Network mobility support in IPv6

Thierry Ernst

Network Mobility Support in IPv6

(Le Support des Réseaux Mobiles dans IPv6)

A thesis presented by

Thierry ERNST

in fulfillment of the requirements for the degree of

Doctor of Philosophy

to the

Department of Mathematics and Computer Science

France

October 29, 2001

committee

M. Andrzej DUDA, INPG	*President*
M. Hossam AFIFI, INT	*Referee*
M. Bernard TOURANCHEAU, Sun Labs	*Referee*
M. Laurent TOUTAIN, ENST	*Referee*
M. Roland BALTER, UJF	*Director*
M. Claude CASTELLUCCIA, INRIA	*Advisor*
M. Hong-Yon LACH, Motorola Labs	*Advisor*

prepared at

INRIA Rhône-Alpes (PLANETE Team) and Motorola Labs Paris

THESE

pour obtenir le grade de

Docteur

de l'Université Joseph Fourier de Grenoble

(Arrêté ministériel du 30 mars 1992)

Spécialité : Informatique

présentée et soutenue publiquement par

Thierry ERNST

le 29 octobre 2001

Le Support des Réseaux Mobiles dans IPv6 (Network Mobility Support in IPv6)

Jury composé de

M. Andrzej DUDA, INPG	*Président*
M. Hossam AFIFI, INT	*Rapporteur*
M. Bernard TOURANCHEAU, Sun Labs	*Rapporteur*
M. Laurent TOUTAIN, ENST	*Rapporteur*
M. Hong-Yon LACH, Motorola Labs	*Examinateur*
M. Roland BALTER, UJF	*Co-Directeur*
M. Claude CASTELLUCCIA, INRIA	*Co-Directeur*

Thèse préparée au sein du projet PLANETE de l'INRIA Rhône-Alpes
et de Motorola Labs Paris

Network Mobility Support in IPv6

Abstract: This thesis is devoted to the study of *network mobility support* in IPv6, the new generation of Internet Protocol. *Network mobility support*, unlike *host mobility support*, is concerned with situations where an entire network changes its point of attachment in the Internet topology. The purpose is to provide continuous and optimal Internet access to all nodes located in the mobile network. *Network mobility support* must be considered separately from *host mobility support* because it raises a number of new issues concerning the question of addressing, locating and routing. Our first contribution is the definition of a taxonomy that is used to describe all existing *host mobility support* schemes. Our second contribution is the definition of a new terminology, set of issues and set of requirements specifically targeted to *network mobility support*. Among possible approaches, we focus on Mobile IPv6, the IETF *host mobility support* standard and we study its ability to support mobile networks. We consider both short and long-term solutions. We propose extending Mobile IPv6 with new features as an immediate solution to address its shortcomings. As for the long term, we address the question of scalability to a large number of mobile networks communicating with a potentially large number of nodes. We propose the use of two distinct multicast techniques as solutions to reduce signaling incurred by Mobile IPv6. The performance of our multicast extensions is evaluated by simulation. We conclude this dissertation with a prospective architecture framework which combines our multicast extensions with a number of other techniques identified during this study.

Key-words: IPv6 - Network Mobility - Host Mobility - Mobility Support - Mobile Networks - Network in Motion - Mobile Routers - Routing - Multicast - Internet - Simulation - Mobile IPv6

Le Support des Réseaux Mobiles dans IPv6

Résumé: Cette thèse est dédiée à l'étude du support des réseaux mobiles dans IPv6, la nouvelle génération du protocole qui régit les communications dans l'Internet. Les travaux traditionnels dans ce domaine se préoccupent de fournir une connectivité permanente pour les stations mobiles. En revanche, l'objet de la présente étude est de traiter séparément le cas d'un réseau tout entier qui migre dans la topologie Internet, ce qui pose un certain nombre de nouveaux problèmes. Nous étudions tout d'abord l'État de l'Art dans le domaine traditionnel du support de la mobilité des stations mobiles. Cette étude nous permet de définir une taxinomie des propositions. En second lieu, nous définissons une nouvelle terminologie dédiée au support des réseaux mobiles, ainsi que leurs caractéristiques et les problèmes spécifiques causés par leur mobilité. Parmi un ensemble d'approches envisagées, nous nous consacrons tout particulièrement à l'usage de Mobile IPv6, le standard de l'IETF pour le support des stations mobiles. Dans un premier temps, nous proposons un certain nombre d'extensions nécessaires à ce protocole. Pour le long terme, nous proposons de réduire le coût des messages de contrôle induit par ce protocole au moyen de deux techniques multipoint. La première, dite traditionnelle, établit un arbre de distribution entre le réseau mobile et ses correspondants. La deuxième enregistre directement la liste des correspondants dans le message de contrôle. La performance de ces extensions multipoint est évaluée par simulation, et nous concluons cette dissertation par une vue d'ensemble d'une nouvelle architecture de gestion de la mobilité rassemblant diverses techniques, dont nos extensions multipoint.

Mots-Clés: IPv6 - Réseaux - Support Mobilité - Réseaux Mobiles - Routeurs Mobiles - Routage - Communications Multipoint - Internet - Simulation - Mobile IPv6

Acknowledgments / Remerciements

This Ph.D. has been supported by Motorola Labs Paris and ANRT (French "Association Nationale de la Recherche Technique") under CIFRE convention number 97595. This CIFRE convention associates: a firm, Motorola Labs (Paris, France); a French academic laboratory, INRIA Rhône-Alpes[1]; and a French university, UJF[2]. I would like to thank all people in the above cited organizations that made this Ph.D. possible, and particularly my two advisers, Hong-Yon Lach from Motorola Labs (NAL team), and Claude Castelluccia from INRIA (PLANETE team) for their direction, enlightenment, advices and support. I would also like to thank my board of examiners for their precious time and comments: Hossam Afifi[3], Roland Balter (UJF), Andrzej Duda[4], Bernard Tourancheau[5], and Laurent Toutain[6]. Last, but not least, I thank all NAL and PLANETE members. But since this Ph.D. has been prepared in France, I am saving my more cheerful and personal acknowledgments to write them in French, my mother tongue.

Tout d'abord, et une fois n'est pas coutume, je vais commencer par ne pas remercier les voleurs de vélo, empêcheurs de tourner en rond et autres briseurs de rêves qui n'ont pas cru en mes motivations, un jour ou l'autre au cours de mon long parcours, depuis les premiers pas d'écolier jusqu'à ma soutenance, et qui ne liront certainement jamais ces mots. Paradoxalement, il est probable que sans eux je n'eusse jamais suivi ce parcours et que je n'eusse jamais entamé une thèse ... Alors, au fond peut-être m'ont-ils aidé à surmonter les obstacles.

En face des briseurs de rêves, il y a tous ces gens qui m'ont émerveillé, soutenu, guidé, sensibilisé, entraîné; ces gens qui transportent l'enthousiasme et le communiquent. Je voudrais leur adresser mes plus vifs remerciements. Pour commencer, mon professeur d'anglais au collège, sans l'enthousiasme de laquelle je n'aurais jamais suivi ce chemin, et sans doute pas écrit ma thèse en anglais. Elle ne le sait pas, mais c'est un peu grâce à elle que tout a commencé, en partant de Strasbourg pour arriver à Grenoble en passant par Middlesbrough et Sophia-Antipolis. Les chemins les plus courts ne sont pas forcément les meilleurs. Merci au passage à tous ceux, sur ces étapes, qui ont rendu ce périple possible. Merci à Walid Dabbous pour ses cours de réseau enchanteurs, c'est un conteur tel qu'on en trouve généralement pas dans une discipline scientifique et technique. Sans ses cours et son dynamisme, je n'aurais peut-être pas découvert le monde passionnant des réseaux. Merci ensuite à toutes les équipes de l'INRIA qui m'ont accueilli. Merci à mes deux superviseurs, Claude Castelluccia pour l'INRIA, et Hong-Yon Lach pour Motorola. Merci à Claude, dont j'ai manqué le premier rendez-vous que nous nous étions donnés et qui m'a tout de même pris en thèse et mis en contact avec Hong-Yon. Merci à Hong-Yon pour sa compréhension lorsque j'ai préféré les montagnes grenobloises aux embouteillages parisiens. Merci à tous les deux pour m'avoir fait confiance et encadré durant cette thèse, merci pour votre patience lorsque les simulations promises ne venaient pas. Merci à Roland Balter pour l'ingrate gestion administrative. Merci à tous les autres membres de mon jury pour avoir accepté, de gré ou de force, d'en faire partie, pour votre temps, vos commentaires, et pour les tracasseries administratives habituelles ou inhabituelles. Merci aux membres des équipes NAL et PLANETE pour leur enthousiasme, leurs commentaires, leur aide, leur soutien et leur amitié, ainsi qu'aux autres dispersés dans d'autres labos. Merci aussi à tous les autres, tous ces diffuseurs de connaissance, professeurs, élèves et étudiants, camarades, ou collègues, au cours de mes études, stages, et de ma thèse, qu'il serait impossible de tous nommer.

Enfin, merci à tous ces amis et êtres chers, grenoblois, parisiens ou d'ailleurs pour leur soutien et leur compréhension, les séances photos acrobatiques, les absences glacières, les couleurs primaires, les sorties ski et luge, les gîtes et les randos, les cinés et les soirées, les tartiflettes et les vendanges tardives, les sambas et les séances musicales, les sushis et les sashimis, les valises en cuir et les photos numériques, les Père-la-Chaise, les gags en tous genres, les pannes d'essence, la fourrière les jours de marché, les déménagements express, les sourires, qui sont autant de surprises, de moments de rire, de joie et de partage, au cours de la

[1] INRIA: Institut National de Recherche en Informatique et en Automatique (French national institute for research in computer science and control), Grenoble, France

[2] UJF: Université Joseph Fourier, Grenoble, France

[3] INT: Institut National des Télécommunications, Evry, France

[4] INPG: Institut National Polytechnique de Grenoble, France

[5] Sun Microsystems Laboratories, Grenoble, France

[6] ENST: Ecole Natioanle Supérieure des Télécommunications de Bretagne, France

préparation de cette thèse. Merci tout spécialement à celui qui pour une fois ne sera pas le premier nommé et à qui j'ai laissé la lourde tâche de régler bien des détails suite à mon départ pour un autre monde.

A tous, grand merci.

Contents

Foreword and Naming Conventions

In order to spread to a more important number of people, this Ph.D. dissertation is written in English, although prepared in France (INRIA Grenoble and Motorola Labs Paris) and defended in a French University (Université Joseph Fourier Grenoble). Thus, the author apologizes for all the remaining spelling and style mistakes due to the lack of experience in writing English. For the French reader, an extended French summary details each chapter and can be found in the appendix.

The following naming conventions and styles are employed throughout this document: literal words, protocol names and well known acronyms are `emphasized using this style`. New terms and terminology peculiar to mobility are introduced whenever required and usually emphasized using this style. Other topics or terminology that need to be emphasized *will appear in this style*.

Introduction

Motivations and Objectives

As we see in today's life, geographical mobility of people is increasing. This is the result of the pressure of the professional life and the family scattering, which impact the social life, this in turn generating a need for more mobility. In these conditions, anyone would wish to benefit from the same social and professional environment without restriction of the current geographical location. The Era of digital information could in a way achieve this wish. More and more executives or representatives are expecting to transfer files from their workplace file system, to obtain on-line information, to communicate with their customers and providers as if they were at their office in front of their computer. Similarly, a traveler would like to stay in touch with his family and friends, sending them photographs and sounds, while listening to its favorite music. As a result from this, there is a continuous interest in the Internet, the most appropriate media for digital information exchange, while cellular telephony gives people the opportunity to be reachable anywhere. Despite this, the cellular network is currently tuned to carry voice only although there is also a desire to transmit other types of data, whereas the Internet doesn't allow effective mobile communications as in cellular telephony.

At the same time that mobility of people is required, recent advances in computer miniaturization and wireless technology promise increasingly powerful, light, small and functional wireless devices. The hardware miniaturization together with the improvement of wireless communication technologies drive the need for even more mobile communications. As more and more people are traveling with a laptop, a PDA, a WAP or i-mode phone, a digital camera, or any other high-tech device, there is a desire to connect it to the Internet from anywhere, at anytime, and to remain permanently connected to it without any disruption of service. No one should be abstained from using its usual computing resources and Internet access while moving, especially when traveling by train or by plane. However, the Internet it is not tuned to allow mobility in the midst of data transfers because protocols used in the Internet are not conceived for devices that frequently change their point of attachment in the Internet topology. Basically, something similar to cellular telephony as compared to fixed telephony is needed in the Internet. The Internet must be upgraded with *mobility support*.

Indeed, *mobility support* is not only concerned with mobile devices. There are situations where an entire network could migrate in the Internet topology, which we refer to as a *mobile network*. Applications include networks attached to people (Personal Area Network or PAN) and networks of sensors deployed in aircrafts, boats, cars, trains, etc. For instance, an airline or a train company could provide permanent on-board Internet access, allowing passengers to use their laptop, PDA, or mobile phone to connect to remote hosts, download music or video, browse the web, etc. During an international fare, the aircraft or the train changes its point of attachment to the Internet and gets Internet access from distinct Internet Service Providers. Similarly, a coach, the metropolitan public transport, or the taxi company could allow passengers to connect their PAN to the Internet via the embarked network, therefore ensuring, while on-board, an alternative to the metropolitan cellular network, in terms of price or available bandwidth, access control, etc. Meanwhile, a number of Internet appliances deployed in the mobile network are used to collect traffic and navigation data from the Internet while sensors within the mobile network collect and transmit to the Internet live information, like the current number of

passengers, expected time to arrival, the amount of petrol left in the tank, etc. For a number of reasons (network management, security, performance,...), it is desirable that Internet appliances deployed in cars, trains, busses, etc. do not connect individually and directly to the Internet, therefore exhibiting the need to displace an entire network.

Traditional work turning around mobility is to provide continuous Internet access to *mobile hosts* only (*host mobility support*). Despite this, there is currently no means to provide continuous Internet access to nodes located in a *mobile network* (*network mobility support*). Indeed, the question of moving networks that frequently change their points of attachment in the Internet has not yet or rarely gained the deserved attention from the Internet community, although there is a perceived need for it. It is therefore the purpose of this dissertation to address this lack. We study further aspects of mobility in the Internet and we do the spade-work on the question of routing packets to and from nodes located in *mobile networks*.

The following analogy may help to clarify the generic mobility problem addressed by this study. In fixed telephony, every customer is traditionally allocated a phone number that both identifies the phone, and its location in the phone network: the prefix, according to its length, determines the country, the area in the country, the city, and the district. The problem arises when a subscriber moves and wants to be reachable by its usual correspondents wherever he happens to be. In this situation, the same number cannot be used anymore. Before the advent of mobile telephony, the user had to pick up a new phone every time he moved to a new place and to provide the new phone number to all its correspondents. This necessitated to agree on a protocol in order to decide when, how and who is going to advertise the new phone number to all the correspondents, and how long the phone number could actually be used. In mobile telephony, a permanent phone number is attributed to the subscriber who is reachable anywhere, in its country, as long as he carries his mobile phone with him. Ongoing communications are maintained without disruption and transparently to the subscriber and its correspondents. All the burden of the mobility management is on the telecommunication network that keeps track of the location of the mobile phone.

In the Internet, there is typically a change of the physical IP address each time an Internet appliance changes its point of attachment. This results in losing packets in transit and breaking transport protocols connections if no specific services are added to handle mobility. Mobility support is therefore to give the ability to cross networks in the midst of data transfers without breaking the communication session and without increasing the network load and delays. In addition, *network mobility support* is to address the question of providing a continuous network access to all nodes located in a *mobile network*. The aim is to route packets from and to *mobile networks* optimally,, efficiently, and transparently to upper layer protocols although this might cause a change of the IP address at the network layer. The long-term goal of this research project is to provide a network-based solution to support network mobility efficiently, and with minimum signaling overhead.

The Internet Protocol is the protocol that effectively allows any two Internet nodes to communicate with one another. There presently exists two versions of the protocol, the getting old IPv4, and the new generation IPv6. We have decided to focus our study on IPv6. A number of reasons has favored the choice of IPv6 over IPv4. First, IPv6 is meant to replace IPv4 and to address its shortcomings. IPv6 has built-in features that allow to support the new services requested by recent applications. This includes support for mobility, multicast, traffic reservation, security, etc. It is thus easier to bring extensions to IPv6 rather than to IPv4, all the more since deployment of IPv6 has not yet started. A second reason is that IPv6 offers a generous number of addresses compared to IPv4. Asian countries, which were not granted a significant part of the IPv4 address space, but which represent a significant part of the world population, and thus an important market, thanks to China and India, are potentially more concerned by this, and are showing a tremendous interest in IPv6. Cellular operators are also concerned by the available address space and the built-in features of IPv6 since they would like to connect each mobile phone to the Internet. Basically, any kind of device, like a fridge, micro-wave, watch, may be connected to the Internet, this driving the need for an even more important number of addresses. On a more personal basis, we advocate that it's not very wise to bring new functionalities to a protocol that has exhibited its limits. It's time now to move to IPv6, to focus the effort on IPv6, to bring to IPv6 the applications it deserves and to realize the dream everyone is looking for.

As it is not the intend of this dissertation to elaborate security aspects, we will not enter into the details although this topic deserves a lot of interest. It would probably deserve a dissertation of its own and we are lacking time and space for it. We will only have a few words about it from place to place to outline issues.

