobility binding. If the L2 triggers prove to be trustworthy, the ARP-based move detection can be stopped altogether. The use of ARP for move detection was suggested by H. Lewkowetz and S. Vaarala in [29] to be used in private address spaces that require NAT traversal.

7.4 Behaviour at foreign networks

The heterogeneous handoff algorithm acts whenever an interface obtains L2 connection to a foreign network. DHCP is currently the only realistic way to obtain CoA. The algorithm assumes that the DHCP client can be commanded interface-wise to lease and release the dynamic IP address in the visited subnet. When the lease is acquired, the DHCP client triggers the *coa_acquired* event for the algorithm. The

lease is renewed automatically if not told otherwise. The *coa_fail* event is triggered if the lease is expiring and it can not be renewed. Similarly, if the IP address lease is not granted at the first place, the coa_fail is triggered. IP leases should always be released if possible.

L2 handoffs within a subnet do not require a Mobile IP registration, at least in theory. Regretfully, DHCP poses limitations. Consider two interfaces connected to the same subnet, the first has a CoA, and the other has only a L2 association. Handoff to the latter interface could be done with gratuitous ARP, which would "hijack" the CoA from the first interface. Unfortunately, this would violate the DHCP state machine. The further renewals of the IP address lease would likely confuse the DHCP server, since the MAC address of the requesting interface would have been changed.

True optimization is possible with a IEEE-802.11 NIC making reassociations with access points connected to a single subnet. When the mobile node learns which access points connect to the same subnet, the same CoA can be used between L2 handoffs. Usually the mobile node is unaware of underlying subnet of an access point. In this case, the DHCP client can be triggered, although the interfaced subnet may be the same.

If the subnet remains the same, the access points executing IAPP or similar proprietary protocol will issue a Proxy ARP to update possible bridges connecting the subnet. As the bridges learn the MAC address of the mobile node, the data flow is resumed, and the mobile node can conclude that the L2 handoff has occurred within the subnet. In the absence of data traffic, a DHCP discover is broadcasted. Choosing the same DHCP server as prior to the L2 handoff, the mobile node contributes that the IP address would not change. For the same requesting MAC address, the DHCP server is likely to offer the same IP address as formerly. Because there has been lease left for that IP address, it can not be offered to another host meanwhile.

Regardless of the offered IP address, a Mobile IP registration is started after an IP lease is acquired. RFC3344 [46] base on globally routable IP addresses, but there is a work in progress to introduce NAT to the Mobile IP protocol. With NAT and private networks, two different subnets might offer exactly the same IP addresses for the use of the mobile node. The registration is needed, because the mobile node hasn't got knowledge of its IP address after the NAT box. There is no harm, even if the subnet was the same, and the registration is done in vain. The data flow is continuous, and by following the DHCP specification, the mobile node ensures that bridges of the subnet become updated with ARP.

To conclude, without better information the mobile node starts a DHCP process upon each L2 handoff. Formerly used DHCP server, if present, is preferred in order not to change the IP address (CoA) in case of intra-subnet L2 handoffs. In case of future NAT traversal, the Mobile IP registration is started always after the IP lease has been obtained. The result is a robust and simple functionality without real drawbacks. In practise, the mobile node may decide not to lease IP address for IEEE-802.11 NICs, if there is a stable fixed connection present at the same time. The handoff would cause somewhat longer interrupt when the user decides to move at the first time, but unnecessary IP address leases would be avoided.

GPRS is an exception to IEEE-802.11 and IEEE-802.3. The mobile node is not responsible to handoff management or leasing IP addresses. The host acquires IP address with PPP (Linux) or NDIS (Windows) from the GPRS card. The GPRS connection can be seen as a virtual link, that can be straightforwardly used when present. GPRS differs from other transfer media by its cost of bandwidth. For this reason I leave the choice of dialing GPRS to the user. Without active GPRS connection, the heterogeneous algorithm just leaves GPRS out of consideration. Usually, the GPRS PDP context remains in the active state for a long enough time for user not to be disturbed by the effort of periodically triggering the GPRS dial-up.

7.5 Velocity adaptation

The Scan log is used for rough velocity adaptation. This will remedy a pathologic situation, in which the mobile node is in rapid motion and tries to associate to IEEE-802.11 access points. If the Mobile IP registrations have not enough time to complete between the changing L2 connections, no connection at all can be acquired. Instead, GPRS should be used during the movement.

There is not much to be concluded from IEEE-802.11 signal measurements, which are the only way to perceive velocity with the current hardware. RSSI measurements are relative and do not reflect real received signal power. Different transmit power levels of access points and obstacles in the propagation path of the signal render the power measurements useless for determining the distance to the access point. With accurate and real-time open terrain power measurements, the mobile node would be able to deduce distances to access points and thus the velocity from the attenuation of the signal.

My velocity adaptation method goes around these problems with a cost of inaccuracy. It is based on multiple access point signal level measurements to achieve a statistical proof of the velocity. the absolute velocity is not measured. Instead the average duration of access point coverage will be used to determine if it is worth to try IEEE-802.11 connections.

The quantity representing the expected maximum time that the mobile node resides in the coverage area of an access point is called a *cell duration*. The cell duration is periodically determined in the following fashion:

Step 1. A portion of the Scan log is selected for examination. Parameter *adaptation_period* determines the examination period that is calculated from the end of the Scan log (the most recent values). A sufficient value for the adaptation_period is 30 seconds.

Step 2. All the BSSIDs occurring, i.e. having positive RSSI values, during the adaptation_period are considered.

Step 3. If no access points have been encountered during the period, the cell duration is set to undetermined. Abort.

Step 4. For each heard BSSID, the duration of continuous coverage is calculated. The coverage is considered to be continuous if there are positive RSSI values with maximum of 5 seconds gaps of zero RSSI. Note that one BSSID may produce multiple coverage durations in case of back and forth movement of the mobile node. Continuous coverage duration extending past the start of the adaption_period, is added to the measure. These are candidates for the cell duration.

Step 5a. If there is a previous, valid cell duration: The new cell duration is the maximum of the largest candidate and the previous cell duration. Abort.

Step 5b. In absence of previous cell duration: At least two access points in consideration must have been completely passed i.e. positive RSSI values do not extend to the beginning or to the end of the Scan log. This ensures that the access point has really been passed by. The largest of the candidates is chosen to be the cell duration.