Organization of this Dissertation

This Ph.D. dissertation thus investigates issues arising when an entire network changes its point of attachment in an IPv6 wide-area network. The document is structured as follows:

Part 1: We first define the terminology that we are going to use in this dissertation, and then we describe the TCP/IP protocol suite and particularly the *network layer*, in charge of node-to-node communication. We focus on IPv6 and we also detail multicast routing on which our solutions will be based later in this dissertation. We then detail the general problem caused by mobility, and why the network layer cannot handle it efficiently. As a result of mobility, a new route must be found. Since the IP address must reflect the location in the Internet topology, mobility usually generates a change of the physical IP address every time a node is attached to a new link in the Internet topology. This poses two questions: how to advertise the new topological location and how to handle the change of address at the transport layer where the IP address is used as an identifier. Once the mobility problem is defined, we study the *State of the Art* in the area of *host mobility support*. This study is essential in order to investigate how current *host mobility support* schemes could be applied to *network mobility support*. Their study shows that all schemes make use of some well-defined components and functions. Schemes could be summarized by the location of these components in the network architecture, and the functions they perform. According to an abstraction model which we define, all proposals are then fetched in a few set of frameworks which each exhibit some specific characteristics of the proposals. We conclude this part by a taxonomy of the frameworks.

Part 2: We address the question of *network mobility support* specifically. We highlight some characteristics peculiar to mobile networks, and then we list a number of issues and requirements. A few potential approaches we went through during the course of this study and derived from the study of the *State-of-the-Art* are investigated. As a solution based on Mobile IPv6 (the existing IETF standard for *host mobility support* seems promising, the following chapters discuss its ability for *network mobility support*. Mobile IPv6 shortcomings are identified and explained. Then, we propose some extensions that could gradually address these shortcomings. We particularly address the question of scalability of Mobile IPv6. The point of novelty of our approach is to combine Mobile IPv6 with multicast to distribute the topological location of a *mobile network* to its correspondent nodes. We propose two distinct multicast techniques as a solution to reduce Mobile IPv6 signaling. The first technique relies on traditional multicast protocols and is most suitable for a large number of correspondent nodes. The second one is a new multicast technique that records the list of destinations in the packet itself and is more suitable for a small number of correspondent nodes due to the packet overhead. A prospective architecture that combines our multicast extensions with other mechanisms concludes this part.

Part 3: The performance of our multicast extensions is evaluated by means of simulation. We first start by a presentation of our simulation tool, based on NS-2, the configuration of our simulations and our metrics used for the evaluation. Then, we conclude by the performance analysis that validates our solutions.

Conclusion: We conclude by discussing our contributions and the perspectives of this study.

Part I

Mobility in the Internet

Chapter 1

Terminology

During the course of this dissertation, we will make use of a number of terms pertaining to networking and mobility. Our terminology adopts some of the terminology already defined in the IPv6 (RFC 2460 [Deering and Hinden, 1998], RFC 2461 [Narten et al., 1998]) and Mobile IPv6 [Johnson and Perkins, 2000] specifications, and more generally in [Tanenbaum, 1996; Stewart, 1998; Huitema, 1998].

The first section defines the usual terminology for describing the architecture of the Internet whereas the second section defines the usual terminology for describing mobility. As far as mobile networks (section 1.3) are concerned, we need to define our own terminology because no one presently exists to define the issues, goals, architecture elements, problems and requirements pertaining to *network mobility support*. Finally, the fourth section defines distinct types of mobility.

1.1 The Internet

The Internet is a collection of heterogeneous networks hierarchically organized and running under distinct administrative policies. A *network* is simply speaking a collection of *nodes* and *links*. The Internet terminology distinguishes two kinds of nodes: a *router* is a node that forwards packets not explicitly addressed to itself whereas a *host* is any node that is not a router. We will refer to the term *end-node* as the node that initiates or terminates the transmission of a packet, i.e. the source or the destination of the packet. Any router that forwards the packet closer to the destination on the path between the source and the destination will be referred to as an *intermediate router*. A node's attachment to a link is termed *interface*. Nodes may have any number of interfaces, and each interface may be attached to distinct links[1]. All nodes connected on the same communication link form what is usually term a *subnetwork* (typically, an Ethernet link, or a *802.11b WLAN*). Subnetworks are interconnected by means of routers. Thus, a router typically has at least two interfaces and routers are primarily used to forward traffic between subnetworks.

Typically, the Internet is partitioned into different administrative *domains*. A domain shall usually correspond to a large organization operating an internal corporate network or an Internet Service Provider (ISP), UUNET, MCI, Sprint, Wanadoo, AOL are instances of ISPs[2]. A domain may be further divided into a number of *sites* or independent networks. Although the term *site* has no common definition, in this document we will refer to a *site* as a subdivision of a bigger network belonging to an organization and limited to a geographical area, or to a particular branch of the organization. A campus shall be a good instance of a site. Sites are interconnected by means of a high speed *transit network* (a *backbone*). A *stub network* is a *leaf network*, i.e. network that does not forward packets that does not originate or terminate within itself. The gateway between networks under distinct administrative policies

[1] In this situation, we say that the node is *multi-homed*

[2] The two latest are *dial-up* ISPs

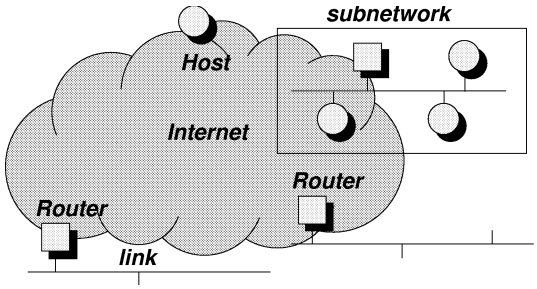

Figure 1.1: Symbols

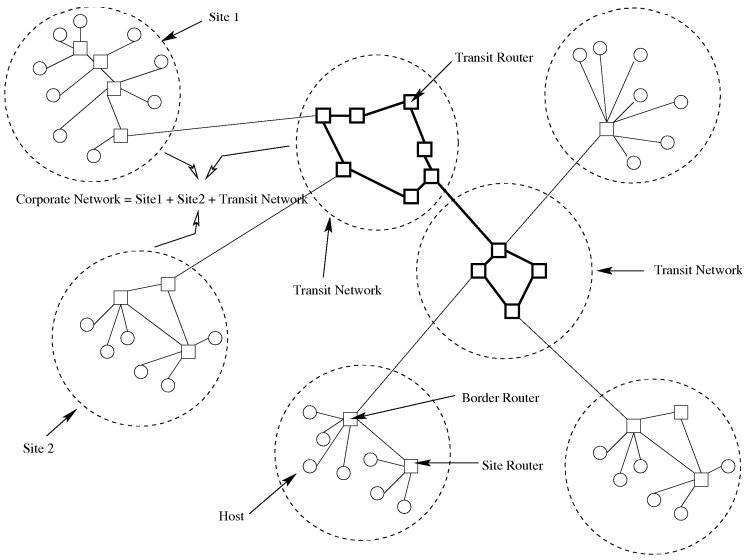

Figure 1.2: Instance of a corporate network partitioned into sites

is a *Border Router* (BR). Several BRs may be attached to the same transit router. We shall therefore refine a *network* as a set of subnetworks that are attached to the `Internet` through a common border router. Whatever the term employed, (domain, site, or network), a set of subnetworks under a common administrative policy is usually referred to an *Autonomous System (AS)* as far as routing is concerned.

To illustrate this terminology in our figures, we will make use of a symbolic representation as illustrated on fig.1.1. Figure 1.2 further illustrates our terminology by showing an instance of a corporate network partitioned into several sites.

The role of *internetworking* is to interconnect all the networks that form the `Internet` so that any two nodes can communicate with each other. As a result from this, the `Internet` is not specific network-technology-dependent, allowing a global network of unlimited scope and reach. This has largely accounted for its success. Internetworking is performed by the `TCP/IP` protocol suite, a *packet-switched* technology. Unlike *circuit-switched* technologies like `ATM` or telephone networks, `TCP/IP` it relies on the *connectionless* concept. In this concept, routers cooperate to determine the path toward the destination and carry packets between the two nodes. The forwarding decision called *routing* is made on a per-packet basis. The intelligence is indeed put at the edge of the network (i.e. end-nodes), whereas the purpose of the network infrastructure is only to provide internetworking. This allows an easy deployment of new functionalities without need to upgrade the network infrastructure.

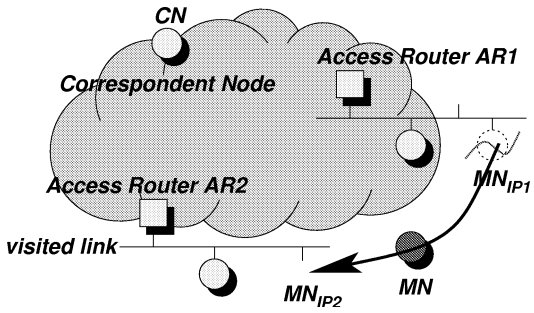

Figure 1.3: Mobility Reference Model

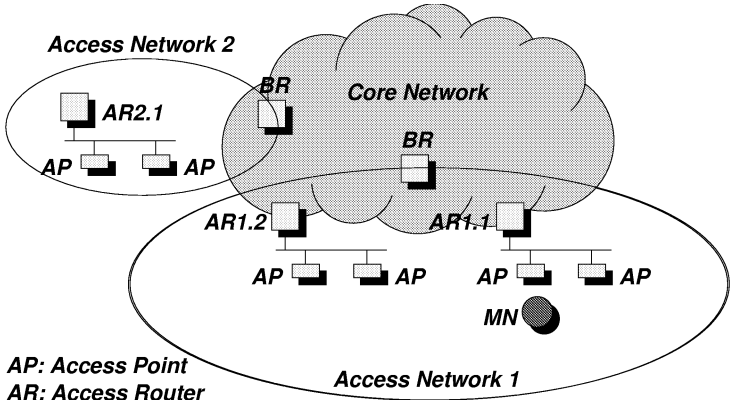

Figure 1.4: Access Network

1.2 Nodes in Motion

We shall refer to a mobile node as an `Internet` node that changes its point of attachment to the network topology, i.e. a node that moves from a subnetwork to another. We refer to visited links as the subsequent subnetworks where a mobile node is attached to. The routers that serve the visited link and provide `Internet` access to mobile nodes are termed *access routers* (ARs). To illustrate this, figures in this dissertation will make use of the mobility reference model as illustrated on fig.1.3 which shows a mobile node moving between two subnetworks. Note that a box filled with red typically means that the node performs a mobility management function. The terminology that usually applies to cellular mobility is illustrated on figure 1.4. The *access network* is a cellular network that provides `Internet` access to wireless nodes. The *access point* (AP) is the link-layer attachment point that interfaces between a wireless technology and the subnetwork.

1.3 Networks in Motion

We refer to a *mobile network* as a *network* whose border router dynamically changes its point of attachment to the `Internet` and thus its reachability in the topology. Our study is concerned by concrete instances of mobile networks that may be deployed in the near future and for which there already exists a tremendous need. Those includes trains, aircrafts, cars, buses that want to offer permanent `Internet` access to `Internet` appliances carried by passengers and fixed appliances deployed within the mobile network.

As an example of a mobile network, an airline company could provide permanent on-board `Internet` access, allowing passengers to use their laptops, PDA or mobile phone to connect to remote hosts, download

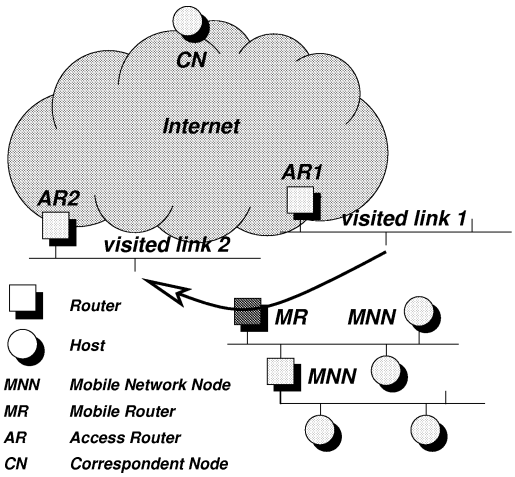

Figure 1.5: Terminology for Mobile Networks

music or video, browse the web, etc (this scenario is mentioned in [Tanenbaum, 1996] in section.1.2.4 and section.5.5.8). At the same time, air control traffic could be exchanged between the aircraft and air traffic control stations (this scenario has been investigated by `Eurocontrol` - European Organization for the Safety of Air Navigation - since 1998 [Quinot, 1998]). During a transatlantic flight, the aircraft changes its point of attachment to the `Internet`. Over the oceans, the aircraft gets connected to the `Internet` through a geostationary satellite; over the ground, it's through a radio link. Handoffs do typically not occur very often (a radio link may cover 400-500 kilometers), but it may happen between distinct `ISPs`. In our scenario, we may consider that passengers could themselves carry a network, for example `Personal Area Networks`, i.e. a network composed by all `Internet` appliances carried by people, like a `PDA`, a mobile phone, a digital camera, a laptop, etc. Another similar scenario involving ships and aircrafts is mentioned in `RFC 1726` ([Partridge, 1994], section 5.15).

To describe such kind of scenarios, we need to define a new terminology in addition to the already existing terms. We therefore introduce the following new terms relevant to mobile networks. First, we refer to the border routers that attach the mobile network to the rest of the `Internet` as the mobile routers (MRs). A mobile router has at least two interfaces, the first attached to the visited link, and the other attached to an internal link of the mobile network. We call *mobile network node* (MNN) any host or router located within the mobile network, either permanently or temporarily. A MNN may be any of a mobile router, a local fixed node, a local mobile node, or a visiting mobile node. All MNNs share a common and permanent `IP` prefix that we call the mobile network prefix. The mobile network prefix is a bit string that consists of some number of initial bits which identifies the set of subnetworks that compose the mobile network. It also identifies the topological location of the mobile network when the mobile router is attached to its home link. In addition, we call *correspondent node* (CN) any external node that is communicating with one or more MNNs. This new terminology is summarized in fig. 1.5 and detailed hereafter:

- mobile network: a set of nodes composed by one or more `IP`-subnets attached to a mobile router and mobile as a unit, with respect to the rest of the `Internet`, i.e. a mobile router and all its attached nodes. The mobile router changes dynamically its point of attachment to the `Internet` and thus its reachability in the topology. All nodes in the mobile network share the same `IP` prefix: the mobile network prefix.
- mobile IP-subnet: a mobile network composed of a single `IP`-subnet.
- mobile network node (MNN): any host or router located within the mobile network, either permanently or temporarily. A mobile network node could be any of a mobile router, local fixed node, local mobile node, or visiting mobile node.

- mobile router (MR): the border router which attaches the mobile network to the rest of the Internet. The mobile router has at least two interfaces, an external interface, and an internal interface. The mobile router maintains the Internet access for the mobile network. It is used as a gateway to route packets between the mobile network and the fixed Internet.
- local fixed node (LFN): any host or router permanently located within the mobile network and that does not change its point of attachment.
- local mobile node (LMN): a mobile node that belongs to the mobile network and that changes its point of attachment from a link within the mobile network to another link within or outside the mobile network (the home link of the local mobile node is a link within the mobile network).
- visiting mobile node (VMN): a mobile node that does not belong to the mobile network and that changes its point of attachment from a link outside the mobile network to a link within the mobile network (the home link of the VMN is not a link within the mobile network). A VMN that attaches to a link within the mobile network obtains an address on that link.
- node behind the MR: synonym for a mobile network node (MNN).
- correspondent node (CN): any node located outside the mobile network that corresponds with one or more MNNs. CNs corresponding with MNNs located in the same mobile network are said to be CNs of this mobile network.
- home prefix: a bit string that consists of some number of initial bits of an IP address which identifies the home link within the Internet topology (i.e. the IP subnet prefix corresponding to the mobile node's home address, as defined in Mobile IPv6).
- foreign prefix: a bit string that consists of some number of initial bits of an IP address which identifies a foreign link within the Internet topology.
- mobile network prefix: a bit string that consists of some number of initial bits of an IP address that is common to all IP addresses in the mobile network (i.e. all MNNs have the same IPv6 network identifier). For a mobile network restricted to a single mobile IP-subnet, the mobile network prefix is the network identifier of this subnetwork. In some circumstances, the mobile network prefix may be that of the home prefix or the foreign prefix with a longer number of bits, but not necessarily, as this will be developed later in this study.
- external interface of a MR: the interface attached to the home link if the mobile network is at home, or attached to a foreign link if the mobile network is in a foreign network.
- internal interface of a MR: interface attached to a link inside the mobile network. This interface is configured with the mobile network prefix.
- nested mobility: a mobile network that comprises visiting mobile nodes or even mobile networks. For instance, a bus is a mobile network whereas a passenger is either a VMN in a mobile network if it carries a mobile phone or a mobile network in a mobile network if it carries a PAN.
- multi-homing: a mobile network that has two or more active interfaces connected to distinct parts of the Internet. This could either be a single MR with two interfaces simultaneously connected to the Internet, or the mobile network may be connected to the Internet via two or more MRs. In the first case, we could think of a unique router used to connect a car both to the cellular phone network and to a navigation satellite. In the second case, we may think of a PAN where a GSM phone is used to connect the PAN to the cellular phone network whereas a Bluetooth PDA is used to collect bus timetables from the city bus network. In this situation both the phone and the PDA are mobile routers.
- idle mobile network: a mobile network that does not engage in any communication outside the mobile network may be considered as idle from the point of view of the fixed Internet, although there may be internal traffic between any two MNNs.
- idle mobile network node: a MNN that does not engage in any communication.