The step 3 ensures that if there is no reinforcing information about the movement of the mobile node, the velocity adaptation is stopped after an adaptation_period amount of silence. Five second allowance for zero RSSI in the step 4 eliminates influence of temporary shadows that would result to a too short cell duration. The velocity adaptation should be triggered only in case of real movement. Step 5a ensures that the cell duration is accumulated in reaction to the hints that the movement has stopped. Prolonged coverage of any access point is a satisfactory hint. The velocity adaptation is triggered at the first time in the Step 5b. The duration of availability of a single access point varies by transmit the power level of the access

point and by the shape of the cell. Furthermore, the mobile node's trajectory of movement may pass through the border of the coverage area. All these issues would lead to a smaller duration of availability than with an access point in open terrain and the mobile node passing the circle shaped coverage area diagonally. For this reason, a sufficient amount of access point must be involved in the measurements. I suggest two, but if the reliability has to be increased, more access points should be required. On the other hand, if too many access points are required, the velocity adaptation is seldom triggered, even if it should. The probability of correct conclusion raises with the number of simultaneously monitored access points. It is enough that one of the access points is passed by close enough and that there is no substantial cornering effect. The relative velocity (cell duration) is concluded from the highest duration of availability of any of the access points. One way to improve the reliability of the cell duration is to require that a certain RSSI threshold is reached. For example 50 % of the maximum RSSI. This excludes the case that the mobile node enters the border of the coverage area and leaves immediately. Figure 8 illustrates the determination of the cell duration.

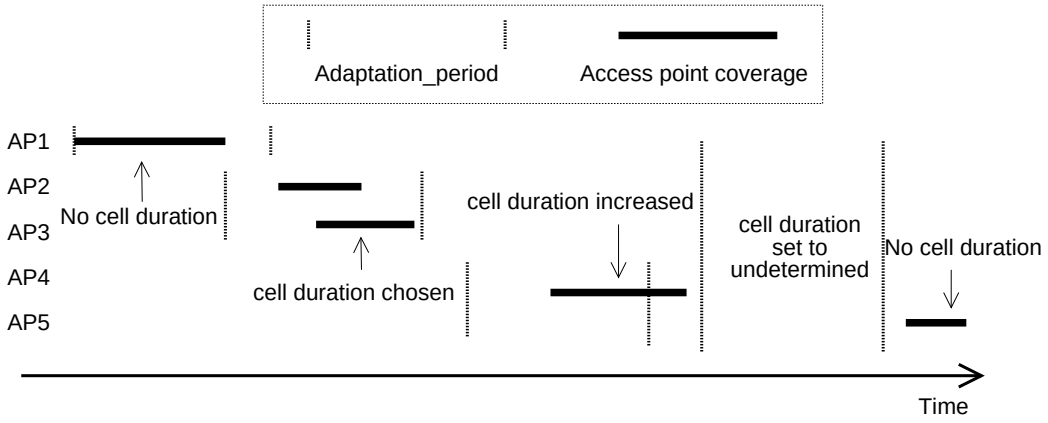

Figure 8: Cell duration

The cell duration gives the amount of time that a forthcoming access point can be expected to be available, provided that the velocity does not change. Increase in the speed of the mobile node diminishes the duration of coverage of all the encountered access points and thus the cell duration. The cell duration, expected registration

latency, and DHCP latency can be used for calculating the expected time of availability of subsequent access points. This is only an approximation also for another reason. IEEE-802.11 L2 handoffs can not be controlled. Efficient connection time is wasted if the adapter postpones the L2 handoff when a new access point becomes available. If the L2 handoffs could be controlled, the cell duration could be used for adjusting the L2 handoff strategy. With shortening cell duration, more recent access points should be associated, despite of the RSSI values at time.

It is unlikely that the L2 handoff algorithm of the network adapter would start the L2 handoff instantly if there is another access point available with greater RSSI. At least the threshold- and hysteresis-margin-based L2 handoff algorithms (see Section 5.4) would postpone the L2 handoff. However, instant L2 handoffs are assumed for robustness of the velocity adaptation method. *Expected communication time* at the next access point is:

Cell duration - registration latency - DHCP latency

The cell duration diminishes as the velocity increases. I define that if the expected overall handoff latency becomes 20 percent or greater of the cell duration, GPRS should be used instead of IEEE-802.11. This rule is called *interrupt-%*. Later, this rule is used also in other contexts. When the continuous access point coverage ends, the switch to GPRS is inevitable despite the velocity adaptation. I use 300 ms for the average DHCP latency (4 messages and their processing). Alternatively, the average DHCP latencies could be recorded in the AR history. The expected registration latency can be determined by using the AR history. The average of the last five successful registrations independent of the used access routers is calculated.

The interrupt-% is a simple, but not totally intuitive rule. With very short handoff latency and cell duration, the user experienced QoS with IEEE-802.11 could be better than with using GPRS instead. Frequent, but fast handoffs would be almost unnoticeable. However, there would be extensive registration traffic. This would waste bandwidth and burden the home agent. When both the cell duration and the

registration latency are quite large. The use of GPRS is justified for another reason. While the connection is good for the majority of time, rare handoffs introduce intolerably large latencies with the interrupt-% rule.

An example: The velocity of the mobile node is 20 m/s. The terrain is clear, and surrounding access points have coverage areas of 50 meters radius. The registration latency have been 700 ms on average. Consequently, the cell duration is 5 seconds, and each L2 handoff inflicts 1 second handoff latency not including the latency of the L2 handoff itself. The expected latency reaches 20 % of cell duration. Switch to GPRS is committed.

The velocity adaptation method remedies pathologic registration behaviour of a mobile node moving fast in near continuous WLAN coverage. The method is not perfect. The mobile node has to be in motion at least for an adaption_duration amount of time to acquire a usable cell duration value. Acceleration or braking can not be noticed. Even if the mobile node has stopped, WLAN is used only after the cell duration has been increased to the value satisfying the interrupt-% rule. RSSI of a single access point can not be used for detecting accelerations. A rapid increase in RSSI might indicate a substantial speed up or just a cornering effect. Likewise, RSSI values are steady if the mobile node moves along a circle-shaped trajectory around the access point. in this case, the velocity adaptation indicates no movement at all. The velocity adaptation method goes around problems involved in the quality of the L2 information. The cell duration is immune to the non real-time RSSI values. Errors in the beginning and at the end of the cell duration cancel each other. The same holds also for the different sized averaging windows of RSSI values.

The velocity adaptation should be aborted if wired connections are noticed.

7.6 Ping-Pong detection

Ping-pong is a situation, in which a IEEE-802.11 link layer handoff algorithm rapidly forms associations between multiple access points. The connection oscillates between access points causing frequent handoffs. Registration latencies lead to worsened QoS, and GPRS should be used instead.

The Scan log can be used for ping-pong detection. Since L2 handoffs can not be controlled, the handoff decisions of the IEEE-802.11 NIC have to be monitored. I somewhat extend the concept of the ping-pong: Frequent heterogeneous handoffs between GPRS and WLAN are alike unwanted. This happens when the coverage of WLAN is teetering. Each time WLAN becomes available, a registration is performed. Soon after, the signal fades and GPRS must be used instead. This causes yet another registration.