An Ad-hoc network as defined in the IETF MANET Working Group is not to be confused with a mobile network. An ad-hoc network, is an autonomous system of mobile nodes (i.e. routers) connected by wireless links. The routers are free to move randomly and to organize themselves arbitrary. In a mobile network, some routers may effectively move arbitrary, but this not a common case. However, an ad-hoc network connected to the Internet which changes its point of attachment may be considered as a special instance of a mobile network and may exhibit common issues. Ad-hoc networking is principally concerned with routing packets between any two nodes in the ad-hoc network.

1.4 IP-Layer Mobility

IP-layer (or network-layer) mobility arises when a portion of the Internet changes its point of attachment in the IP hierarchy. We will speak about *host mobility* when a host changes its point of attachment to the Internet topology. We will speak about *network mobility* when the *router* that connects an entire network changes its point of attachment to the Internet topology. We shall use the term mobile node alternatively for a mobile host or a mobile router as long as we don't pay attention to potential nodes behind the mobile router.

IP-layer mobility occur in situations where a node is plugged from one subnetwork to another or preferably where a wireless node connects to the Internet by means of any wireless technology, for instance 802.11b WLAN, Bluetooth, satellite link, GSM, etc.

We note that a topological displacement does not necessarily preclude a geographical displacement. This may for instance be the case when a mobile node is able to connect to the Internet by means of two or more wireless technologies or when it switches from one ISP to another that offers better prices. Similarly, a geographical displacement does not preclude a change of the point of attachment to the Internet topology. This may arise when a mobile node is, for instance, attached to a wireless access point which spans a very large geographical area or when a mobile node switches from one access point to another that belongs to the same subnetwork (for example a node that switches from a 802.11b WLAN AP to a GSM AP). In this situation, the mobile node is still attached to the same subnetwork. From a *network layer* point of view, there is no change of topological location, and no change of IP address either. This type of mobility is best referred to as *link-local mobility* and is best handled at the *link-layer*. It is therefore out of scope of the present study and will be left out throughout this report.

Two subnetworks may be geographically very close but topologically distant. Given the fact that topologically distant sections of the Internet usually belong to distinct domains or sites, mobility could be classified according to the following two definitions:

- Local-Area Mobility refers to mobility within a single administrative domain, i.e. between subnetworks topologically close in the IP hierarchy. In the literature, and depending on the definition of "closeness", this is also termed intra-site mobility, intra-domain mobility, local mobility or micro-mobility. As an instance of Local-Area Mobility, the displacement of a node within a limited vicinity of adjacent subnetworks, like in a campus, that belong to the same organization or between ARs that belong to the same ISP.

- Wide-Area Mobility refers to mobility across domain boundaries, i.e. between subnetworks topologically distant in the IP hierarchy. In the literature, and depending on the definition of "remoteness", this is also termed *inter-site mobility*, *inter-domain mobility*, or *global mobility*, or macro-mobility. As an instance of Wide-Area Mobility, displacement of a node between distinct ISPs or organizations, or between widely separated sites of a single organization.

Both types of mobility could be illustrated on figure 1.4 where a mobile node is moving between two distinct sites separated by a large internetwork depicted by the cloud. When the MN moves from the

Where	Topologically close mobility	Topologically distant mobility
Administrative Domain	intra-domain mobility	inter-domain mobility
Geographic Domain	intra-site mobility	inter-site mobility
IP Cellular networks	micro-mobility	macro-mobility
IP hierarchy	local-area mobility local-mobility	wide-area mobility global-mobility

Table 1.1: Types of Mobility: Terminology

subnetwork served by AR 1.1 to the subnetwork served by AR 1.2, it performs Local-Area Mobility. When it moves from this latter subnetwork to the subnetwork served by AR 2.1, it performs Wide-Area Mobility.

Tab. 1.4 summarizes the wording used in the literature. Intra-domain / Inter-domain mobility mostly refer to as mobility within and between distinct domains of administration. The former refer to as mobility between subnetworks that belong to a single administrative domain (or `autonomous system`) whereas the later refer to as crossing domain boundaries. Intra-site / Inter-site mobility is very close to this definition, but usually refers to mobility within and between geographical areas that belong to a single administrative domain. Micro / Macro mobility is more associated with cellular telephony and respectively refer to as mobility between ARs that belong to the same cellular network and between ARs that belong to distinct cellular networks. The last two have a more generic meaning.

Chapter 2

The TCP/IP Reference Model and Addressing

The first section introduces the TCP/IP *reference model* and its addressing scheme. Section 2.2 describes the two versions of the IP protocol; IPv4 (RFC 791) [Postel, 1981a] is briefly presented, while we pay more attention to IPv6 (RFC 2460) [Deering and Hinden, 1998]. This section also outlines other protocols that assist the network layer. Section 2.3 focuses on routing, we describe the mechanism, the routing techniques, and the protocols, for both unicast routing and multicast routing. Multicast is particularly detailed because our solutions to support networks in motion is based on multicast protocols, as this will be described in chapter 7. Section 2.4.1 discusses the mobility paradigm, i.e. why the TCP/IP addressing prevents mobility of IP nodes and concludes that a number of mobility support services are required.

2.1 The TCP/IP Reference Model

The Open Systems Interconnection (OSI) [ISO, 1984] reference model separates networking in seven distinct layers [Tanenbaum, 1996]. It is defined by the International Standards Organization (ISO). Each layer of the model performs a well defined function. The model only specifies what are the tasks of each layers whereas the specification of the exact services and protocols is left aside. The purpose of a separation into layers is to render layers independent one from another. This facilitates the enhancement of one layer without after-effects on the other ones.

The TCP/IP *reference model* defines the protocol suite used for data exchange between hosts in the Internet specifically. As its name stands for, it is named after its two main protocols, Transport Control Protocol (TCP) [Postel, 1981b] and Internet Protocol (IP). It is very similar to the OSI *reference model*, but does not map well into the seven layers as defined in the OSI *reference model*. It has fewer layers than its OSI counterpart as the use of intermediate abstraction layers between the *transport layer* and the *application layer* was not perceived.

2.1.1 TCP/IP Layers

We briefly describe the TCP/IP layers. We will pay more attention to the *network layer* in the forthcoming sections.

- *Network Access Layer*: defines the characteristics of the hardware which carries the communication

signal and actually allows communication between nodes on the same subnetwork. It includes the OSI Data Link Layer and the OSI Physical Layer. This layer deals wit pure hardware and access methods. This layer includes Ethernet, IEEE 802, etc. A 48-bit MAC (Medium Access Control) address is used for all communications at this layer.

- *Network Layer*: provides internetwork-wide reachability between any two subnetworks. It roughly corresponds to the OSI Network Layer and is entrusted the tasks of addressing and routing. It is mainly made of the Internet Protocol (IP) assisted by a number of other protocols like routing protocols.

- *Transport Layer*: corresponds to the OSI Transport Layer and is entrusted end-to-end datagram delivery between a source and its ultimate destination end-node. TCP (Transport Control Protocol) (RFC 793) and UDP (User Datagram Protocol) (RFC 768) are the most widely used transport protocols. TCP is a reliable connection-oriented protocol which ensures datagrams delivery, in the right order, and without corruption between two end points whereas UDP is its unreliable connection-less counterpart.

- *Application Layer*: sits right on top of the transport layer in the TCP/IP reference model. It includes the OSI Presentation Layer and the OSI Session Layer and the OSI Application Layer. Protocols like HTTP, FTP, TELNET, DNS sit there.

2.1.2 The TCP/IP Addressing Scheme

Protocols at each layer require identifiers. Thus, the TCP/IP reference model defines the concept of *IP addresses and port numbers*.IP addresses are allocated to *interfaces*, not to nodes, because a node may have several interfaces on the same subnetwork. The purpose of the IP address is to identify the current topological location of an interface within the Internet (i.e. its network's point of attachment), and the interface itself. As for the port number, it identifies the application running on the node that owns the interface.

The implementation choice for identifying TCP and UDP communication flows is to make use of IP addresses and port numbers of both sender and receiver. The IP address is needed to perform identification checks. When packets are sent, the transport protocol header contains the source and destination port numbers, and the IP header contains the source and destination IP addresses. In other words, TCP and UDP make use of network, transport and application layer identifiers, this resulting in a $< source\ IP\ address$, $source\ port\ number$, $destination\ IP\ address$, $destination\ port\ number >$ 4-uplet communication flow identifier.

In summary, the *IP address* is used at both the network and transport layers whereas the *port number* is used at both the transport and application layers. At the network layer, the IP address is used to determine the subnetwork where resides the interface, and to identify the interface itself. At the transport layer, the IP address is used to identify the node. We therefore note that the IP address serves a dual semantic:

1. identifying the topological location of an interface in the Internet.

2. identifying the node to which the interface belongs to.

Conforming to the OSI terminology, the IP address corresponds to the Network Service Access Point (NSAP) while the {IP address, port number} 2-uplet corresponds to the Transport Service Access Point (TSAP). The NSAP makes the service association between the network layer and the transport layer; the TSAP makes the service association between the transport layer and the application layer.

2.2 IP and Related Protocols

The purpose of the Internet Protocol (IP) is to carry packets between end-nodes over an interconnected system of networks and transparently to higher layers. IP mainly defines a packet format and an addressing scheme. IP is a connectionless protocol in the sense that each IP packet contains all the information it needs to get from a source to its destination. It relies on the *best effort* principle; whereby there is no guarantee that the packet gets actually received by the right destination. In case of non-delivery, another protocol, ICMPv6 [Conta and Deering, 1998], is used to give a diagnostic back to the source. The reliable delivery of packets is the responsibility of the upper layer protocols.

There currently exists two versions of this protocol, IPv4 [Postel, 1981a], and IPv6 [Deering and Hinden, 1998] [Lee et al., 1998]. A major difference between IPv4 and IPv6 is the expansion of the address space from 32 bits to 128 bits. We will not detail IPv4 as the focus of our study is on IPv6.

2.2.1 IPv4: Internet Protocol Version 4

IPv4 was first developed by the Department of Defense's Advanced Research Projects Agency (ARPA). It was expected to be a research network for ARPA only and surely not meant to become such a worldwide standard, and this explains why it doesn't provide space for improvement. The IP header is illustrated in figure 2.1 (a).

IPv4 addresses are 32-bits large. The *network number* is a prefix in the 32 bits address. Its length expressed in a number of bits is varying according to the Class of the address (there are four classes, each class has a distinct network prefix length that determines a maximum number of machines in that network), but never exceeds 24 bits. It identifies the subnetwork. The remaining bits form the *host number* which demultiplexes between hosts in the same subnetwork. As a result of this address format and length, an host can not solely be identified by the *host number* field, the full IP address is used instead.

2.2.2 IPv6: Internet Protocol Version 6

IPv6 [Deering and Hinden, 1998] is the new version of the Internet Protocol, as a replacement of IPv4. Many reasons have motivated the design of a new version, particularly the lack of available IPv4 addresses[1], better embedded security, provision for *Quality of Service*, the ability to specify additional information in extension headers, native multicast, built-in mobility support, ... For the novice reader, we advice [Huitema, 1998; Cizault, 1999; Lee et al., 1998].

2.2.2.1 IPv6 Header

The IPv6 header is illustrated on fig. 2.1 (b). The IPv4 header is also shown as a matter of comparison for the header size and complexity. The main visible differences between the two versions is the size of addresses. IPv6 addresses are now 128-bits large. As we see on the figure, the IPv6 header is only 16 bytes longer that its former version although each IPv6 address nearly accounts for the size of an IPv4 header. The IPv6 header has in fact been much simplified. The flow label is reserved to be used for *Quality of Service*.

Any number of additional headers (Extension Headers) can be added in order to assist routing or to provide security or more detailed information to the destination node. Currently, the following Extension Headers are defined:

[1] It is important to note that most of the available IPv4 address space id allocated to American institutions and vendors. Then, the lack of IPv4 addresses is not really perceived in the US, while this strikes a need for IPv6 deployment in the rest of the world.

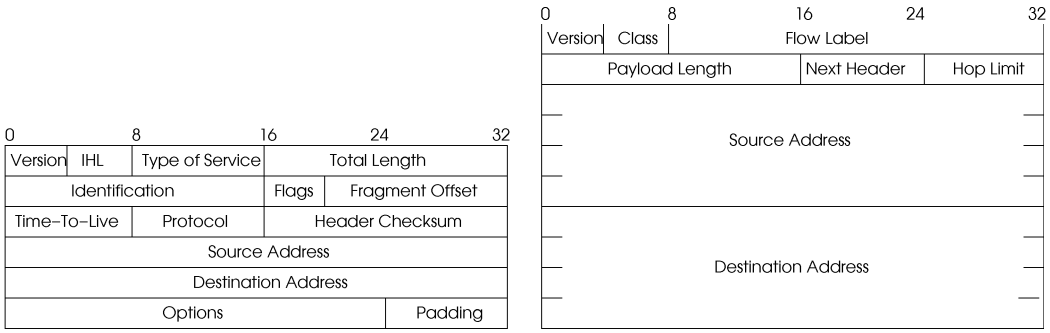

(a) IPv4 Header

(b) IPv6 Header

Figure 2.1: IP Headers

- The `Destination Options Header` carries additional information that is processed by the destination node only. This header is very important to extend `IPv6` functionalities, particularly mobility.
- The `Routing Header` carries a number of addresses that specify a sequence of nodes where the packet must be routed to before actually reaching the final destination (source routing). The destination address specified in the original `IPv6 Header` is the address of the first intermediate router. Once it gets there, the destination address is permuted with the next address in the extension header. The final destination performs security checks based on the initial order of addresses in the `IPv6 Header` and `Routing Header`.
- The `Hop-by-Hop Options Header` carries additional information that must be processed by each intermediate router, like for instance an alert option.
- The `Encryption Security Payload (ESP) Extension Header` carries information that allows the receiver to authenticate the source and decrypt the payload.
- The `Authentication Header (AH)` carries information that allows the receiver to authenticate the source.

`IPv6` also defines *encapsulation* as a means to force a packet to take a different route. This is performed by enclosing the original packet as the payload of a new packet, and by appending a new `IPv6 Header` specifying the new destination. The former header becomes the *inner header* whereas the latter is the *outer header*. The original packet becomes the payload of the new packet and get routed to the destination specified in the outer header. At the destination, the inverse process is performed (*decapsulation*). This mechanism is also called *tunneling*.

2.2.2.2 IPv6 Addressing

`IPv6` addresses are 128-bits large. In theory, `IPv6` offers enough bits to allocate an address to billions of nodes on every square meter on Earth[2], but in practice it is much less due to the hierarchical nature of an address. There is nevertheless enough address to embed an `IPv6` stack in every electric device[3] on Earth. The `IPv6` address format is defined in [Hinden and Deering, 1998]. It allows for several address

[2]It can support up to 2^{128} interfaces, i.e. $340,282,366,920,938,463,463,374,607,431,768,211,456$ addresses, which is approximately $665,570,793,348,866,943,898,599$ addresses per square meter of the surface of the Earth. On the other hand, `IPv4` can "only" support 4.2 billions of hosts

[3]This is not limited to computers, cameras, phones, etc, but also includes washing machines, TVs, micro-waves, fridges, etc

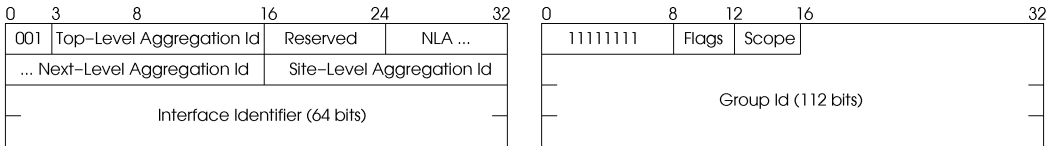

(a) Aggregatable Global Unicast address format

(b) Multicast Address

Figure 2.2: IPv6 Address Formats

architectures to coexist simultaneously. The first three bits determine the address format. A few address formats have been defined so far and new ones can be added at will. Most of the address space is still unassigned, this leaving scope for new and better address architectures.

Aggregatable Global Unicast address format The `Aggregatable Global Unicast Address Format` [Hinden et al., 1998] is the most common address format. It is intended to facilitate scalable `Internet` routing.

- The first set of fields form the `network identifier` and is 64-bits wide. The `network identifier` plays the role of a *topological location dependent identifier*. It is used by routing protocols which aim to deliver datagrams to their destination subnetwork via the most optimal path. Because a flat addressing would require each router to keep an entry for each single subnetwork, the `network identifier` is hierarchically structured so that a single entry in the routing table provides a route for an aggregation of subnetworks. An aggregation of networks have a common network identifier prefix; the deeper in the hierarchy, the longer the prefix length.

- The last 64 bits make the `interface identifier` which is used to identify the interface on a link. Basically, it allows to demultiplex between the nodes on the same subnetwork. This identifier is constructed in `IEEE EUI-64` format [IEEE, 1997] and may or may not be globally unique (provisions are made to construct `IEEE EUI-64` addresses from `IEEE 48bit MAC`; this identifier may have a global scope within the `Internet` provided the 'u' bit is set). If it is made globally unique, the interface can be permanently identified by the sole `interface identifier` whatever subnetwork it is located in, otherwise it has to be used in coordination with the `network identifier`. In common practice the `interface identifier` is a number randomly generated.

This address format is illustrated on fig. 2.2 (a). 48 bits are used for the public topology, the remaining 80 bits belong to the private topology. There are 16 bits in the `Site-Level Aggregation Id` field, which means that any organization can contain up to $65,000$ subnetworks.

Taken individually, none of the identifiers are globally unique. There is therefore a close relationship between the two parts of the address. Ideally, the `IPv6` address would be split in two independent fields. When the `TCP/IP reference model` was first designed, the address space was restricted to 32 bits. This did not provide a sufficiently large address space to physically split addresses into two distinct fields. On the other hand, `IPv6` addresses are large enough. The trend toward a globally unique interface identifier was raised after the `GSE` (`Global, Site and End-System Designator`) address space format proposal [O'Dell, 1998; Crawford et al., 1998] which outlined a need for it but was finally rejected although it provided valuable ideas (see section 3.3.1). This choice has never been made nor favored, although abundantly debated. Indeed, globally unique `interface identifiers` face privacy concerns. The `IP` address is always in the clear; as a result anyone could track the movements of a user by monitoring the source and destination addresses of all packets issued by or intended to the same node.

Multicast address format Internet-wide deployment of the multicast capabilities is ensured by the multicast address format that must be recognized by all routers. It is illustrated in fig 2.2 (b). *Flags* are currently unassigned, but the last one, the *T flag*. The *T (Transient) flag* indicates that the multicast address is not permanent, which is usually the case. Only well-known addressed assigned by a global Internet numbering authority have this flag unset. The *scope* field confines the propagation of the packet to a limited area. The *group id* field uniquely identifies the group.

2.2.3 Neighbor Discovery

Neighbor Discovery [Narten et al., 1998] is the protocol used by IPv6 nodes on the same link to discover the presence of other nodes and their link-layer addresses, to find routers, and to maintain the connection with the neighbor subnetworks. The protocol defines the following messages:

- Neighbor Solicitations: used to determine the link-layer address and to perform Duplicate Address Detection (DAD).
- Neighbor Advertisements: used to respond to Neighbor Solicitations and to advertise a new link-layer address.
- Router Solicitations: used to determine what routers are on-link.
- Router Advertisements used by routers to advertise their presence and also to respond to Router Solicitations. They provide a list of prefixes.

2.2.4 Address Configuration

Each interface is identified by a topologically correct address obtained on its subnetwork by means of either stateless [Thomson and Narten, 1998] or stateful DHCPv6 Address Autoconfiguration [Bound and Perkins, 1999].

Stateless Address Autoconfiguration Addresses are configured by listening to network prefixes advertised by Neighbor Discovery Router Advertisements. The address is the concatenation of the network prefix and a suffix. The suffix could alternatively be the node's MAC address or any bit string automatically generated. Duplicate Address Detection is performed to insure that the autoconfigured address is unique on the subnetwork.

2.2.5 Security

Securing packet exchanges at the network layer is a base requirement for IPv6. IPSec and related protocols must be supported by all implementations. Security at the network layer is mainly concerned with *authentication*, *encryption*, *authorization* and *privacy*. Among other things, the network layer must guarantee that packets are not forged, not eavesdropped, not sent by malicious hosts. To achieve this, the sender of a packet must be authenticated by the receiver (for instance by means of the AH Extension Header that carries information necessary to prove the identity of the sender, or by means of the ESP Extension Header that conveys information necessary to decrypt the packet). As a result from this, packets cannot be sent on behalf of another node, and they cannot be rewritten by intermediate nodes either. There are still a lot of ongoing discussions at the IETF, particularly concerning key exchange mechanisms, and a means to convey authorization information (to ensure that a node is granted permission to perform some operation) seems to be missing.

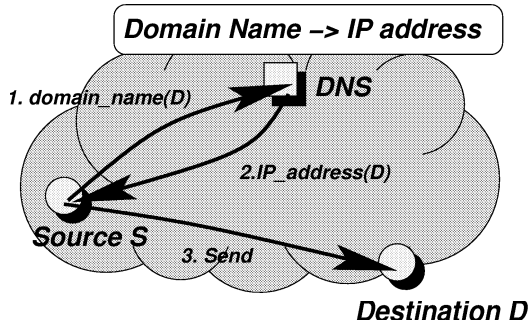

Figure 2.3: Domain Name System definition

2.2.6 DNS

Prior to sending a packet, as requested by the application layer, the source needs the IPv6 address of the destination. This address is usually not known, and must be retrieved in some database. The Domain Name System (DNS) [Mockapetris, 1987a,b] is a distributed and duplicated database designed for the purpose of retrieving addresses. It records bindings between *fully qualified domain names* and IP addresses. The *domain name* is a variable-length ASCII string which globally and uniquely identifies a host in the Internet topology and provides a location independent identifier of the host. Once a host has a pending packet to send, the application layer first checks if it knows the IP address of the destination. If it doesn't have the address, the DNS is queried with the domain name of the destination. The IP address corresponding to the requested domain name is returned by a Name Server. Once the IP address of the destination is found, the network layer at the source constructs the packet and transmits it. This mechanism is illustrated in fig. 2.3. The Name Server replying to the query could be the authoritative repository for this record (Primary Name Server) or a cache (Secondary Name Server). The mapping between the *domain name* and the IP address is cached at the requester for the lifetime specified in the reply. This avoids calling the DNS each time a communication flow is established with the same destination.

2.3 Routing

Routing protocols aim at routing datagrams to the relevant destination node by the most optimal path. The actual forwarding of packets from a sender node to a destination node is based on routes computed by the routing protocols. Each router is required to run at least an instance of a *unicast routing* protocol while running a *multicast routing* protocol is optional. Unicast routing protocols are used to route packets between any tow nodes whereas multicast routing protocols are used to optimize bandwidth consumption when there are multiple destinations for a given packet. Multicast routing is a means of minimizing bandwidth use by sending only one copy of a packet on a particular link when there is more than one recipient reachable through that link. Hence, the aim of multicast routing is to avoid duplicate information flowing over the same link. The sections below first introduce unicast routing before describing *traditional multicast* and then *Small Group Multicast*, an orthogonal and more recent multicast technique. We conclude this section with a comparison of the two multicast techniques.

2.3.1 Unicast Routing

In the unicast model, the purpose of the routing protocol is to update topological changes. It maintains a routing table used to determine the path toward any part of a network. The routing table is computed by a routing algorithm according to some *metrics*. The best route may be determined in terms of minimum cost of delay, bandwidth overload and probability losses and may differ depending on some local policy. Once

an incoming packet arrives, the routing table is searched for a route to the destination as specified in the IP address destination field of the IP header. The routing information in the table is hierarchical and records the next hop toward a host (*host-specific route*) or preferably to a network or set of networks, i.e. a network prefix (*network-specific route*). The table is searched for the longest prefix match and the next hop toward the destination is returned. The packet is then forwarded to the next hop and so on until it reaches the node corresponding to the IP destination address.

2.3.1.1 Unicast Routing Techniques

Routing protocols are essentially classified according to the two techniques currently in use, *Link State (LS)* and *Distance Vectors (DV)*.

Link State Routing Protocols Each router determines its own connectivity which is thereafter flooded to the entire network. Protocols that use this technique are seen in the literature [Huitema, 1995] as very powerful protocols since routes are computed at each node with a total knowledge of the topology, based on (1) the exchange of databases between adjacent routers and (2) the flooding of *Link State Updates*. This is necessary to provide synchronized full topology knowledge to all nodes. However, these two operations are "acknowledged". In the database exchange process, *Database Description* packets sent by one of the routers (*the master*) are acknowledged by the *slave* with responses containing a summary of its link state data. As for *Link State Updates*, they are acknowledged by sending *Link State Acknowledgment* packets back to the sending neighbor.

Distance Vector Routing Protocols This technique is very simple and easy to implement. In this technique, a router E receiving that R can reach X with metric n deduces that E itself can reach X with metric $n + m$, where m is the metric of the RE link in both directions. In turn, E broadcasts a list of *Distance Vectors* to its neighbors all the destinations that can be reached by E.

2.3.1.2 Unicast Routing Protocols

Since the Internet is growing continuously, it is not possible to compute all possible routes. Routing tables wouldn't scale. We therefore distinguish two categories of routing protocols. *Intra-Domain Routing Protocols* compute routes for networks running under the same routing policy (i.e. same domain of administration), whereas the purpose of the *Inter-Domain Routing Protocols*) is to route packets between distinct domains. They advertise reachability for an aggregation of networks that share the same address prefix. They usually rely on policy routing in the sense that the propagation of the routing information is restricted to the domain policy.

Intra-Domain Routing Protocols The two most commonly used protocols are OSPF [Moy, 1997], and RIP [Hedrick, 1988; Malkin, 1994]. OSPF makes use of the *Link State* technique whereas RIP makes use of the *Distance Vector* technique.

Inter-Domain Routing Protocols The most popular protocol that falls into this category is Border Gateway Protocol (BGP) [Rekhter and Li, 1995; Rekhter and Gross, 1995; Stewart, 1998]. It makes use of *Path Vectors*, a technique based on *Distance Vectors*. In this technique, complete paths to each destination are advertised in routing exchanges messages instead of counting to infinity as in the *Distance Vector* technique. Paths are therefore not limited to the infinity value and various metrics can be used. Each router only advertises paths it itself uses similarly to *DV* protocols. This advertisement is performed by means of TCP connections between any two BGP neighbors. Upon reception of a new path, a BGP node

checks whether it is shorter than the path already recorded in its table and only advertises shorter paths. A router E receiving that A can reach E via nodes B, C and D with metric m, deduces that E can reach A with the same metric m.

2.3.2 Traditional Multicast Routing

The traditional concept of multicast relies on the *multicast model*, as defined by Deering [Deering, 1991; Deering and Cheriton, 1990]. In this model, a *multicast address* is assigned to a collection of nodes that form a *multicast group*. A *multicast routing protocol* construct a *multicast delivery tree*. Groups are *open*: the source does not know about members, the source does not need be member and the source only knows the multicast address of the group. Groups are *dynamic*: new members can join and leave at any time and do not need to register or to negotiate their participation with a centralized group management entity. Usually, a group membership protocol is associated with the multicast routing protocol to gather with information about the existence of group recipients for a given multicast group. `IGMP` is the protocol used in `IPv6` for this purpose. It informs a given router that there exist subscribers to a given group on its attached subnetwork. Then, packets sent to the multicast address are duplicated by routers whenever the next hop toward members of the group differ.

Only a minority of the routers actually deployed in the `Internet` are multicast-enabled. Consequently, multicast routing is ensured by the `Mbone`, a virtual multicast network where connectivity between two multicast-enabled routers is ensured by point-to-point tunnels. These routers run the `mrouted` daemon.

We commonly distinguish two kinds of *multicast delivery tree*, the *Shortest Path Tree* (`SPT`), and the *Shared Tree*, or *Core-Based Tree* (`CBT`). The `SPT` is a *minimum spanning tree* rooted at the source. Each source in the group has its own `SPT`. The `CBT` is a single delivery tree built per multicast group, and is shared by all senders in this group. This tree is rooted at a single *core* router.

Multicast protocols are classified in the two following categories:

- *Dense-Mode Protocols*: this category is also known as *broadcast-and-prune* and always use a *Reverse Shortest Path Tree* rooted at the source (source specific `SPT`). Data packets are periodically flooded on the distribution tree, and routers that don't have receivers prune the branch of the tree. Pruning ensures that packets are not transmitted on branches where there are no subscribers. This category performs better when the topology is densely populated by group members since routers are less likely to prune the branch of the tree. Every router keeps state information for every source ((S, G) entries), regardless there actually exists members for the group.
- *Sparse-Mode Protocols*: this category is also known as *explicit-join*. It either uses a `SPT` or a `CBT`. A router acting as a *Rendez-Vous Point* (*RP*) or *core* is used as a meeting place to bring sources and receivers together. Members are expected to send explicit join messages to the *RP*. The source sends data to the *RP* which relays along the multicast distribution tree. This category is more efficient for a few widely distributed group members. Finding an optimal *RP* for the group is a NP-complete problem and requires the knowledge of the whole network topology.

2.3.2.1 Intra-Domain Multicast Routing Protocols

- `Distance Vector Multicast Routing Protocol` (`DVMRP`) [Waitzman et al., 1988] a *Dense-Mode Protocol* based on the `Reverse Path Forwarding (RPF)` algorithm. The multicast tree is a `Reverse Shortest Path Tree` created using `broadcast-and-prune`. The source broadcast the packet and routers perform a `RPF` check in order to see if the packet was routed from the shortest path from the source. If so the router forwards the packet to all its neighbors unless they receive an explicit prune from their neighbor down the tree. Otherwise, the packet is discarded. Leaf routers check for the existence of members on their attached subnetworks by means of `IGMP`.

If there is no members, they send a prune message toward the source. The broadcast-and-prune is repeated periodically.

- Core-Based Tree (CBT) [Ballardie et al., 1993] is a *Sparse-Mode protocol*. As it name stands for, it makes use of a single *Core-Based Tree* rooted at a *core*. The source sends the data to the core and the members send explicit join messages to the core. The multicast distribution tree is bidirectional. This is more efficient when packets from the source cross the branches of the tree. In this case, packets are not only sent up to the core, but also down the tree. However, this also adds more complexity. In practice, only a few vendors support CBT.

- PIM-SM [Estrin et al., 1998]: a group has only a single *RP* and share a single shared tree rooted at the *RP*. The *RP* must be discovered by all routers, using a bootstrap protocol (a bootstrap protocol is included in version 2), that also provides robustness on case of failure of the *RP*. Members send explicit join messages to the *RP*. As a result of these messages, forwarding state is created in each router between the member and the *RP*. The source encapsulates data to the *RP* where the encapsulation header is stripped off the packet. Packets are then forwarded along the shared tree. If there are no forwarding state, the *RP* sends a message (register stop) to the source. The overhead of the encapsulation can be avoided by establishing forwarding state between the source and the *RP*. A particularity of this protocol is the ability to switch from a shared tree to a shortest path tree.

- PIM-DM [Deering et al., 1994] is very similar to DVMRP, with two major differences. First, PIM-DM uses the routing table to perform Reverse Path Forwarding checks, and is independent of the algorithm used to build the routing table. Second, PIM-DM forwards packets on all its interfaces. Neighbor routers on the reverse path must then prune when the Reverse Path Forwarding check fails. This diminish complexity of the protocol.

- MOSPF [Moy, 1994, 1997] is a *Dense-Mode Protocol*. As its name stands for, it is built on top of OSPF and makes use of its unicast routing table to build the multicast tree.

2.3.2.2 Inter-Domain Multicast Routing Protocols

Many papers that describe *inter-domain multicast protocol* make use of the wording *wide-area multicast routing*. However, the above cited protocols don't actually provide means for multicast group to span several domains and don't scale to a wide-area network. Indeed, *inter-domain multicast routing* is still an active research issue and no protocol has been standardized so far at the IETF. *Inter-domain multicast routing* is technically very complex and functionalities, like security and group membership, are not yet available. This is discussed in an overview paper from Almeroth [Almeroth, 2000] and in other papers [Ramalho, 2000]. In order to give an immediate solution to an immediate need, the IETF is therefore considering both a near-term solution, that does not scale to a large number of multicast sources, and a long-term solution which tries to address the remaining issues.

Near-Term Solution The near-term solution is a straight forward solution based on PIM-SM to establish a multicast tree between domains containing group members. PIM-SM is assisted by two additional control protocols:

- MBGP (Multiprotocol extensions to BGP-4) [Bates et al., 1998] extends BGP (see section 2.3.1) to carry multicast routes in BGP. Each router only knows the topology of its own domain and the paths to reach each of the other domain. A domain advertises reachability for multicast with a message that specifies the list of domains to which the domain has a path. No information about the multicast group is carried in these messages. This information is only used when a *join* message is sent form an *RP* or router toward the source.

- MSDP (Multicast Source Discovery Protocol) is used to inform a *RP* in a domain that there are sources in other domains. The protocol operates over a TCP connection between RPs

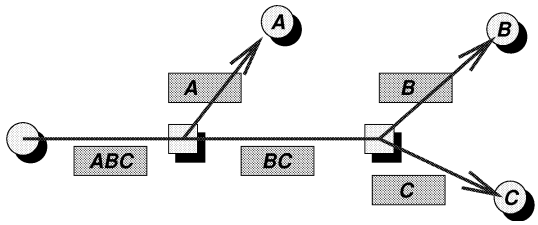

Figure 2.4: Small Group Multicast Forwarding

in different domains. Representatives in each domain announce to other domains the existence of active sources. New sources register with the domain's *RP*. MSDP in the domain detects the new source and advertises it to all directly connected MSDP peers.

Long-Term Solution Work on the long-term solution could be classified in two distinct groups: solutions based on the standard IP multicast model, and solutions that look to change this model. Since this is still a work in progress, we will limit ourselves to a simple citation to the most advance work: BGMP (Border Gateway Multicast Protocol [Kumar et al., 1998] - this protocol constructs a bidirectional shared tree between domains using a single root. It needs a strict address allocation scheme), MASC (Multicast Address-Set Claim) , GLOP (static allocation of multicast addresses to each AS), RAMA (Root Addressed Multicast Architecture [Ballardie et al., 1999] - from the observation that most multicast applications are single-source or have an easily identifiable primary source, this source is chosen as the root of the tree. This prevents the complexity of the core placement), Express Multicast [Holbrook and Cheriton, 1999], and Simple Multicast [Perlman et al., 1999].

2.3.3 Small Group Multicast (or Explicit Multicast)

Small Group Multicast (SGM) [Boivie, 1999; Boivie et al., 2000b], also known as *Explicit Multicast* [Ooms et al., 2000] is a new multicast technique designed to complement traditional multicast. The basic idea is to record the list of recipients of a packet in the data packet itself. This mechanism is illustrated on fig.2.4. A packet is sent to a group formed by recipients *A, B and C*. The first router reads the list of recipients and consequently duplicates the packet and transmits an instance of it on the interface toward A and one instance on the interface toward B and C.