To detect a ping-pong situation, the Scan log is rolled back for a pp_detection_duration amount of time. 15 seconds provide that ping-pong is not concluded to hastily and that the priority of the WLAN interface can be raised fast enough, after the ping-pong situation is over. The ping-pong detection method is valid only if the mobile node is not in motion. If new access points become available, or average RSSI levels between the access points change, history does not guarantee that the ping-pong situation prevails. Therefore, ping-pong is considered to be over instantly when a new access point is heard. An access point is considered new, if it was not heard during the pp_detection_period. On the other hand, if a stationary mobile node experiences teetering WLAN coverage or suboptimal L2 handoff behaviour for a sufficient amount of time, it is reasonable to assume that the situation will continue.

As earlier, I use interrupt-% to determine the required communication time in proportion to the handoff interrupts. The efficient communication time is the whole 15 second pp_detection_period. The sum of handoff latencies must not exceed 20 % of the communication time. The handoff latency is determined as in the veloc-

ity adaptation method, the DHCP latency summed with the average registration latency found in the AR history. Handoff latency occurs upon every L2 handoff to a different subnet, provided that the mobile node has learned which access points connect to the same subnet. Likewise, the handoff latency is added whenever there is a hole in signal reception of the previously connected access point. Figure 9 clarifies the method. The interrupt-% is calculated by dividing the combined length of the gray area with pp_detection_period. Latencies involved in L2 reassociations are negligible.

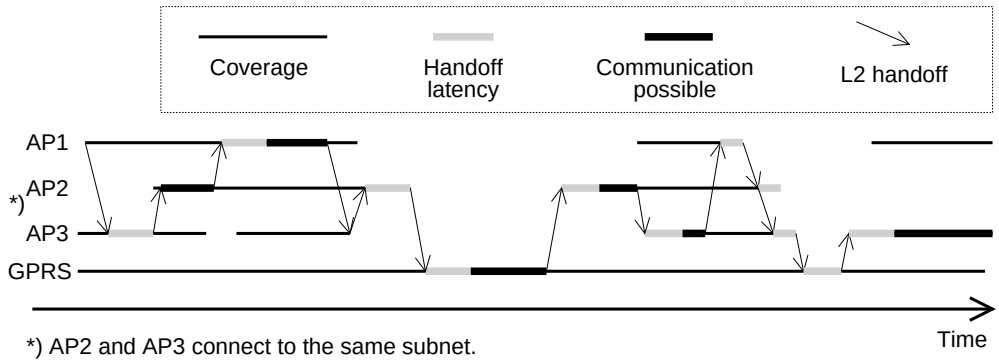

Figure 9: Ping-pong detection

The priority of the WLAN interface must not be raised to the normal level as soon as the interrupt-% becomes over 20 %. In borderline cases, this would inflict unnecessary frequent changes in the interface priorities. There must be a threshold. I define that the portion of interrupts must be lower than 10 %, before the ping-pong situation is considered to have settled down.

The ping-pong detection and velocity adaptation differ each other by various aspects. The velocity adaptation is not based on L2 handoff decisions of the IEEE-802.11 NIC. Instead, the cell duration tells if it is possible to communicate at all, even with optimal L2 handoffs. Moving mobile node can not make assumptions about the behaviour of the L2 handoff algorithm as the IEEE-802.11 surroundings change. For the ping-pong detection, the mobile node has to be stationary. The cell duration allows fast adaptation, when the movement has ceased. The cell duration is increased simultaneously, as the duration of coverage of any access point becomes longer. No

threshold is needed as in the ping-pong detection.

7.7 Design

The complete heterogeneous handoff algorithm is composed of the registration scheme state machine discussed in Section 6.2.1 and the events that control the interface priorities. The set of interfaces that participate in the registration scheme at a time is called the *registration set*. The registration set varies upon changes in the L2 connections and possessions of CoAs. An interface can be part of the registration set only if it has a valid CoA acquired in the interfaced network. Figure 10 presents the heterogeneous handoff algorithm as a UML activity diagram. The circled numbers in the right upper corner of each state or activity are used merely for explanation of the functionality.

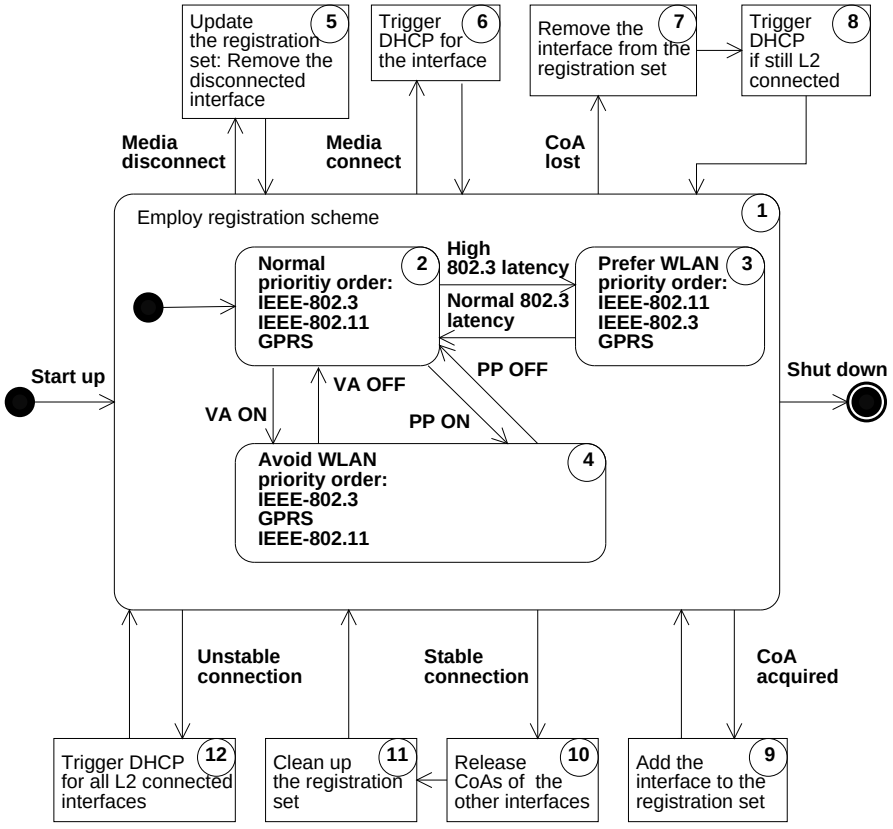

Figure 10: Heterogeneous handoff algorithm

Employ registration scheme (1) is a super state that corresponds for the initial

and basic registration schemes. The activity diagram does not show, whether the mobile node has an active mobility binding or not. The state means only that the applicable registration scheme at the time is followed with the interfaces in the current registration set. Section 6.2.1 showed that the registration schemes acquire mobility binding as fast as possible. Only a *shut down* event causes this state to be exited. In case there are no interfaces in the registration set, the algorithm just waits idle.