The operation of routers is indeed very similar to unicast. The forwarding table is checked for the next hop toward any of the destinations contained in the header and the packet is duplicated as many times as there are distinct next hops. Hence, in this approach, there is no notion of group address and group membership. Further, there is no requirement for a multicast routing protocol. This contrasts with traditional multicast where the data packet is sent to a group address that identifies all the recipients.

As an attempt to standardize an initial set of propositions made at the IETF [Boivie et al., 2000a; Ooms et al., 2000; Imai, 2000], a BOF took place at the 48th IETF meeting (Pittsburgh, August 2000), but didn't lead to a new IETF working group. A taxonomy of proposals can be found in [Ooms, 2000], whereas [Braun, 2000] overviews the different techniques. [Boivie et al., 2001] is a more up-to-date document and proposes a basic specification.

2.3.4 Traditional Multicast versus Small Group Multicast

The properties of *traditional multicast* and *small group multicast* are summarized in tab.2.1. The intuitive comparison between the two techniques shows that *small group multicast* seems more appropriate for a large

Metric	Traditional Multicast	Small Group Multicast
Group Members	unknown to the source	known to the source
Identifier	group of nodes is identified by a single address	no group identifier
Membership Management	a protocol is needed	no group membership management
Distribution Tree	a multicast routing protocol is responsive for distribution tree establishment	no need to build a tree, use standard unicast routing table
Multicast Address	multicast address discovery	no multicast address discovery
Packet Overhead	no overhead	list of group members self contained in the packet
Packet Processing	check for one forwarding interface	check for n forwarding interfaces
Multicast Signaling	large amount of signaling	no signaling
Multicast Memory	State in each router	no state in the routers
Scalability	large number of group members	small number of group members

Table 2.1: Multicast Properties: Traditional Multicast vs Small Group Multicast

number of multicast groups with a short number of members, whereas *traditional multicast* is more appropriate for a large number of group members. Both techniques are indeed complementary to one another since a "one size fits all protocol seems unable to meet the requirements of all applications". Applications of *small group multicast* include *narrowcast-like* (or *few-to-few*) applications (*IP telephony*, collaborative applications), whereas traditional multicast is targeted to *broadcast-like* (or *one-to-many*) applications (TV and radio programs, weather forecast, ...).

2.4 TCP/IP Addressing and Mobility

Although `IPv6` allows for improvements and upgrading and should in theory handle mobility well, mobility support still rather look like an eyesore on top of a well-designed framework. This comes from the fact that the design of the `TCP/IP` protocol suite has been based on the assumption that end-nodes are fixed. In the following sections, we demonstrate that mobility breaks down the `TCP/IP` *reference model*. We particularly explain the impact of mobility on the `TCP/IP` addressing scheme and discuss why the role played by `IP` addresses prevents easy and efficient support of mobility. From this discussion, we are able to list a number of services required to achieve mobility.

2.4.1 Mobility Paradigm

Intuitively, there seems to be two main issues when a node changes its topological point of attachment.

1. how to determine the current topological point of attachment of a mobile node: *locating*.
2. how to route packets to the current topological point of attachment of the mobile node: *routing*

However, once we get in the context of TCP/IP, we see that the problem caused by mobility it is not only a question of *locating* and *routing*, which would merely be the task of the network layer, but also a question of *addressing*. Indeed, the TCP/IP addressing design choices don't facilitate the network layer's ability to support mobility. First, the IP address of a node is bound to the topological location of that node. Second, the IP address has a dual semantic and is used at different layers. As a result from this, there is a trade-off between retaining the IP address which fails routing and changing the address which breaks upper layer connections [Bhagwat et al., 1996; Ioannidis et al., 1991]. This is developed in the following sections.

2.4.1.1 Effect of Mobility on TCP/IP Addressing

Fixed nodes are permanently attached in the same subnetwork and are identified by a permanent IP address which determines the subnetwork where they are attached to. Unlike fixed nodes, mobile nodes changes their point of attachment in the Internet topology. They are moving from subnetwork to subnetwork and are reachable at different locations in the Internet topology. As far as the IP-layer is concerned, a number of considerations drive a change of the IP address when a mobile node attaches to a new subnetwork. When a host is moving, its address no longer reflects the current point of attachment. Since conventional routing is based on the IP address, retaining the IP address would result in routing the packet to the visited link where the mobile node has obtained this address. As a result of the displacement of a node from one subnetwork to another, packets cannot reach that node. As far as security is concerned, routers are requested to filter packets so that source and destination addresses in the IP header are compliant with the address space served by the router. Packets with a topologically incorrect source address are discarded (*ingress filtering*). Basically, a node using an address got from another subnetwork would never get any datagrams, and never succeed to send any.

IP-layer mobility therefore implies two things:

1. A change of topological location: packets intended to a particular mobile node are routed via a different path over time. This requires to find out a new path toward the current topological location of the mobile node.

2. A change of the routable IP address: any node must be identified by a topologically correct IP address that determines the current subnetwork where the interface is attached to. A node that changes its point of attachment is therefore identified by distinct IP addresses.

As seen in section 2.2.2.2, the address format currently in use in IPv6 is Aggregatable Global Unicast Address Format. In this address format, addresses are made of the network identifier and the interface identifier fields. Actually, there are reasons for changing both fields when the mobile node enters a new subnetwork:

1. The network identifier field is used for routing to the current subnetwork. It must identify the current subnetwork where resides the node, otherwise routing protocols can't route packets to their destination.

2. The interface identifier field is used to identify the node on the subnetwork. The MAC address of the node may optionally be used for this, but nothing guarantee that the MAC address is itself globally unique. Moreover, another node on the subnetwork may already be allocated the same interface identifier. Thus, the mobile node may not be able to retain the same interface identifier even if it wishes so. In addition to this, the use of a permanent interface identifier is not recommended for a mobile node due to privacy concerns.

2.4.1.2 Effect of Address Change

In section 2.1.2 we have depicted that the IP address has a dual semantic and serves both to identify the topological location of an interface and to identify the interface itself. Moreover, those two functions are embedded in an address meaningful at two different layers. At the network layer it is principally used as a location dependent identifier. At the transport layer, it is used as a permanent node identifier.

As we take a look at the UDP or TCP headers, we see that these protocols make use of the IP address solely as a *permanent and invariant node identifier* to identify peers of the communication flow. For historic reasons, TCP/IP makes the implicit assumption that nodes always resides in the same subnetwork. At the time of this design choice, mobility was not perceived, nodes where not assumed to migrate between subnetworks, and the IP address was not assumed to change during a connection's lifetime. As mobility in the Internet implies a change of IP address, any change of location at the network layer is perceived by both network and transport layers. This breaks TCP connections since the protocol cannot authenticate an incoming packet as coming from the same communication flow while subsequent communications get rejected. As a result, the IP address cannot be seen as a location invariant and independent identifier at the transport layer anymore. This clearly violates the OSI layer independence concept which aims to process design changes at each layer without impact on other layers.

2.4.1.3 Conclusion

This chapter shows that IP addresses are bound to topological locations and that nodes that change their topological location must also change their IP addresses. The coordination of the dual IP address semantic and its dual use at different layers fails in offering a location independent and invariant node identifier. Not only the dual semantic and use of the IP address breaks the OSI Reference Model, but it also prevents mobility. For short, connections are broken because TCP/IP violates the layer concept and not because hosts are moving. We advocate that mobility cannot be supported efficiently while the change of the IP address has an impact on upper layer protocols.

We conclude that the TCP/IP addressing scheme is leading to serious limitations on TCP/IP's ability to support mobility. Upper layers must make use of a location independent identifiers while mobility support is mainly an IP-layer issue best achieved at the IP-layer, and transparently to the transport layer [Bhagwat et al., 1996; Ioannidis et al., 1991].

2.4.2 Mobility Support Services

From this discussion, it is clear that mobility cannot be achieved without the help of specific support services, what we term *mobility support*. Without mobility support, all on-going connections are broken as a result of the migration from one point of attachment to the other. The purpose of mobility support is thus to provide permanent and uninterrupted Internet access to mobile nodes and to route packets optimally between any two given nodes when one or both nodes are mobile nodes. The aim of mobility support could be refined as to provide efficient *addressing*, *locating* and *routing* when the destination of the packet changes its topological location. This allows us the define the following mobility support services. Note that we do not advocate where and how these services should be implemented. This will be discussed in chapter 4, after we have presented the existing mobility support schemes.

First of all, a terminology is needed to distinguish the node one wants to talk to, and the location of that node. We therefore introduce the following two *abstract* terms that we will extensively use in this study:

- node identifier: it permanently, invariantly and uniquely identifies the node[4]. It is assigned to the

[4] *node* is not to be confused with *interface*, which has a very specific meaning in the TCP/IP addressing terminology (sections 2.1.2

node itself and must not change if the node changes its point of attachment in the Internet topology.

- location identifier: it identifies the location of the node in the Internet topology, and is used to route packets to the current point of attachment of the node.

By using the terms node identifier and location identifier, we do not assume any underlying addressing scheme nor mechanism to bind the node identifier to a location identifier. Thus, a mobile node may have at least one node identifier and one location identifier. In order to determine the topological location of a mobile node, we first need to bind a permanent and invariant node identifier to a location identifier. Then, we need to retrieve this binding and to actually deliver packets to the current point of attachment. Therefrom, mobility support can be defined as a set of three additional network services in charge of updating the topological location of a mobile (Location Update), locating the mobile node (Location Lookup) and optimally routing packets to its current point of attachment:

Location Update: establishes bindings between the node identifier and the location identifier. In practice, Location Update tracks the succession of subnetworks visited by a mobile and therefore the succession of transient IP routing addresses acquired by the mobile node.

Location Lookup: retrieves the location identifier corresponding to a node identifier. In practice, Location Lookup finds out the current topological location of the mobile node in the topology, i.e. the IP address that identifies its current point of attachment.

Routing: determines which route must be taken by a packet given the location identifier of the destination. It aims to determines the most optimal path between the source and the destination.

and 2.2.2.2), nor to the *host number* defined in section 2.2.1. The transport layer does not care about the interface used to reach a particular node although we said that IP addresses are allocated to interfaces. We note that the selection of the right interface as the source or destination of packets also brings some interesting issues as debated in the IPv6 working group, but this is not pertaining to mobility

Chapter 3

Mobility Support: State of the Art

This chapter presents a number of *mobility support* schemes that achieve the services highlighted in section 2.4.2. They can all fit in a few distinct frameworks as this will be depicted in chapter 4. We begin our study with the official IETF standard or work in progress, namely namely Mobile IPv4, Mobile IPv6, and Hierarchical Mobile IPv6. Other proposals are more or less detailed according to the available information and their relevance to this present study. Some non-IP propositions are outlined since IP proposals are also made on the experience gained from other protocol suites. However, due to lack of space and time, it is impossible to list all proposed schemes, all the more since new ones appear on a regular pace. Some propositions made recently at the IETF or in conference papers are therefore missing from this study and we apologize for this. We have also left out mobility management in cellular telephony, like GSM, although a comparison between *circuit-switched networks* and *packet-switched networks* would have been interesting. This probably deserves a *State of the Art* of its own. The second section lists a number of addressing schemes that worth mentioning.

3.1 IETF Mobility Support Schemes

3.1.1 IETF Mobile IP

Mobile IP is the official IETF standard for *host mobility support*. It is developed in the Mobile IP working group for both IPv4 and IPv6. The first section describes features common to IPv4 and IPv6, and then we detail the protocols.

3.1.1.1 Mobile IP fundamentals

Mobile IP can be seen as a sub-layer that provides additional services between the network and transport layers. It introduces *two-tier addressing* as the solution to the conflicting dual semantic and use of IP addresses depicted in section 2.4.1. *Two-tier addressing* associates a mobile node with two distinct addresses, a permanent home address, and a temporary careof address. An address translation mechanism offers migration transparency to upper layers and insures backward compatibility with transport protocols. Connections are not disrupted as a result of mobility. This solves the question of mobility without changing the mobile node's IP address.

The home address is obtained on a link in the home network (home link) and serves as a location invariant node identifier. It is configured with the home prefix. The careof address is obtained on the link in the visited network (foreign link) and serves as a location identifier, i.e. a routing directive which reflects the

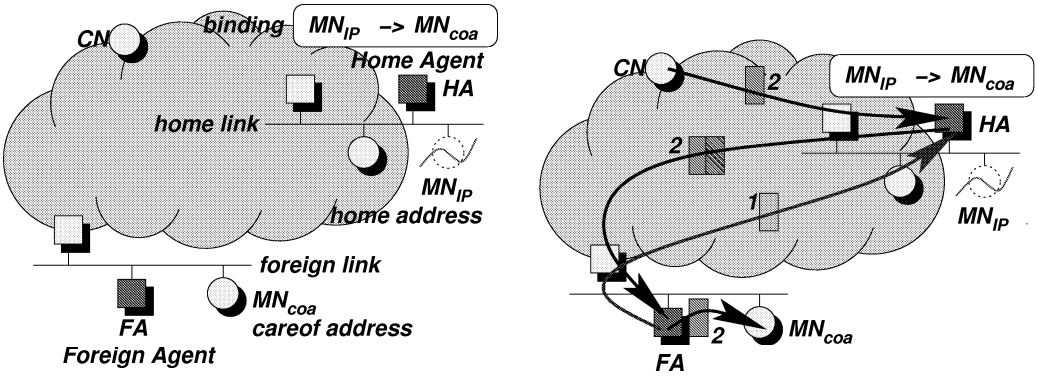

(a) Mobile IPv4 Terminology

(b) Registration with the HA and Triangle Routing

Figure 3.1: Mobile IPv4

current point of attachment to the Internet. It is configured with the foreign prefix. This terminology is summarized on fig. 3.1. The binding between the home address MN_{ip} and the careof address MN_{coa} is registered with the home agent (HA), a special router[1] on the home link able to intercept packets intended to the MN. A correspondent node willing to communicate with a mobile node first calls the DNS which returns the home address of the mobile node. Packets are then routed to the home link where they are intercepted and encapsulated by the HA to the careof address.

3.1.1.2 Mobile IPv4

Mobile IPv4 (RFC2002) [Perkins, 1996a] is the official IETF standard to support mobility in IPv4. It was designed from the experience gained by other IPv4 mobility proposals, particularly [Johnson, 1993]. When roaming, the MN detects its movement by listening to *agent advertisements* sent by the foreign agent (a dedicated Mobile IPv4 access router on each foreign link). When it attaches to a new foreign link, the MN first obtains a new careof address. This careof address can alternatively be a co-located address (i.e. this address is obtained through DHCP) or a forwarding address (i.e. this address of the foreign agent) Then, a Registration Request containing the binding between the permanent and the temporary addresses is sent to the HA. The HA acknowledges with a Registration Reply, and records the binding in a table (Binding Cache). There is no routing optimization in this RFC, so packets sent by CNs always get routed to the home link of the MN where they are intercepted by the HA. The HA performs a lookup in its Binding Cache and encapsulates the packets to the MN's careof address. The packet is whether decapsulated by the foreign agent or the mobile node itself.

3.1.1.3 Mobile IPv4 and Mobile Networks

A very brief section in the Mobile IPv4 specification proposes a solution to support single mobile IP-subnets as standard mobile nodes (see RFC2002 [Perkins, 1996a] section 4.5). A commercial implementation of this has been announced very recently by Cisco Systems [Leyden, 2001]. The mobile IP-subnet is no more than a subnetwork attached to a mobile router MR. The MR performs Mobile IPv4. It has a permanent home address on its home link and gets a new careof address on each subsequent foreign link where it attaches. As a usual mobile node, a Registration Request is sent to MR's home agent (fig.3.2) to instruct it to intercept and tunnels packets to its careof address.

[1] The HA is not necessarily a router, but since it forwards traffic not intended to itself, it could actually be considered as a router

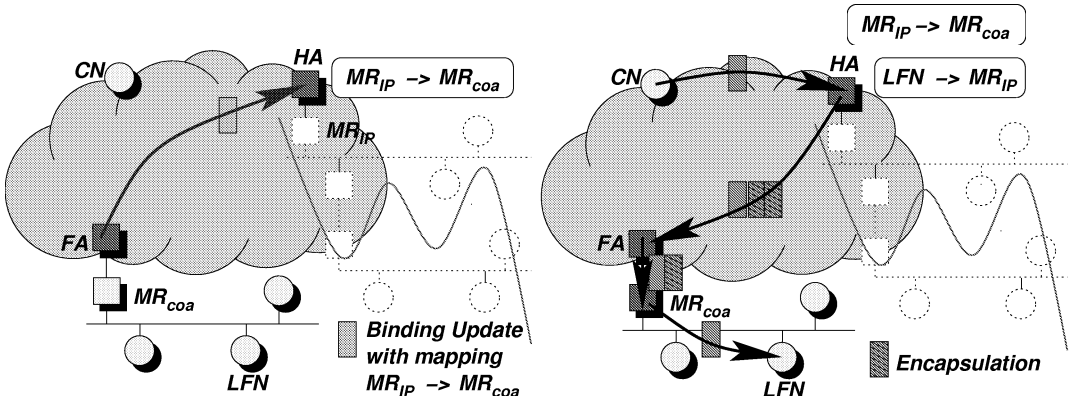

(a) Registration (b) Data Transmission

Figure 3.2: Mobile Networks in Mobile IPv4

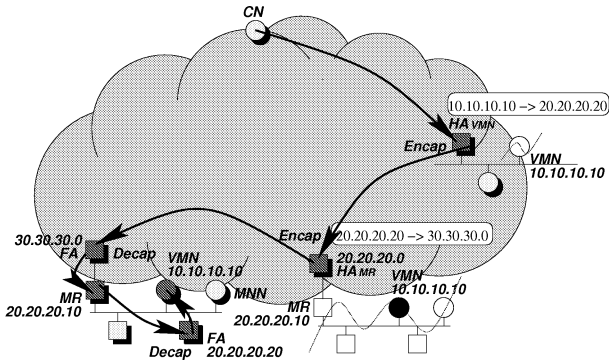

Figure 3.3: VMN in Mobile IPv4

Connectivity In order to intercept packets intended to LFNs[2], two means are suggested, but not detailed. In the first one, the HA is configured with a permanent registration for each LFN that indicates MR's home address as the LFN's careof address. Datagrams sent by CNs are intercepted by the HA and encapsulated to the careof address of the mobile IP-subnet where it is decapsulated by the FA and forwarded back to the LFN. In the second one, `Internet` access to the mobile network is advertised by the MR through a bi-directional tunnel using normal `IP` protocols.