At the start up (2), the interface priorities are set in order that IEEE-802.3 is the most preferable, then IEEE-802.11 and the GPRS interface is the last one. This priority order is likely to provide the best QoS. Initially, no contrary information is available. The registration schemes pursue to arrange connection via the preferred interface. If this fails, other interfaces are tried in priority order.

The *media disconnect* event occurs when an interface loses its L2 connection or if GPRS is disconnected. This causes the interface to be removed from the registration set (5). The *media connect* event occurs if the IEEE-802.11 interface changes access point or if the IEEE-802.3 cable is connected. The DHCP client is triggered instantly to acquire a new CoA (6). In case of IEEE-802.11, the interface has to be removed from the registration set if the underlying subnet changed. The media connect and disconnect events are L2 triggers originating from the network adapters.

CoA acquired is an event, which the DHCP client sends to the algorithm. It indicates that an IP address lease has been acquired for an interface, and that the DHCP client will maintain the lease until further notice. The interface is appended to the registration set (9). The event occurs also when the GPRS interface is set up. Upon activation of the PDP context, the GPRS interface communicates the IP address to the upper protocol stack. The *CoA lost* event occurs if the DHCP client is unable to renew the IP address lease. This event occurs slightly before the actual expiration of the lease. The interface is removed from the registration set (7), but the other interfaces have enough time to register before the data flow ceases. If the interface is

still L2 connected, the DHCP client is triggered again to obtain another IP address lease (8). The CoA lost event occurs only if the event is unexpected. The algorithm may decide not to reserve CoAs for every interface and release some of the leases. This does not generate a CoA lost event. In the absence of L2 triggers, the CoA lost event occurs also if the ARP-based movement detection method concludes that the subnet has been changed.

the event *stable connection* leads to release of the IP address leases of the inactive interfaces (10). This is done in order not to unnecessarily consume IP address space. The cost is that without ready CoAs, the first heterogeneous handoff after the stable connection inflicts an interrupt to the data flow. The stable connection event is triggered if either the IEEE-802.3 or IEEE-802.11 interface has had an unbroken mobility binding for a 5 minutes. The *unstable connection* event occurs if the stable connection is interrupted by a L2 handoff or by the CoA lost event in the active interface.

Three sub-states inside the employ registration scheme state control the interface priorities. The IEEE-802.11 interface is avoided in the ping-pong and rapid movement situations (4). IEEE-802.11 is preferred over other interfaces, if the latency through the IEEE-802.3 interface is substantially higher (3). Several secondary events control the state transitions.

PP ON or "ping-pong detected" is an event generated by the ping-pong detection. The event is triggered if conditions for the interrupt-% are not met. That is, more than 20 % of the efficient communication time is wasted in handoff latencies. This locks the priority of the IEEE-802.11 interface to the lowest, preferring IEEE-802.3 and GPRS interfaces (4). *PP OFF* or "ping-pong ceased" event comes alike from the ping-pong detection. *VA ON* occurs if the velocity adaptation has concluded that the movement is too rapid for IEEE-802.11 connections. *VA OFF* is triggered upon increase in the cell duration so that the interrupt-% rule is satisfied again. Similarly, the media connect event in IEEE-802.3 interface ends the velocity adaptation. A

fixed connection proves that the mobile node has stopped moving.

Event *high 802.3 latency* adjusts the interface priorities (3) to prefer IEEE-802.11, if the IEEE-802.3 interface suffers from an unexpectedly high latency. The registration latency depends heavily on the used access router. All the access routers in the subnet have to be tried for registration, before the event is triggered. The experienced latency is compared to the latencies in the other recently visited subnets or to the past latencies in the same subnet. The AR history is used for latency comparisons. The latencies of the five past registrations are averaged for comparison. The latency difference should be substantial to overcome the better characteristics of IEEE-802.3 and the uncertainty involved in the latency measurements. I recommend the fixed connection must suffer from 300 milliseconds higher latency than expected, before the event is launched. In all, the high 802.3 latency is triggered if following conditions are met:

- The latency of the IEEE-802.3 interface proves to be 300 milliseconds higher than the expected latency.
- The IEEE-802.11 coverage has been uniform for the period of latency measurements. No L2 handoffs are allowed. Otherwise, it would not be clear that using WLAN would improve the QoS. RSSI level has to be over 50 % for the whole period. The IEEE-802.11 connection must be stable.
- The best access router in the subnet that is connected with the IEEE-802.3 interface is in use.

Event *normal 802.3 latency* normalizes interface priorities when it becomes evident that the latency in the subnet of the IEEE-802.3 interface is not abnormal. Several reasons trigger this event.

- The average of the latency samples in the IEEE-802.11 connected subnet turn out to be only 150 ms better than in the IEEE-802.3 connected subnet. The

150 ms threshold is to prevent frequent updating of the interface priorities in border-line cases.

- The association to the access point is lost or 25 % RSSI level is reached. Inferior signal strengths are not tolerated, because of increased transmit errors.
- IEEE-802.3 is connected to a new subnet.
- IEEE-802.11 has been used long enough, and latencies in the IEEE-802.3 connected subnet may have changed. I recommend one hour.

Latency measurements change interface priorities rather slowly. That is because sufficient amount of latency samples must be generated to smooth the variations. This makes the latency comparison an optimization that can be done only if the mobile node has been stationary for a sufficient amount of time. Registration frequency can be accelerated for faster operation.

Events that control the interface priorities and the registration set can be classified to the primary and secondary events. A primary event may inflict the triggering of a secondary event. For example media_connect event from a IEEE-802.3 interface triggers the VA OFF event. Each event updates either the registration set or interface priorities. If a primary event triggers a secondary event, both events should be processed before the changes are communicated to the registrations scheme. This ensures that the registration is not started with parameters that would be changed in the next instant.

8 Analysis of results

Developing a heterogeneous handoff algorithm for Mobile IP revealed to be a multifarious task. Extensive work have been committed by others (see Chapter 5) to improve the quality of service for a host that moves frequently in the Internet. Un-

fortunately, only a small amount of the improvements was applicable in the case where foreign agents are not available in general. I aimed to discover an algorithm that can be implemented on the current platforms and that is able to operate with widely available network interface hardware. The quality of the link layer information and triggers is currently the most decisive aspect that hampers efficient handoff management. It is not realistic to assume that L2 handoffs could be controlled by an external entity (see Section 7.1). I identified methods that can be used in improving handoff management in the situation assumed in Chapter 7. Within these limitations, the heterogeneous handoff algorithm was developed. It can be stated that the reached efficiency in handoff management is inadequate for real time applications. Sufficient QoS can not always be guaranteed. However, the algorithm improves overall QoS in many ways. Especially the situations that normally would lead to pathetic handoff management and unacceptable QoS, can be avoided with various methods I developed. Following sections evaluate the required criteria and discuss about the tradeoffs that were made. Many of the justifications of the design have been discussed in earlier chapters.

8.1 Minimal interference to data flow

The minimal interference to data flow is the most important goal in way to improve the user experienced QoS. Handoff latencies cause interrupts in the data flow. Likewise, it is important to use the best available link layer transfer media at a time. Independent of the transfer media, perceived QoS is also dependent on the characteristics of the route to the home agent and to the corresponding nodes. The routes may differ substantially between the network interfaces.

Inter-media handoffs are always seamless, provided that the connection via the previous interface stays up until the new mobility binding is ready. This originates from the facts that the mobile node can receive packets simultaneously via multiple interfaces and that the home agent starts to forward traffic immediately after

accepting a Registration Request. Difficulties arise when the connection is lost in the source interface, and the handoff must be committed reactively. Intra-media handoff caused by a reassociation of IEEE-802.11 to a different subnet similarly inflicts interrupt to the data flow. This is impossible to avoid without special foreign agents to forward traffic or foreign agent hierarchies to localize the registration. I focused to make handoffs faster to minimize the damage. Handoff latency consists of the L2 handoff latency, registration latency, and the DHCP latency for acquiring the CoA. To minimize the DHCP latency, the inactive interfaces should also acquire IP address leases. These leases are released if the mobility binding is considered stable enough i.e. future handoffs are unlikely. The use of link layer triggers allows the mobile node to instantly notice L2 handoffs between subnets. This shortens the handoff latency usually by several seconds compared to the standard mobile IP move detection methods.

In Chapter 6, I showed that forming a mobility binding via multiple, prioritized interfaces is not trivial. I developed the initial registration scheme and the basic registration scheme, which shorten the registration latency substantially in case that the registration would fail via the most preferred interface. With timestamp-based replay protection, the registration latency is also shortened by at least one round trip to the home agent in case the Registration Request of the preferred interface would be lost. The registration schemes schedule the registrations of multiple interfaces to fast reach the mobility binding, but also to avoid flooding the home agent. The registration schemes take care that the mobility binding is efficiently formed via the highest priority interface and the lowest latency route. Without explicit registration scheme, the mobile node would behave catastrophically if the subnet of the preferred interface would not provide route to the home agent at all. The worst case is that there would be no mobility binding at all until leaving the subnet.

The velocity adaptation and ping-pong detection methods both remedy the situation that the IEEE-802.11 interface would not reach sufficient connection under any

circumstances. GPRS is raised in the priority order. From the QoS point of view this is a considerable optimization. Otherwise, the mobile node would be locked in a state that no connection is available as the registration is continually tried and failed using the IEEE-802.11 interface. I defined that the connection is unusable if registration latencies will exceed 20 % of the efficient communication time. With very low registration latencies, this rule prevents excessive registration traffic, even while the overall connection might be usable. For applications that communicate with TCP, even small but frequent interrupts would inflict excessive retransmissions. Therefore, the interrupt-% rule is applicable independent of the length of the handoff latency.

The velocity adaptation method (see Section 7.5) is based on the signal strength measurements of the surrounding access points. The method is coarse because of the quality of the L2 information. The signal strength information is not real-time with all the IEEE-802.11 adapters. Relative signal strength values, varying transmit powers of access points, and asymmetric coverage areas prevent distance measurements to the access points. These matters make the velocity adaptation difficult. My velocity adaptation method goes around problems of L2 information (see Section 7.1). It still has its limitations. In certain WLAN surroundings, high velocity may be concluded falsely. The probability of errors can be diminished by claiming more evidence of the high velocity i.e. by increasing the required number of access points that are quickly passed by. On the other hand, this postpones the conclusion of the high velocity. The velocity adaptation adjusts interface priorities to avoid registrations via the IEEE-802.11 interface, if the velocity is too high.

The ping-pong detection (see Section 7.6) method suffers from the same kind of difficulties as the velocity adaptation. Sufficient amount of evidence must be acquired, before conclusions can be made. This postpones the updating of the interface priorities. Speeding up the process would lead to increasing number of false decisions. It is reasonable to be absolutely sure that WLAN is unusable, before GPRS is raised

in the priority order. The bandwidth of GPRS is order of magnitudes lower than IEEE-802.11. For many user applications this is the same that the connection would be lost.

The velocity adaptation and ping-pong detection involve many parameters. I have suggested values for these parameters. The recommended values are by no means optimal. The proof of that would necessitate the knowledge of the characteristics of the average IEEE-802.11 surroundings that the mobile node will encounter. One possibility to fix this is to dynamically adjust the parameters according to the experienced surroundings.

The AR history enables latency-based policy for the handoff target network selection. The latency in the current subnet is compared to the expected latency calculated using the data in the AR history. Currently, the average registration latency is the only QoS parameter that is measured. However, it is not used alone in the handoff target selection. The characteristics of the link layer transfer media are also taken into account. In practice, this means selecting between the IEEE-802.3 and IEEE-802.11 interfaces. The wireless connection has to be stable, interrupts in the data flow would make comparison of latencies irrelevant. Furthermore, the AR history provides an optimization if a subnet is visited multiple times. The access router that provided the lowest latency to the home agent can be chosen at the first place. Networks that provided no route to the home agent can be avoided pro-actively.

8.2 Detailed algorithm design

I have presented the heterogeneous handoff algorithm and the involved external entities in detail that should give sufficient guidelines for the implementation. Parameters that determine the behavior of the handoff management were explicitly stated and the recommended values were given. Events that control the interface priorities and the registration set were explained by their origins and by the ac-

tions they trigger. State machines of the heterogeneous handoff algorithm and the registration schemes were presented to give the reader a clear picture of the functionality. Each design choice was justified as they were presented, and no noteworthy ambiguities should remain.

8.3 Implementational feasibility

This goal was one of the most determining factors of the design choices that were made. The topmost idea was to develop an algorithm that can really be implemented and that operates in current networks as they are. This implies the absence of foreign agents, operation system considerations, and adaptation to the limitations of the current network adapters and their operating system interfaces. The resulting algorithm makes no unrealistic demands for neither the operating system nor networking hardware. These issues where covered in Chapter 7. Several handoff improvement methods had to be dropped out of the design for they would be extremely hard or impossible to implement. The heterogeneous handoff algorithm can be readily implemented and deployed in mobile nodes. The goal for the implementational feasibility can be considered reached.