Nested Mobility In the scenario illustrated on figure 3.3, a visiting mobile node VMN (see definition in section 1.3) enters a mobile IP-subnet. The VMN operates `Mobile IPv4` as usual mobile nodes. VMN obtains a careof address VMN_{coa} from a router serving as a FA in the mobile network and registers it with its HA. This careof address is configured with the mobile network prefix. Datagrams sent by CN are routed to the home address VMN_{ip} and then encapsulated by the VMN's HA to the VMN_{coa}. If the mobile IP-subnet has moved, datagrams are intercepted again, this time by the HA serving the MR, and encapsulated to its careof address MR_{coa}. The FA serving the MR decapsulates the datagram and forwards it to the MN_{coa} where it is decapsulated by the FA serving the VMN. As we note, triangle routing occurs two times.

[2] see our terminology in section 1.3

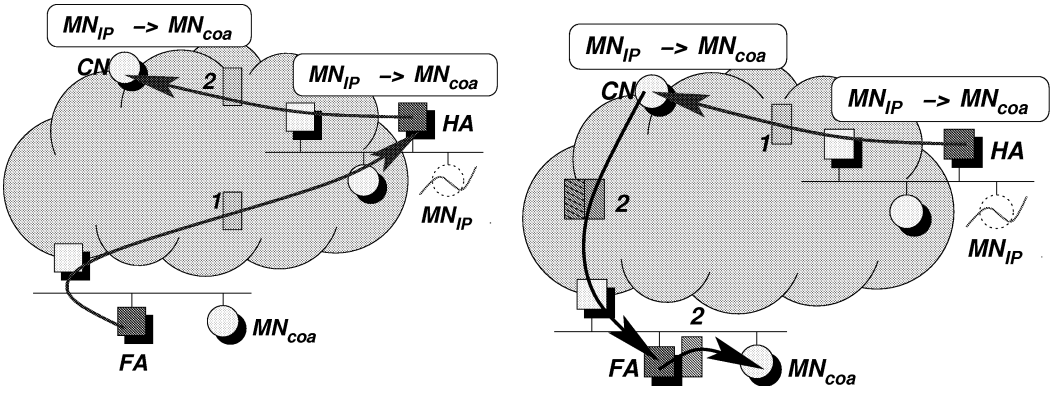

(a) Registration with the HA and the CN

(b) Optimal Routing

Figure 3.4: Mobile IPv4 with Route Optimization

3.1.1.4 Mobile IPv4 with Route Optimization

`Mobile IPv4 with Routing Optimization` [Perkins and Johnson, 2000] (fig. 3.4) is a work in progress that proposes extensions to `Mobile IPv4` in order to avoid triangular routing via the HA. The Binding Update (BU) is a new message that advertises the MN's careof address to the CN. BUs are sent to CNs by the HA on behalf of the MN. In order to determine if a BU should be sent to a CN, the HA monitors incoming packets intended to the MN. CNs may also directly inquire for a BU by sending a Binding Request to the MN's home address. Following packets can then be encapsulated directly from the CN to the MN's careof address, therefore bypassing the HA.

3.1.1.5 Mobile IPv6

`Mobile IPv6` [Johnson and Perkins, 2000; Perkins and Johnson, 1996] (fig.3.5) is adapted from `Mobile IPv4 with Routing Optimization` and takes advantage of the enhanced features of `IPv6` over `IPv4`. It is still a work in progress but should become an `IETF Proposed Standard` in a short future, when security issues are solved. Although it is not yet standardized, every `IPv6` node is in principle required to implement `Mobile IPv6`, thus ensuring wide support of mobility.

`Mobile IPv6` defines two `Destination Extension Header Options`: the Home Address Option and the Binding Update Option. When roaming, the MN detects its movement and obtains a new careof address MN_{coa} on each subsequent foreign link it visits. The careof address is obtained using either `stateless` [Thomson and Narten, 1998] or `stateful DHCPv6 Address Autoconfiguration` [Bound and Perkins, 1999]. The MN may own several careof addresses at anytime, one of which is selected as the primary careof address.

The registration of the binding between its home address MN_{ip} and the primary careof address MN_{coa} is performed by means of a Binding Update (BU) message. The BU is a datagram that contains a Binding Update Option which records the MN_{coa} (unless it is already specified in the `IP` header source address field), and a Home Address Option which specifies the MN_{ip}. All packets carrying a Binding Update Option must also contain an `AH` [Kent and Atkinson, 1998a] or an `ESP Extension Header` [Kent and Atkinson, 1998b] used for authentication[3]. In order to bypass ingress filtering, the source address of packets emitted by the MN is usually set to the MN_{coa} while the MN_{ip} is inserted in a Home Address Option of the `Destination Extension Header`.

[3] Authentication and other security aspects of `Mobile IPv6` are currently under revision

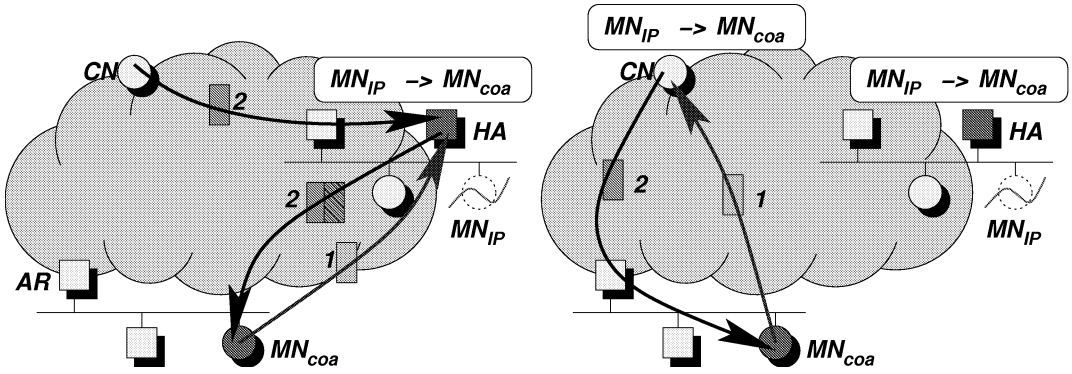

(a) Registration with the HA and Triangle Routing

(b) Registration with the CN and Optimal routing

Figure 3.5: Mobile IPv6

Once it receives a valid BU, the home agent records in its Binding Cache the binding between the home address and the careof address. This home address is used as the key for searching the Binding Cache. As a result of this registration, the home agent adds a *host-specific route* for the mobile node's home address (i.e. for a 128-bit `IPv6` address) via its careof address through a tunnel. Then, the home agent uses "gratuitous" `Neighbor Advertisement` messages [Narten et al., 1998] to intercept all datagrams intended for the MN and encapsulates them to the current careof address (fig. 3.5 (a)).

At this point, the MN may also send a BU containing its primary MN_{coa} to some or all CNs recorded in its Binding List to avoid triangle routing via the HA (fig.3.5 (b)). The CN authenticates the packet by means of the `AH or ESP Extension Header`. Forthcoming packets are directly sent to the MN_{coa} using an `IPv6 Routing Extension Header` containing the MN_{ip}. BUs could be piggybacked in payload datagrams or sent alone in separate packets containing no payload.

BUs are resent periodically whether or not the MN sends or receives any actual traffic. Though, the MN must not send BUs more frequently than one per second. Typically, the MN sends 5 consecutive BUs at this rate just after forming a new careof address, if it is going to be used as the primary careof address. This ensures quick update of the Binding Caches and avoids packets to be sent to the former point of attachment in case some BUs get lost. After these 5 consecutive BUs, the MN may keep sending BUs, but at a lower rate (typically every 10 seconds) in order to refresh the Binding Caches.

3.1.1.6 Conclusion

A fundamental difference between `Mobile IPv4` and `Mobile IPv6` is the routing optimization between the CN and the MN. Optimal routing from CNs to MNs is made possible by sending the primary careof address to the CNs by means of a BU. This topic is absent from `Mobile IPv4;` in `Mobile IPv4 with Routing Optimization,` BUs are sent by the HA whereas in `Mobile IPv6` BUs are directly sent by the MNs. It also worth mentioning that `Mobile IPv4` messages are `UDP` packets whereas `Mobile IPv6` messages are `IP` packets. Also, there is no foreign agent in `Mobile IPv6` unlike `Mobile IPv4,` and the careof address is always a co-located address.

3.1.2 IETF Hierarchical Mobile IPv6

`Hierarchical Mobile IPv6` [Soliman et al., 2001] is a recent `IETF` work in progress in the `Mobile IP` working group. It extends `Mobile IPv6` and separates Local-Area Mobility from Wide-Area

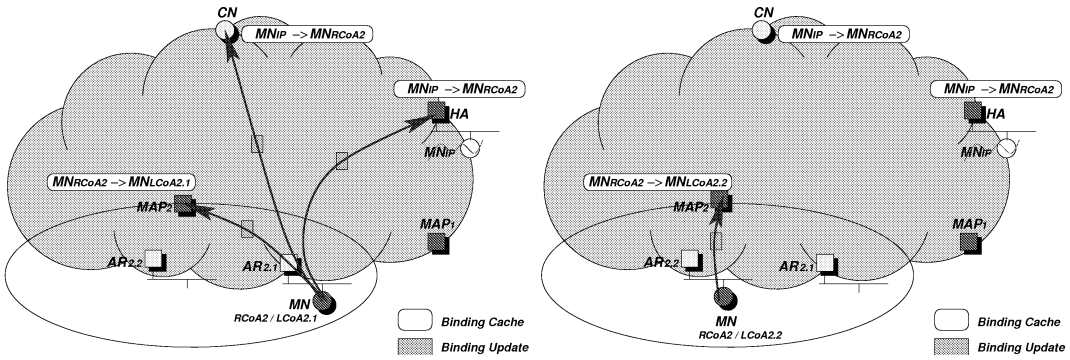

(a) Wide-Area Mobility: MN enters in site 2 (b) Local-Area Mobility MN moves within site 2

Figure 3.6: Hierarchical Mobile IPv6 Basic Mode

Mobility. The main benefit of this proposal is to render Local-Area Mobility transparent to CNs and to limit `Mobile IPv6` signaling in the backbone. This work is based on some former work developed at `INRIA` as early as in 1997 (which will be presented in section 3.2.10). `Hierarchical Mobile IPv6` introduces a new entity, the Mobility Anchor Point (MAP), which is an enhanced HA. A MAP is servicing a domain and receives all packets intended for mobile nodes located in its area of administration. The specification proposes two modes of operation, the `Basic Mode` and the `Extended Mode`.

3.1.2.1 Basic Mode

A MN that performs `Basic Mode` has two careof addresses. The regional careof address (*RCoA*) is received from the MAP (i.e. the *RCoA* is a forwarding address on the MAP's subnetwork; it's not a topologically correct address for the MN) and is kept as long as the MN remains located in the same administrative domain. The MN also gets a local careof address (*LCoA*) on each visited link. The MN establishes the binding between the current *RCoA* and the *LCoA* with the MAP which acts as a kind of local HA. The MN also registers the binding between its home address and the *RCoA* with its HA and CNs. All packets intended to the MN are therefore sent to the *RCoA* using a `Routing Extension Header`. Packets get to the MAP's subnetwork where they are encapsulated by the MAP to the current *LCoA*.

The registration is illustrated on fig. 3.6. As we see, Local-Area Mobility within the site is transparent to the HA and CNs. Local-Area Mobility is only perceived by the MAP which keeps and up-to-date entry between the *RCoA* and the current *LCoA*. As in `Mobile IPv6`, BUs must be sent periodically to the HA to refresh the binding between its home address and its *RCoA*.

3.1.2.2 Extended Mode

The recent `Extended Mode` work in `Hierarchal Mobile IPv6` is seen as a solution to support visiting mobile nodes (see definition in section 1.3). In this case, a hierarchy of MAPs is deployed. There is a MAP in the visited domain, and the MR is acting as the MAP for nodes visiting the mobile network. The `Extended Mode` provides a topologically correct address to the VMN when it enters a mobile network. The MR, as a mobile node, performs `Basic Mode` and obtains a *RCoA* from the MAP in the visited domain and a *LCoA* on each visited link. As a MAP, it advertises its *LCoA* in the MAP Option. A VMN that enters the mobile network obtains a local careof address *LCoA* on the visited link and listens to MAP advertisements. It uses the MAP's current local careof address as its *RCoA*. The VMN first registers the binding between its home address and its *LCoA* with its MAP (MR), and then registers the binding between its home address and its *RCoA* (behind the MAP's local careof

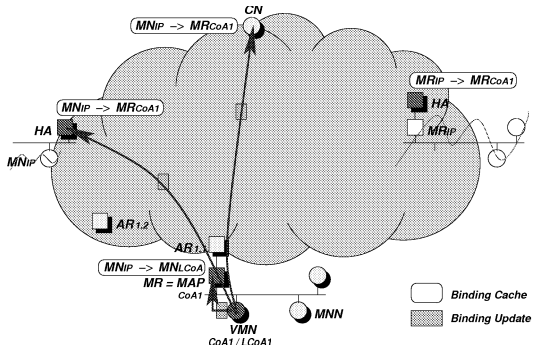

Figure 3.7: HMIPv6 Extended Mode - support for mobile networks (VMNs)

address) with its own HA and correspondent nodes. This is illustrated on fig.3.7.

3.2 Other Mobility Support Schemes

3.2.1 Sunshine and Postel (1980)

The question of mobility has been mentioned as early as in the eighties in [Sunshine and Postel, 1980] where mobility is considered in the ARPA Internet environment. The proposed scheme is based on two methods: an extended interpretation of the network address and the use of source routing. The MN is named directly from a set of reserved network numbers which provides a recognizable address, used as a permanent node identifier. In addition, the MN uses the address of a *forwarder* in the visited network as the routing directive. The MN maintains the binding between the two addresses in some globally centralized or distributed database. For greater efficiency, the CN and the previous forwarder could also be notified the forwarding address. In this case, the CN that wants to communicate with the MN queries the database asking for the forwarding address of the MN given its name. Subsequent packets are then directly sent to the MN via the forwarding address by means of *source routing*.

3.2.2 LSR

Loose Source Routing from Carnegie Mellon University and IBM is one of the initial IPv4 proposals made by Johnson [Johnson, 1993] and then kept up by Perkins. This scheme aims in minimizing changes required to the existing IPv4 implementations. It makes use of the IPv4 Loose Source Routing and Record IP option (LSRR), an option that causes the packet to be routed through a series of intermediate routers.

This option is used instead of tunneling and allows the MN to retain the address obtained on the home link, used as a node identifier. When the MN moves to a visited link, it registers itself with the access router on the visited link (*foreign gateway* or FG) and uses the address of this FG as its new location identifier. The MN then notifies this address to the access router on the home link (*home gateway* or HG) and also deregisters from the previous FG. Initial packets sent from the CNs are routed to the native network. The HG inserts a LSRR option in this packet which causes it to be redirected to the current FG. The HG also provides the location identifier of the MN to the CN, where it is cached, therefore insuring optimal routing between the CN and the MN for subsequent packets. Once a packet gets to the visited link, it is simply forwarded by the FG on this visited link. In the event that ongoing packets are sent to the previous FG, the LSRR option is processed by the FG which causes them to be redirected to the home link.

This scheme is not actually applicable due to a number of security concerns, mainly because the header of

the packet is modified by intermediate routers. Nevertheless, this proposal served as a foundation for the Mobile IPv4 specification.

3.2.3 VIP Sony

Virtual Internet Protocol (VIP) from Sony [Teraoka et al., 1991, 1992; Teraoka and Tokoro, 1993; Teraoka, 1993; Teraoka et al., 1994] is one of the first mobility support schemes targeted to IPv4. It aims to realize *host migration transparency*. The authors introduce the notion of *virtual network* as a sub-layer between the network and transport layers, and the *propagating cache method*. The virtual network is a logical network that hides the architecture of the physical underlying network. A host that migrates in the physical network does not migrate in the virtual network.

A MN has two addresses: a permanent *virtual address* (or *VIP address*) taken from its *native subnetwork*, used as a location independent and invariant node identifier, and a temporary *physical address*, a location identifier used as a routing directive. After migration, the MN obtains a new *physical address* and sends the binding between the two addresses to its native subnetwork where it is broadcast to all hosts. Each packet sent by the MN contains the *VIP address* recorded in a header inserted between the IP header and the transport header as an IP Option.