Controlling the L2 associations of IEEE-802.11 proved to be unrealistic with current network adapters employing their own L2 algorithms (see Section 7.1). Windows XP operating system's Wireless Zero Configuration service would make the L2 handoff behaviour even more unpredictable. Scanning of access points and their corresponding RSSI levels is often time consuming. In addition, the real-timeness of this information is not guaranteed. The quality of L2 information depends heavily on the implementation of the device driver and the firmware of the network interface card. This excludes the handoff prediction methods that try to estimate the future RSSI from the past values. There is also another reason for not to predict L2 handoffs of IEEE-802.11. Characteristics of the IEEE-802.11 signal modulation provide sufficient data rates even with very modest signal strengths. After the turn-

ing point, the error probability increases with quickening pace and the connection deteriorates fast. The fine line of a feasible connection and no connection at all makes it difficult to make accurate predictions whether the connection is lost or not. Usually, an alternative connection must be sought via GPRS. For bandwidth and latency sensitive applications, this is effectively the same that the connection would be lost. Therefore, the costs of a false prediction are too severe to overcome. It is better not to predict, if it can not be done accurately and reliably.

Knowing the QoS requirements of the applications used at a time would be quite beneficial. Handoffs could be scheduled to the moments that the interrupts inflict minimum harm. The interfaces could be prioritized more precisely according to the required bandwidth and maximum latency. Currently this kind of information is impossible to obtain. The resulting heterogeneous algorithm is suboptimal in this respect.

8.4 Conformance to Mobile IP

The heterogeneous handoff algorithm and handoff management issues presented within this thesis fully conform with the Mobile IP specification. No new protocol messages or deviant use of the messages were introduced. The ARP-based move detection method is a supplement to the protocol. It is a link layer method that is not part of the RFC 3344 [46].

The absence of the foreign agents was the main design guideline for the handoff management. The intention was to enable as wide as possible operating region for the mobile node. Foreign agents are scarce in current networks. However, the resulting algorithm has no features, that were inapplicable when registering via a foreign agent. The initial and basic registration schemes do not pay attention, whether the registrations are done via foreign agents or not. Whenever a foreing agent is present, the ARP-based move detection method can be substituted with

the standard Mobile IP move detection that is based on Agent Advertisements. The ping-pong detection and velocity adaptation use IEEE-802.11 signal levels to make conclusions. They are not dependent on foreign agents whatsoever.

8.5 Efficient use of bandwidth

This criteria advocates that the mobile node should use network resources sparingly. The basic registration scheme limits the number of registrations (see Section 6.2.1). Each registration failure diminishes the rate of the future registration attempts via the interface. The interface is changed fast if the registration pends. This limits the registration traffic. The ARP-based move detection wastes bandwidth. It can be partially justified, because ARP traffic is limited in the subnet. If greater than one second move detection latencies are tolerated, the sending of ARP requests can be done more seldom. If the mobile node decides that L2 triggers can be trusted i.e. they always occur before the move detection would notice the movement, the ARP based move detection is unnecessary. One way to spare network resources is to release DHCP leases of the inactive interfaces. This is done after the current connection is seen stable enough. This is a compromise between wasting the IP address space and capability of faster handoffs. The use of Mobile IP simultaneous bindings was rejected. It duplicates traffic all the way from the home agent to the mobile node. Inter-media handoffs are seamless anyways. If the simultaneous binding is not supported by the home agent, the mobility binding is changed to the new CoA, whether it is wanted or not. The active scanning of access points would consume bandwidth of the air interface. However, it was left out of the design for its inefficiency and because the implementations of some IEEE-802.11 NICs prevent data transmit during the scanning. It is up to the NIC to determine how access points should be scanned.

8.6 Minimal user interaction

The heterogeneous handoff algorithm requires no actions from the user other to occasionally trigger the GPRS dial-up. The decision to use GPRS is left to the user, because of the high cost involved in the use of the media. This user interaction has modest impact to the usability. Once dialed, GPRS remains active for a long time. There are tradeoffs in freeing the user from configuring the parameters of the heterogeneous handoff algorithm. More intelligent handoff decisions between interfaces could be committed with precise QoS requirements. The required QoS depends heavily on the applications that are in use. For example voice-over-IP does not tolerate high latency, whereas file transfer is mainly sensitive to throughput. The most important QoS parameters that the user might determine are the maximum tolerated latency, the maximum length of interruption in the data flow, and the required bandwidth. In addition, the user might be given control over which type of interfaces should be preferred. These parameters could be collected to different user profiles. The registration schemes, velocity adaptation, inter-media handoff decisions, move detection, and ping-pong detection would all reach more optimal results with clearly defined QoS requirements. I chose not to disturb the user with these issues. The usability of the software suffers too much, if the user has to be aware of the requirements of the applications in use and of the surrounding networking environments.

8.7 Security

The security of the mobile IP is not weakened by the heterogeneous handoff algorithm, since the protocol messages are not used in questionable ways. Spoofed Registration Requests (see Section 6.2.1) might have had security impacts, but they were rejected.

8.8 Future work

The heterogeneous handoff algorithm developed within this thesis has a limited lifetime. The emergence of the IPv6 networks will render part of the work obsolete. However, current forecasts indicate that the algorithm has its place for several forthcoming years.

Future developments in IEEE-802.11 network interface cards may allow external entities to have complete control over the link layer handoffs. Similarly, more sophisticated and accurate link layer information may become available. This would give more possibilities for the handoff management. The heterogeneous handoff algorithm could be extended to start registrations via another interfaces, before losing the link layer connection of the active interface. Seamless handoffs would be achieved. Ping-pong situations could be avoided, and more precise velocity adaptation would become possible. Link layer handoffs between access points in the same subnet could be preferred to avoid latency involved in the Mobile IP registration. The L2 handoff management could be combined with the upper layer heterogeneous handoff algorithm. This includes the prioritization of available access points by their received signal strengths. Velocity adaptation could be used to change the L2 handoff strategy for fast moving mobile nodes. More eager associations to newly heard access points would allow longer connection times and thus better QoS.

The use of the AR history and the Scan log can be extended. By analysing the past movements and the handoff decisions made, the parameters involved in the handoff management can be adjusted to fit better to the profile of the particular mobile node. People often move along routine paths and operate in familiar networking surroundings. This is the case for example when travelling between work and home. More optimal handoff decisions are possible by predicting forthcoming connection possibilities by matching the movement pattern to earlier movements.

I used average latencies to the home agent to compare connections via otherwise

equal interfaces. Other criteria can be added to the cost function determining the interface priorities. An external module to observe TCP retransmissions would provide information about the packet drop in the route. Monitoring of the congestion control mechanism would add information about throughput. With these improvements, the cost function would be more reliable and improve the selection of the interface.