In the *propagating cache method*, every node in the network has a cache called the *Address Mapping Table (AMT)* that records the binding between the two addresses. This cache is used for address conversion. The *AMT* is updated by snooping on the header and the option of data and control packets sent by the MN as they traverse the network (lazy update). A CN willing to communicate with MN does not know its physical address. Packets are therefore sent to the *VIP address* toward the native subnetwork. The router handling the native subnetwork translates the destination address to the *physical address* and forwards the packet. If a router on the path has an entry in its cache, it translates the address directly to the physical address and forwards the packet, therefore avoiding triangle routing via the native subnetwork.

3.2.4 LINA

LINA (Location Independent Network Architecture) [Ishiyama et al., 2001] proposes a new addressing model, called *embedded addressing* that achieves the separation between the node identifier and the location identifier (respectively termed *node identifier* and *interface locator* by the authors) as a solution to overcome the dual use of the IP address. This is basically performed by dividing the network layer into two sub-layers, the *identification sub-layer* and the *delivery sub-layer*. The application layer specifies the destination whether by its *node identifier* or its *interface locator*. In the first case, mobility is supported, and the application communicates with a node regardless of its topological location. In the second case, mobility is not supported and the application communicates with a specific interface (i.e. at a particular point of attachment) regardless of the identity of the node. The mapping between the *node identifier* and the *interface locator* is performed by the *identification sub-layer* whereas the *delivery sub-layer* is in charge of effectively transmitting the packet to the current point of attachment.

In practice, an *ID-embedded Locator* is constructed from the *node identifier* and the current *interface locator* (*embedment operation*). The inverse operation is termed *extraction*. The binding between the *node identifier* and the current *interface locator* is maintained at a Mapping Agent. A mobile node registers the new mapping with its Designated Mapping Agent when it changes its point of attachment, and periodically.

When a correspondent node wants to communicate with a mobile node, the application layer only knows the *node identifier*. A *generalized identifier* is constructed from the *node identifier* and a well-known fixed *dedicated locator* (typically a virtual *interface locator*), by means of the *embedment operation*. This *generalized identifier* is then transmitted to the *identification sub-layer*, which performs *extraction* to determine the *node identifier*. It then calls the *Mapping Agent* for the current *interface locator* corresponding to the *node identifier*. Then. embedment is performed again, this time with the *node identifier* and the

current *interface locator* instead of the *generalized identifier*. The *ID-embedded locator* is obtained and the packet can be passed to the *delivery sub-layer*. At the MN, the *extraction* operation is performed by the *identification sub-layer*, then the *embedment* operation with the *dedicated locator*.

LIN6 is the IPv6 implementation of LINA to support node mobility. LIN6 makes use of the Aggregatable Global Unicast Address Format [Hinden et al., 1998]. The *node identifier* is 64-bits large. The *dedicated locator* is a fixed IPv6 predefined 64-bit prefix, known by all IPv6 nodes, that does not identify any specific subnetwork. This 64-bits prefix is therefore concatenated with the 64-bits *node identifier*] to form the *generalized identifier*. The use of the DNS acting as a *Mapping Agent* is discussed, and authors conclude that the *Mapping Agent* should better be a dedicated server. The DNS is however use to return the address of the Mapping Agent corresponding to a particular node. The current *interface locator* is the AGUAF IPv6 address assigned to the interface of the node. The lower 64-bits bits of this IPv6 address is used for the *node identifier*. LIN6 maintains compatibility with the existing IPv6. Nodes that do not implement the new features use the standard IPv6 address of the destination, and mobility is of course not supported. The authors also compare their solution with GSE [O'Dell, 1998] (see section 3.3.1) and advocate that is does not face the same issues outlined in [Crawford et al., 1998].

3.2.5 Columbia University

The proposal designed at Columbia University [Ioannidis et al., 1991; Ioannidis and Maguire Jr, 1993; Ioannidis, 1993] is targeted to address Local-Area Mobility in IPv4, and Wide-Area Mobility, in some extent. Mobility is managed by a set of cooperating routers within a campus, called *Mobile Support Stations (MSSs)* [Ioannidis et al., 1991] or *Mobile Support Routers (MSRs)* [Ioannidis and Maguire Jr, 1993]), each servicing a logical or geographical segment (cell) of the campus. A *MSR* must be located in each subnetwork that may be visited by MNs.

Mobility is performed by means of two new IP-layer protocols: the *Mobile Internetworking Control Protocol (MICP)* is defined to register the MN and to locate other *MSRs* in the network; the *IP-within-IP Protocol (IPIP)* is defined to tunnel packets between two *MSRs*.

The MN does not need to change its address as a result of its motion. It retains its permanent address obtained on the native subnetwork and used as a node identifier. When a MN moves to a new subnetwork, it registers with the local *MSR*, and the previous *MSR* is informed of the *MSR* currently handling the MN (forwarding address, used as the location identifier). CNs send packets directly to the permanent address of the MN. If the local *MSR* does not have an entry for the MN, it queries all the other *MSRs*. If multicast is supported, the query could be sent to the multicast group of all *MSRs*, otherwise the query would be flooded to all neighbors *MSRs*. The reply, a binding between the MN's address and the address of the *MSR* handling the MN, is cached. Subsequent packets are then encapsulated to the *MSR* handling the MN.

Wide-Area Mobility is also slightly considered in this proposal, though considered unusual. In this case, a *MSR* is needed in each potential visited subnetwork. When the MN enters in a new network, it acquires a temporary address in the visited network and registers it with the *MSR* in its home network. This *MSR* is then able to tunnel packets to the temporary address.

3.2.6 Cellular IP

The Cellular IP proposal from Columbia University (COMET) and Ericsson [Valkó, 1999; Campbell et al., 1999a,b] defines a new routing protocol to handle Local-Area Mobility (the term used in the papers is *micro-mobility*) in an IP cellular network. It relies on Mobile IPv4 to provide Wide-Area Mobility. The usual unicast routing protocols are replaced by Cellular IP. A MN entering a new domain is assigned a careof address, no change of address is required when the MN changes its point of attachment within the domain. Cellular IP supports fast handoff and paging techniques. It integrates location management and handoff support with routing. To minimize control messaging, regular

data packets transmitted by MNs are used to refresh host location information and to maintain reverse path routes from the MN to the domain border router. In order to extend battery life and to reduce traffic on the air interface, MNs do not have to update their location upon each handoff. The location of idle MNs is tracked only approximately by Cellular IP. When there is a pending packet for an idle MN, this one is paged, and the MN updates its location.

The interested reader may also read [Shelby et al., 2000] for the IPv6 adaptation, and the web page [COMET].

3.2.7 HAWAII

The HAWAII [Ramjee et al., 1999a,b] protocol from Lucent Technologies defines a routing protocol to handle Local-Area Mobility and relies on Mobile IPv4 to provide Wide-Area Mobility. A MN entering a new domain is assigned a careof address. It retains its careof address while moving within the visited domain, thus the HA does not need to be notified unless the MN moves to a new domain. Router in the domain maintain *host-specific routes* for each MN in the domain. The routing information is created, updated and modified by explicit signaling messages sent by MNs. A multicast protocol is used to page the MN when incoming data packets arrive and no recent routing information is available.

3.2.8 INRIA CBTM

The Core-Based Tree Mobility (CBTM) management scheme from INRIA [Castelluccia, 1997, 1998b; Castelluccia and Jacquemin, 1998] is targeted to IPv6 and supports both Local-Area Mobility and Wide-Area Mobility. Mobile IPv6 is advocated for supporting Local-Area Mobility while multicast is used to support Wide-Area Mobility. This proposal aims in reducing signaling in the core network, and handoff latency. The key idea is to allocate a multicast address to each MN and to route packets to the MN along a multicast tree. The author proposes to rely on a *sparse-mode multicast protocol* (namely PIM-SM), which is more appropriate to this situation. The benefit of the multicast address is to provide a location independent address which can therefore be used as a permanent node identifier. Local-Area Mobility is hidden to CNs.

The MN obtains a careof address on each visited link. Each time this occurs, the MN joins the multicast group with the new careof address. Indeed, the MN joins the PIM Rendez Point (RP) servicing the visited domain. The same RP is retained as long the MN keeps visiting the same domain, otherwise a new RP is chosen. A new RP distribution algorithm is defined to determine the current MN's RP (i.e. the current visited domain). Data packets are sent by the CNs to the multicast address of the mobile node. They are first encapsulated to the closest RP, therefrom decapsulated and re-encapsulated to the current domain's RP, therefrom decapsulated and forwarded along the multicast tree up to the MN. If the MN has registered with more than one careof address, packets are delivered simultaneously to all its careof addresses, which provides smooth handoffs. On the opposite direction, the MN must insert its multicast address in all data packets it emits in order to be identified by its CNs. This is done by means of a Multicast Source Address Option, a new IPv6 Destination Header Option.

3.2.9 MSM-IP from University of Illinois

The Timely research group at the University of Illinois [Mysore and Bharghavan, 1997] proposes the use of multicast for addressing and routing data packets to mobile nodes, in IPv4. The authors advocate that both multicasting and mobility involve similar issues of location independent addressing, address translation, packet forwarding and location management. Basically, a message sent to a multicast group does not embed the location of the destination. Hence, location independent addressing and routing mechanisms are necessary to support both multicasting and mobility support.

MNs are identified by a unique multicast address which plays the role of a location independent and invariant node identifier. Multicast routers are deployed in every subnetwork that either supports the transmission or reception of multicast packets. They form a virtual network among themselves, in order to exchange group membership information, to perform location discovery and to forward packets. The MN, as a receiver, must advertise its presence via IGMP on each new visited link. The multicast router on the subnetwork then joins the group. In addition, the MN may update a *location server* with a more recent multicast router used to join the distribution tree. CNs do not need be members of the group to transmit to the group, packets are simply sent to the multicast address. The local multicast router in the subnetwork of the CN picks up the packet and attempts to discover how and where to join the tree. The *location server* is queried for this purpose and returns the address of a multicast router. The requester then joins the multicast distribution tree with this information and packets are forwarded along the multicast tree.

3.2.10 INRIA HMIPv6

Hierarchical Mobility Management from INRIA [Castelluccia, 1998c,a, 1999, 2000] proposes another hierarchical mobility support scheme, targeted to IPv6, and that handles Local-Area Mobility and Wide-Area Mobility differently. IETF Hierarchical Mobile IPv6 is principally based on it. The author proposes to use Mobile IPv6 for supporting Wide-Area Mobility when the MN crosses domain boundaries whereas Local-Area Mobility may be supported by Mobile IPv6 or any mobility support protocol. Any number of levels could be envisioned, though in practice only one is needed.

A *mobility subnetwork* is deployed in every domain and serves as an address space reserved to assign temporary domain addresses (VCoA) to visiting MNs. A MN that enters the domain gets a new physical careof address on each visited link and negotiates a VCoA from the *mobility subnetwork*. A *Mobility Server (MS)* on this mobility subnetwork is used to register the binding between the home address of the MN and the VCoA. This VCoA is advertised to the HA and the CNs while the advertisement of the physical careof addresses are confined within the domain. CNs send packets to the MN via the VCoA using a Routing Extension Header; the packet gets routed to the MS which encapsulates it down to the MN. Local motion of the MN is therefore not perceived by CNs.

3.2.11 Hierarchical Mobility Management in CLNP

In [Carlberg, 1992], the question of mobility is targeted toward the ISO Connection-Less Network Protocol (CLNP), but with a few changes, it could easily be migrated to TCP/IP. In the OSI addressing architecture, the address has a dual semantic as in IP, thus the problem caused by mobility is very similar. The general idea is to use routing to deal with Local-Area Mobility and to use a *directory service* to store the domain where the MN resides to deal with Wide-Area Mobility. Thus, the *directory service* is used to find out the domain, whereas normal routing take over once the packet reaches the domain.

In OSI, the global network is organized into *domains*, themselves organized into *areas* comprising one or more subnetworks. The MN has a logical address (*logical-NSAP*) used as a permanent identifier to maintain transport connections alive. It is also allocated a routing address in the visited area (*area-NSAP*). Mobility within the domain is supported by the intra-domain routing protocol which is augmented with a new *hello* message. The newly obtained routing address is then flooded to some or all routers in the area and also to the border routers servicing each area in the domain. Subsequent movements within an area does not require a change of the routing address, however, routing updates are propagated in the area. Routing updates are propagated to all border routers servicing the other areas only when the MN enters a new area.

The OSI Directory Information Tree is used in conjunction with Inter Domain Routing Protocol (IDRP) to store and distribute the domain that hosts the MN. When the MN visits a new domain, the border router servicing the domain updates the directory with the identifier (prefix) of the domain that hosts the MN. Traffic from the CN is directly sent to the logical address. If the MN is not in

the same domain, the packet reaches the border router of the source domain which calls the directory for the identity of the domain hosting the MN. Packets are then encapsulated to the border router servicing the domain returned by the directory. Once decapsulated, the packet is forwarded to the MN using standard routing.

3.2.12 Hierarchical Foreign Agents

[Perkins, 1996b] proposes a hierarchy of foreign agents to handle `Mobile IPv4` registrations locally. The network is divided into regions and each region is serviced by a foreign agent. Each foreign agent knows the hierarchy of foreign agents from the top of the hierarchy down to it and advertises it, so that it could be determined by the MN. The HA is at the top of the hierarchy. Tunnels are established along the path of foreign agents that serve the MN. The MN is permanently identified by its home address and has a temporary careof address. Packets are sent to the HA which tunnels them along the chain of intermediate foreign agents.

3.2.13 Hierarchical Mobility Management by Caceres

[Caceres and Padmanabhan, 1996] proposes a hierarchical extension to `Mobile IPv4` and separate mobility in three cases (mobility between base stations on the same subnetwork, i.e. *link-local mobility*, Local-Area Mobility and Wide-Area Mobility). Local-Area Mobility is managed by a hierarchy of foreign agents. Each foreign agent maintains a routing entry for each mobile node under its coverage area. Upon motion between subnetworks in the same domain, the foreign agent on top of this hierarchy must be notified. The HA only keeps track of displacements between domains. Thus, Local-Area Mobility is transparent to both the HA and the CNs.

3.2.14 Concurrent Online Tracking of Mobile Users

The question of locating mobile users or devices and updating their location is theoretically addressed in [Awerbuch and Peleg, 1991]. The authors observe that tracking mobile users relies on two operations: *move* and *find*. *Move* relates to the necessary updates of the information which allows for localization while *find* relates to the process of accessing this information. Based on this observation, the authors describe a hierarchical regional directory which allows tracking dynamically mobile users while they are roaming. It optimizes both the communication access cost for finding a mobile user and updating its current location into the directories where it is registered.

In order to communicate with a mobile node, the correspondent node first need to locate the mobile node in the network. Finding the current location is easy if the mobile node reports all its displacements to some centralized directories in the network, but it is as well expensive to update all this information. This is particularly true if mobile users are allowed to move frequently. The mechanism has to be dynamic enough, and should be as optimal as possible (lessen communication overhead and delays, providing optimal path, ...). There is therefore a trade-off between providing a full information of the mobile node's topological location (cheap to find, expensive to update) and providing no information about the current location (cheap to update, expensive to find).

The authors propose an intermediate partial-information strategy based on a hierarchy of regional directories RD_i organized into i levels. The number of levels is $\log(weighted\ diameter\ of\ the\ network)$. The i^{th} level registers users residing within distance 2^i from it. That is, only directories in the near-by area of the mobile are updated. The lowest level is updated first; the number of levels where the tracking information must be updated is varying according to the distance of the movement. This organization of the regional directory is based on regional matchings (a graph-theoretic structure, based on the concept of sparse graph covers).

When MN moves, it does not update its address at all RD_i but only at the lowest levels. MN leaves forwarding pointers at the old location and maintains locally a list of regional addresses it crossed, and the length l_i of the migration path since it updated RD_i. It decides to update its regional address at RDi if $l_i \geq 2^{i-1} - 1$. In other words, if MN moved a distance d, $\log d$ lowest levels are updated. By updating we mean that the directory will have a binding pointing directly at the new address. The lowest the level, the more accurate is the regional address. Pointers at distant location are updated less often. This implies that the path is less optimal, but the MN can still be reached by the help of the forwarding pointers it left behind. Updates are therefore local and require low communication complexity

The CN first queries the lowest level i of its regional directory. It will get a regional address of the MN from the MN's regional directory situated a distance less than 2^i away. Otherwise, one level upper is queried and so on up to the highest level if necessary, which always succeeds. In case the returned pointer is not up-to-date, the CN tracks the MN along the path of forwarding pointers. As a result, only nearby CNs are able to locate the MN directly.

3.2.15 LAR

The Network Research Group at University Louis Pasteur (Strasbourg, France), is working on an address scheme which allows to identify nodes independently of their topological location. [Noël, 1998; Noël et al., 1998, 2001] proposes a Logical Addressing and Routing architecture (LAR) for IPv6 that considers uniformly both point-to-point communication and multicast communications for either fixed nodes or mobile nodes. This team is currently working on the adaptation of this architecture to IPv6 Cellular Networks.

In this architecture, a *LAR* sub-layer is inserted between the network and the transport layers and must be supported by both a set of *LAR*-speakers routers and end-nodes. In addition, a *Communication Manager* (*CM*) sits on top of the *LAR* layer. The *CM* is only needed by end-nodes in a communication flow. Its purpose it to create and discard *LAR* communications on one side, and to control join and leave on the other side. For fixed nodes, the *CM* may actually be co-located with that node (the *CM* is then called *Local Manager* or *LM*). For MN, the *CM* cannot be co-located. The *CM* is thus a server on the home link of the MN and is called *Delegated Manager* (*DM*). Each node supporting this architecture owns a *LAR* communication address. This address is logical, location independent and globally unique, and transparent to the application. It can either identify a group of nodes (multicast) or a single node (unicast). *LAR* addresses are taken from the un-allocated IPv6 address space.