9 Conclusions

I developed a handoff algorithm for Mobile IP that uses IEEE-802.3, IEEE-802.11, and GPRS media to access networks. The goal was to invent an algorithm that can be readily taken into use. Therefore, the functionality was limited by characteristics of the available network interface cards and their operating system interfaces. Foreign agents are scarce in the current networks. This excluded micro mobility solutions that would have improved the handoff performance. The pursued goals were mainly reached. The heterogeneous handoff algorithm uses link layer information for fast movement detection, velocity adaptation, and ping-pong detection. The algorithm is able to prioritize network connections available at a time to improve overall quality of service. The criteria for the prioritization includes latencies to the home agent and the characteristics of the transfer media. The velocity adaptation and ping-pong detection methods remedy frustrating situations of trying to use IEEE-802.11 connections that are working poorly. Hopeless situations can be detected, and GPRS can used instead. Intra-media handoffs with IEEE-802.11 still cause interrupts to the data flow. To compensate this, the algorithm succeeds to lower the number of handoffs on average.

I identified several problems involved in simultaneous Mobile IP registrations with timestamp- and nonce-based replay protection modes. Based on the findings, I developed a registration scheme that allows efficient registering through multiple

interfaces having dynamic priority order. The registration scheme lowers the registration latency, especially in the situations, in which the highest priority interface fails to register and an alternative mobility binding has to be established fast.

The handoff management scheme uses DHCP to obtain care-of addresses and an ARP-based movement detection method as a back up method for link layer triggers. The ability to function without foreign agents is crucial. However, the heterogeneous handoff algorithm and the registration scheme are also applicable when registering via foreign agents.

References

[1] B. Adoba and M. Beadles. RFC 2486: The Network Access Identifier. Internet Engineering Task Force, January 1999.

[2] N. An, Y. F. Hu, R. E. Sheriff, and J. G. Gardiner. Development and Implementation of a Handover Algorithm in Heterogeneous Mobile Environments. In *Communication Technology Proceedings, 2000. WCC - ICCT 2000. International Conference on*, volume 1, pages 543–546, 2000.

[3] H. Balakrishnan, S. Seshan, and R. H. Katz. Improving reliable transport and handoff performance in cellular wireless networks. *ACM Wireless Networks*, 1(4), 1995.

[4] V. Bharghavan and M. Jayanth. Profile-based next-cell prediction in indoor wireless LANs. In *Proceedings of IEEE Singapore International Conference on Networking*, April 1997.

[5] S. O. Bradner. RFC 2026, BCP 9: The Internet Standards Process – Revision 3. Internet Engineering Task Force, October 1996.

[6] K. Brown and S. Singh. M-udp: Udp for mobile cellular networks. *ACM SIGCOMM Computer Communication Review*, 26(5):60–78, 1996.

[7] R. Cáceres and V. N. Padmanabhan. Fast and scalable wireless handoffs in support of mobile internet audio. *Mobile Networks and Applications*, 3(4):351–363, 1998.

[8] P. Calhoun and C. Perkins. RFC 2794: Mobile IP Network Access Identifier Extension for IPv4. Internet Engineering Task Force, March 2000.

[9] A. Campbell, J. Gomez, C-Y Wan, S. Kim, Z. Turanyi, and A. Valko. draft-ietf-mobileip-cellularip-00.txt: Cellular IP, January 2000.

[10] J. A. Cobb and P. Agrawal. Congestion or Corruption? A Strategy for Efficient Wireless TCP Sessions. *IEEE Symposium on Computers and Communications*, pages 262–268, 1995.

[11] D. E. Comer. *Internetworking With TCP/IP Volume I: Principles, Protocols, and Architecture*. Prentice-Hall, Inc., third edition, 1995. (ISBN 0-13-227836-7).

[12] J. D. Day and H. Zimmerman. The OSI Reference Model. In *Proceedings of the IEEE*, volume 71, pages 1334–1340, December 1983.

[13] S. Deering. RFC 1256: ICMP Router Discovery Messages. Internet Engineering Task Force, September 1991.

[14] G. Dommety and K. Leung. RFC 3025: Mobile IP Vendor/Organization-Specific Extensions. Internet Engineering Task Force, April 2001.

[15] G. Dommety, A. Yegin, C. Perkins, G. Tsirtsis, K. El-Malki, and M. Khalil. draft-ietf-mobileip-fast-mipv6-03.txt: Fast Handovers for Mobile IPv6, July 2001.

[16] R. Droms. RFC 2131: Dynamic Host Configuration Protocol. Internet Engineering Task Force, March 1997.

[17] P. Ferguson. RFC 2267: Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing. Internet Engineering Task Force, May 2000.

[18] A. Fladenmuller and R. De Silva. The effect of mobile IP handoffs on the performance of TCP. *Mobile Networks and Applications*, 4(2):131–135, 1999.

[19] M. Gudmundson. Correlation model for shadow fading in mobile radio systems. *Electronic Letters*, 27(23):2145–2146, May 1991.

[20] E. Gustafsson, A. Jonsson, and C. Perkins. draft-ietf-mobileip-reg-tunnel-07.txt: Mobile IPv4 Regional Registration, October 2002.

[21] J. M. Holtzman and A. Sampath. Adaptive averaging methodology for handoffs in cellular systems. *IEEE Transactions on Vehicular Technology*, 44(1):56–66, February 1995.

[22] J. Hu, L. K. Yeung, S. C.Kheong, and G. Feng. Hierarchical Cache Design for Enhancing TCP over Heterogeneous Networks with Wired and Wireless Links. In *GlobeCom 2000*, pages 338–343, San Francisco, November 2000.

[23] V. Kapoor, G. Edwards, and R. Sankar. Handoff Criteria for Personal Communication Networks. In *IEEE International Conference on Communications*, volume 1, pages 1297–1301, May 1994.

[24] J. Kempf and J. Wood. Analysis and Comparison of Handoff Algorithms for Mobile IPv4. In *Proceedings of the 52nd IETF Meeting*, Salt Lake City, Utah USA, December 2001.

[25] S. Kent and R. Atkinson. RFC 2401: Security Architecture for the Internet Protocol. Internet Engineering Task Force, November 1998.

[26] S. Kerry, D. Bagby, and B. O'Hara. Recommended Practice for Multi-Vendor Access Point Interoperability via an Inter-Access Point Protocol Across Distribution Systems Supporting IEEE 802.11 Operation. IEEE Std 802.11f/D3, DRAFT, January 2002.

[27] Y. Kinoshita and T. Itoh. Performance analysis of a new fuzzy handoff algorithm by an indoor propagation simulator. In *Proceedings of 43rd IEEE VTC*, pages 241–345, 1993.

[28] H. Krawczyk, M. Bellare, and R. Canetti. RFC 2104: HMAC: Keyed-Hashing for Message Authentication. Internet Engineering Task Force, February 1997.

[29] H. Lewkowetz and S. Vaarala. draft-ietf-mobileip-nat-traversal-07.txt: Mobile IP NAT/NAPT Traversal using UDP Tunnelling, November 2002.

[30] K. E. Malki, P. R. Calhoun, T. Hiller, J. Kempf, P. J. McCann, A. Singh, H. Soliman, and S. Thalanany. draft-ietf-mobileip-lowlatency-handoffs-v4-04.txt: Low Latency Handoffs in Mobile IPv4, June 2002.

[31] H. Maturino-Lozoya, D. Munoz-Rodriguez, and H. Tawfik. Pattern recognition techniques in handoff and service area determination. In *Proceedings of 44th IEEE VTC*, volume 1, pages 96–100, 1994.

[32] P. McCann. draft-mccann-mobileip-80211fh-01.txt: Mobile IPv6 Fast Handovers for 802.11 Networks, October 2002.

[33] D. L. Mills. RFC 1305: Network Time Protocol (Version 3) - Specification, Implementation and Analysis. Internet Engineering Task Force, March 1992.

[34] A. Mitra, M. Shin, and W. Arbaugh. An Empirical Analysis of the IEEE 802.11 MAC Layer Handoff Process. CS-TR-4395, University of Maryland Department of Computer Science, September 2002.

[35] M. Miyoshi, M. Sugano, and M. Murata. Performance evaluation of tcp throughput on wireless cellular networks. In *Proceedings of the IEEE Vehicular Technology Conference 2001 Spring (VTC2001 Spring)*, Rhodes, Greece, May 2001.

[36] G. Montenegro. RFC 2344: Reverse Tunneling for Mobile IP. Internet Engineering Task Force, May 1998.

[37] G. Montenegro and V. Gupta. RFC 2356: Sun's SKIP Firewall Traversal for Mobile IP. Internet Engineering Task Force, June 1998.

[38] D. Munoz-Rodriguez, J. Moreno-Cadenas, M. Ruiz-Sanchez, and F. Gomez-Casaneda. Neural supported handoff methodology in microcellular systems. In *Proceedings of 42nd IEEE VTC*, volume 1, pages 431–334, 1992.

[39] S. Hanks nad T. Li, D. Farinacci, and P. Traina. RFC 1701: Generic Routing Encapsulation (GRE). Internet Engineering Task Force, October 1994.

[40] K. Pahlavan, P. Krishnamurthy, A. Hatami, M. Ylianttila, J. Mäkelä, R. Pichna, and J. Vallström. Handoff in Hybrid Mobile Data Networks. *IEEE Personal Communications Magazine*, 7(2):34–47, April 2000.

[41] J. Pan, J. W. Mark, and X. Shen. TCP performance and its improvement over wireless links. In *GlobeCom 2000*, volume 1, pages 62–66, San Francisco, November 2000.

[42] C. Perkins. RFC 2002: IP Mobility Support. Internet Engineering Task Force, October 1996.

[43] C. Perkins. RFC 2003: IP Encapsulation within IP. Internet Engineering Task Force, October 1996.

[44] C. Perkins. RFC 2004: Minimal Encapsulation within IP. Internet Engineering Task Force, October 1996.

[45] C. Perkins. RFC 3220: IP Mobility Support for IPv4. Internet Engineering Task Force, January 2002.

[46] C. Perkins. RFC 3344: IP Mobility Support for IPv4. Internet Engineering Task Force, August 2002.

[47] C. Perkins and P. Calhoun. RFC 3012: Mobile IPv4 Challenge/Response Extensions. Internet Engineering Task Force, November 2000.

[48] D. C. Plummer. RFC 1305: An Ethernet Address Resolution Protocol. Internet Engineering Task Force, November 1982.

[49] G. Pollini. Trends in handover design. *IEEE Communications Magazine*, 34(3):80–90, March 1996.

[50] J. Postel. RFC 768: User Datagram Protocol. USC/Information Sciences Institute, August 1980.

[51] J. Postel. RFC 793: Transmission Control Protocol. USC/Information Sciences Institute, September 1981.

[52] R. Ramjee, T. La Porta, S. Thuel, K. Varadhan, and L. Salgarelli. draft-ietf-mobileip-hawaii-01.txt: IP micro-mobility support using HAWAII, July 2000.

[53] S. Seshan. Low-latency handoff for cellular data networks. Technical Report Technical Report UCB/CSD-96-899, University of California at Berkeley, 1996.

[54] C. Tan, K. Lye, and S. Pink. A fast handoff scheme for wireless networks. In *Second ACM International Workshop on Wireless Mobile Multimedia*, pages 83–90, August 1999.

[55] J. Tourrilhes. The Linux Wireless LAN Howto. HTML Document, Feb 2002. (available at `http://www.hpl.hp.com/personal/Jean_Tourrilhes/Linux/`). accessed 28.1.2003.

[56] J. Tourrilhes. Wireless tools for Linux. HTML Document, Feb 2002. (available at `http://www.hpl.hp.com/personal/Jean_Tourrilhes/Linux/Tools.htm`). accessed 28.1.2003.

[57] N. D. Tripathi. Generic Adaptive Handoff Algorithms using Fuzzy Logic and Neural Networks, Ph. D. Thesis. Virginia Polytechnic Institute and State University, August 1997.

[58] N. D. Tripathi, J. H. Reed, and H. F. VanLandingham. Handoff in Cellular Systems. *IEEE Personal Communications*, 5(6):26–37, December 1998.

[59] H. J. Wang, R. H. Katz, and J. Giese. Policy-Enabled Handoffs Across Heterogeneous Wireless Networks. In *Second IEEE Workshop on Mobile Computing Systems and Applications (WMCSA '99)*, New Orleans, LA, February 1999.

[60] IEEE 802.11 WG. Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications. ANSI/IEEE Std 802.11, 1999 Edition.

[61] M. Ylianttila, R. Pichna, J. Vallström, J. Mäkelä, A. Zahedi, P. Krishnamurthy, and K. Pahlavan. Handoff Procedure for Heterogeneous Wireless Networks. In *Proceedings of Globecom'99*, pages 2783–2787, Rio De Janeiro, Brazil, December 1999.

[62] C. Zhang and V. Tsaoussidis. TCP Real: Improving Real-time Capabilities of TCP over Heterogeneous Networks. In *Proceedings of the 11th IEEE/ACM NOSSDAV 2001*, Port Jeffersson, New York, June 2001.