For group communication initiated by the MN, its *DM* creates a *LAR* communication address for the group name provided by the MN and inserts this binding in the DNS, together with its name. A MN wishing to engage in an existing group communication calls the DNS which returns the *LAR* communication address and the address of the *CM* for this group. A joining message is then sent to this *CM* which determines the closest *LAR* branching router in the *LAR tree*. A branch from this branching router to the MN is created as a result of this. When the node moves, it updates the cache of its *LAR* neighbors. A mechanism is also provided to reconfigure the tree in case the MN gets closer to another *LAR* router.

A unicast communication is processed similarly by considering a group with only two members. This results in creating a tree with a single logical branch between the two peers. This logical branch remains permanent independently of the location of the MN since it is based on the *LAR* address of the MN, and not on its physical address. A CN wishing to communicate with a MN first queries its *LM* which in turn calls the DNS for the IPv6 address of the MN. Instead of returning the IPv6 address of the MN, the DNS returns the *LAR* address of the MN, the *LAR* and IPv6 addresses of the *DM* serving the MN. A communication request is then sent by CN's *CM* to the *DM* and relayed up to the MN. The MN then creates a logical branch between it and the CN using its *LAR* address.

When the MN moves frequently, it determines a *LAR* branching router between its previous visited link and its current visited link. This is performed by sending a request toward its CN, which is intercepted by

the *LAR* branching router on the path. This triggers the update of the cache at the CN which now sees the branching router as its neighbor. The displacements of the MN within the vicinity of this branching router is therefore transparent to the CN. If the MN moves far away from this branching router, the tree must be reconfigured, and the cache at the CN must be updated. This mechanism is also useful when the MN frequently moves between two close ARs, therefore avoiding to register and de-register continuously.

LAR is implemented as IPv6 Extension Headers. The *LAR Destination Option Extension Header* is used for the purpose of carrying the *LAR* address of the source node and the *LAR* communication address. It is used for identifying both ends of the communication flow. It must be inserted in all packets. The LAR Hop-by-hop Option Extension Header is used to search for branching routers. It is sent by the *CM* to the joining node. If a *LAR* node on the path is already on the tree, it grafts the new member in the tree.

3.2.16 Deadalus

Deadalus [Seshan and Balakrishnan, 1995] is very similar to Mobile IPv4. Packets destined to a MN are delivered to the MN's HA and then multicast to the base stations in the neighborhood of the MN as a means to reduce handoff latency and packet loss. A pre-arranged multicast address is used for this purpose. Using beacons and signal strength measurements, the MN determines which access routers should join the group, and which access routers it is likely to move in the near future. This proposal suffers from triangle routing via the HA.

3.2.17 DCM

[Blazevic and Le Boudec, 1999, 2000] defines a new multicast routing protocol called *Distributed Core Multicast (DCM)*. This protocols is designed to scale to a large number of groups with few receivers. Once applied to cellular telephony, it is used to route packets down to MNs. Each MN is assigned a multicast address in every visited domain. The current AR of the MN, and neighboring ARs that anticipate its arrival, joins the multicast tree on behalf of it. The benefit of this approach is to reduce latency and packet loss during handoff.

3.2.18 Helmy

[Helmy, 2000] is yet another scheme that proposes to multicast data packets from correspondent nodes to mobile nodes in IPv6. The objective is to reduce latency and packet loss during handoffs in order to meet the requirements for audio applications. The MN is identified by a multicast group and joins the group from the visited subnetworks. CNs send data packets to this multicast group. The use of multicast is advocated because it is perceived that the movement of the MN is in a geographical vicinity, thus limiting the number of hops necessary to reach the multicast distribution tree.

3.2.19 Mobile Next Generation Internet

[Ernst, 1998] is a prospective geographic addressing and routing scheme for mobile nodes in IPv6, based on Dynamic DNS Updates [Vixie et al., 1998], the GPS[4], and a constellation of *Low Earth Orbit* and

[4] GPS stands for *Global Positioning System*, a constellation of satellites used to determine geographical location. GPS is based on a system of coordinates called the Worldwide Geodetic System 1984 (WGS-84), which is similar to lines of latitude and longitude. The location of a GPS user is given by computing the distance from at least three GPS satellites. The accuracy of the location is on the order of 300 feet for civilians, but an accuracy of 3 feet can be obtained by the army. This accuracy decreases if the device taking measurements is moving.

geostationary satellites for managing mobility. This addressing scheme achieves the separation between the location identifier and the node identifier but fails to consider the whole picture and raises a number of issues like security, ease of deployment, impact on upper layers, etc. Basically, the IPv6 address is divided into the Geographical Location Part (GPA), solely used as the location identifier, and the Global identification Part (IPA) solely used as a node identifier. The GPA is a geographical coordinate provided by the GPS. The MN registers in the DNS information about the current domain where it is located. Once the DNS is queried by CNs, the reply contains the IPA of the MN and the information about the current domain of the MN. The CN then fills the destination address with the IPA of the MN and the GPA that identifies the expected geographical location of the MN.

3.2.20 DNS Updates

[Snoeren and Balakrishnan, 2000] proposes an end-to-end architecture based on dynamic DNS updates. This proposal is targeted to TCP-based applications. The MN obtains a new address on each visited link and updates the DNS mappings for its domain name. A migration process is required to maintain the connection. The transport protocol is aware of the mobility mode during the migration process. This proposal avoids triangle routing but incurs handoff delays due to DNS update and migration delays. A similar approach has also been investigated earlier by [Padmanabhan N. and Katz, 1998].

3.2.21 MosquitoNet

MosquitoNet [Cheshire and Baker, 1996] is very similar to Mobile IPv4. The role of the foreign agent is plaid by the MN. Decapsulation is therefore processed by the MN itself.

3.3 Addressing Schemes

The following addressing architectures are not targeted to handle mobility and don't propose any mobility management scheme as LINA (section 3.2.4) and LAR (section 3.2.15). However, that kind of addressing schemes are potentially useful for locating mobile nodes and worth to be studied.

3.3.1 GSE - Global, Site and End-System Designator

GSE (Global Site End-system designator) [O'Dell, 1998] is an alternate addressing architecture which was proposed for IPv6 in the IETF IPNG working group.

GSE was designed to address the question of re-numbering and multi-homing. The author advocated that the impact of renumbering a site should be small (it should indeed be transparent to end-nodes), and cheap. GSE therefore makes a clear distinction between the identity of the interface, which is a globally unique number in the Internet and its point of attachment to the network. There is therefore a strict boundary in addresses between fields identifying hosts and routing to hosts. This was intended to ease re-numbering and multi-homing of networks and interfaces. The 128-bits IPv6 address is separated into three parts of fixed length. The first two pieces (64 high-order bits) are used for routing and therefore determine the point of attachment of an interface to the network (*routing stuff*). The *Routing Stuff* is split into two fields to distinguish between the global topology where data from everyone is allowed to transit and the local topology where no public data transit since it is a private partition of the Internet. The point of attachment of a private site to the global Internet is represented by the *Routing Goop* (*RG*) field (6 high-order bytes). The *Site Topology Partition* (*STP*) field (2 next bytes) identifies which link within the site

an interface belonged to. The low-order 64 bits form the *End System Designator* (*ESD*) which is used to identify the interface. This field is used alone by the upper layers.

The *RG* is hidden from nodes. As such, the *RG* field in the source address is filled with a number meaning "within this site". Routers use the *routing stuff* portion of the destination address field to route packets toward their destination. Packets destined to an interface outside the site reach site border routers where the *RG* is rewritten. The destination does therefore know in which site the source is located. Once packets enter the site through a site border router, the *RG* field of the destination address is rewritten again with a number meaning "withing this site". Moreover, the DNS the full address of interfaces. The *RG* part of the address is handled with indirection: it is referenced as its name, not its number, and is itself resolvable by DNS.

When a site has to renumber, the *RG* is updated. This doesn't have any repercussions on the address of interfaces within the site since nodes within a site don't know what their *RG* is. Renumbering is therefore done transparently. As for other perceived advantages, it facilitates load-balancing between *RGs* when sites are multi-homed, it facilitates routing aggregation. And, although it wasn't designed for it, it could facilitate mobility. Although separating the location part from the identifier part is attractive at first sight, this poses some security problems, as pointed out in [Crawford et al., 1998]. It indeed facilitates an intruder to forge addresses. To facilitate security (authentication),the *RG* should indeed not be hidden from nodes. Finding the location from an identifier is another point of concern.

GSE was therefore rejected but it demonstrated valuable ideas and contributed to major progress in the definition of the current address format (Aggregatable Global Unicast address format [Hinden et al., 1998]). The former provided-based addressing had fluid boundaries. The IETF IPNG working group agreed in creating fixed boundaries in IPv6 addresses to distinguish between portions used for routing within the public topology, routing within a site (private topology) and identifying the interface. There is therefore a clear distinction in the address between localizing an interface and identifying it. The unicast address format is therefore the following: Top-Level Aggregation Identifier (TLA) (sort of Large Structure), Next-Level Aggregation (NLA) identifier (roughly the Routing Goop), the Site-Level Aggregation (SLA) identifier (about STP), and the interface identifier. GSE has also highlighted the need to ease renumbering a site within DNS [Huitema, 1998].

3.3.2 Geographic Addressing and Routing

[Navas and Imielinski, 1997] proposes a routing and addressing method to integrate geographic coordinates (longitude, latitude) into IP addresses. Geographical addressing relies on the Global Positioning System (GPS). A GPS card is embedded into devices to determine the position on the surface of the Earth from a number of GPS satellites. This architecture aims at offering geographic services such as messaging or advertising to all clients within a specified area of coverage whether they have geographical addresses or not. This is in a way very similar to multicasting. The architecture is composed of geographic routers organized hierarchically, geographic nodes, and geographic hosts. Usual IP routers are not aware of the geographic destination which is specified in a geographic message header. In fact, geographic routing is implemented in the application layer, as a virtual network overlaid onto the current IP internetwork. *GeoRouters* are the only ones capable of understanding the geographic message header.

The destination specified in the packet is the actual coordinate of a polygon, or a circle that identifies an area of coverage. Geographic messages are sent to the closest *GeoRouter*, which checks the geographic header and determines if it serves the geographic destination of the message. Its service area spans the service area of all its child routers. For doing so, it computes the intersection between its own service area of coverage and the polygon represented by the geographic destination, i.e. it performs a polygon intersection. If the destination polygon doesn't fully intersect its polygon, it forwards the packet to its parent *GeoRouter*, and so on. If the *GeoRouter* does services part or the full destination polygon, then it checks which child routers it should forward the packet to. *GeoRouters* cache data that allow to forward following packets destined to the same geographic area. This saves processing time.

In their paper, the authors report the performance of the approach in term of routing time. The evaluation shows that the routing time increases while the destination polygon shape becomes more complex. Even with a destination area limited to a simple point, the performance of this scheme remains about 100 times slower than traditional IP routing.

3.4 Conclusion

The description of all these mobility support schemes teach us a number of things. First, we see that there are several approaches to tackle with the addressing problem depicted in section 2.4.1. Second, a few distinct architectures seem to emerge to implement the mobility support services depicted in section 2.4.2.

3.4.1 Mobility Support Approaches

The literature usually discusses two distinct ways to tackle the question of mobility support in IPv4. This discussion is equally applicable to IPv6. The first one is to redesign the TCP/IP addressing scheme, and the second one is to adapt to the existing protocols while providing additional services that preserve backward compatibility. With the advent of IPv6, we advocate a third one: embedding mobility support directly in the network layer.

3.4.1.1 Redesign of the TCP/IP Addressing Scheme

The first approach advocates a redesign of the TCP/IP addressing scheme in order to offer a clear separation between the layers. The dual semantic of the IP address and layer concept violation wouldn't exist anymore if the node identifier and location identifier could be split distinctively from the IP address and the sole location identifier used at layers above IP. This approach is quite orthogonal to the existing TCP/IP addressing scheme and would impact most protocols that compose the TCP/IP reference model, not only network-layer protocols. It is therefore not realistic for IPv4 since changes are not allowed on the already deployed infrastructure. One would think that the advent of IPv6 is an opportunity to redesign addressing and would make more sense than for IPv4. However, IPv6 cannot be diametrically different from IPv4 since one of the basic requirement of IPv6 is to allow smooth upgrading from IPv4 to IPv6 while providing backward compatibility.

We are therefore imprisoned in the current TCP/IP addressing scheme. It is only possible to make some modifications as long as it does not impact upper layers and the already standardized IPv6-related protocols. Such enhancements to the addressing scheme have been proposed recently, for instance GSE (Global, Site and End-System Designator) [O'Dell, 1998; Crawford et al., 1998] or LINA (see section 3.2.4) [Ishiyama et al., 2001] (see section 3.2.4). NSRG (NameSpace Research Group) at the IRTF (Internet Research Task Force) is also working on the separation between the node identifier and the location identifier, but not for reasons only pertaining to mobility.

3.4.1.2 Sub-Layer between Network and Transport Layers

The second approach is to bring the required mobility support services in a sub-layer between the network layer and the transport layer. The purpose of the sub-layer is to make use of the existing protocol suite at both layers without any changes and to render mobility transparent to both layers. No changes are required in the already deployed infrastructure, apart at the MN itself and at a few other mobility-aware nodes. The drawback is that the sub-layer must adapt to existing functions proposed by the network and transport layers. In this situation, mobility support may not be very efficient.

3.4.1.3 Integrated Support of Mobility in IP

Interestingly, IPv6 provides new features and provisions for future improvements. The advent of IPv6 is therefore the opportunity to embed mobility support directly within the network layer and transparently to upper layers. We are therefore not left with the second solution which more or less proposes to fill the gaps with "crummy goods", as this is done in IPv4. Basically, IPv6 is meant to support mobility well. It effectively proposes new functions that allow a better embodiment of mobility support (such as Mobile IPv6 compared to Mobile IPv4, see section 3.1.1). Issues not already covered by the current standards could effectively be addressed in order to provide long term and efficient mobility support.

3.4.2 Mobility Support Architectures

This study first shows that the current IETF standards are somewhat based on an initial proposal defined as early as in the eighties (section 3.2.1). The effort conducted in the beginning of the nineties at the IETF resulted in a number of proposals that finally served as the foundation for the existing Mobile IP standards (sections 3.2.2, 3.2.3, 3.2.5). Then, later proposals are more or less extensions or adaptation of Mobile IP to meet further requirements like reducing signaling overload, handoff delays, and packet loss during handoffs. A number of other proposals provide valuable ideas but are inadequate for IPv6, mainly due to security concerns and implementation concerns which limit the deployment of a potentially good mechanism, or diminish the optimality of the solution.

Recent work in IPv6 shows that Mobile IPv6 is better perceived as a protocol to solve Wide-Area Mobility rather than Local-Area Mobility. Since the home agent and the CNs must be notified upon every displacement of the MN, Mobile IPv6 is clearly inefficient in terms of signaling overhead for MNs with a high movement frequency between topologically adjacent subnetworks (e.g. while walking in the street or driving a car). Even if displacements are confined in a limited part of the topology, control traffic is propagated over the entire network. In addition, Mobile IPv6 does not provide means to solve open issues when mobility occurs between adjacent subnetworks: *smooth handoff*, *fast handover*, *packet loss*, *handoff delay*, *context transfer*.

Despite its critics, Mobile IPv6 is the most advanced solution. Security aspects are well addressed in the specification, though there are still security holes, as currently debated at the IETF. Thus, extensions to provide for effective performance transparency are being designed, principally in the Mobile IP and the Seamoby (Context Transfer) working groups. Simultaneously, the current work on routing protocols (Cellular IP, HAWAII), which also addresses the above issues, was judged too immature and consequently moved to the IRTF.

To conclude with this section. three main groups of proposals emerge clearly from this study: *hierarchical-based* proposals which led to Hierarchical Mobile IPv6, currently being standardized at the IETF, as a solution for Wide-Area Mobility to reduce signaling load in the core network, *micro-mobility* proposals (Cellular IP, HAWAII, ...) as an orthogonal solution for Local-Area Mobility management, and *multicast-based* proposals which exploit the common points between mobility management and multicast group management to provide a location independent and invariant node identifier.

Chapter 4

Taxonomy

This chapter presents a comparative analysis and a taxonomy of the *mobility support* schemes detailed in the previous chapter. In order to evaluate and compare them and to design potential new ones, we need to define an abstraction model that captures their design choices and their architecture. With this abstraction model, we are able to classify mobility proposals according to the location of the architecture components and the function they perform. Similar proposals can then be gathered into a common framework. We have identified a few distinct frameworks in which all proposals could basically be ranged, though there may exists some similarities between any two frameworks. This chapter uses the abstract terms node identifier and location identifier we defined in section 2.4.2.

4.1 Abstraction Model

In [Bhagwat et al., 1996], the authors define an abstraction model which fits a number of mobility support proposals. It is presented in the first section. We advocate that Bhagwat's abstraction model is too restrictive to capture the granularity of all possible mobility frameworks because it was defined to compare an initial set of host mobility proposals which fall in the same class[1]. Particularly, it doesn't help us to identify where and how are performed the *mobility support services* we defined in section 2.4.2 . We want to outline what are the components of the architecture and their function, and where in the topology they should be implemented. We have therefore defined our own abstraction model. It is described in the second section.

4.1.1 Bhagwat's Abstraction Model

In their paper [Bhagwat et al., 1996], the authors define two *functions*, and four *architecture components*, as outlined in fig. 4.1.

4.1.1.1 Functions

- function f:$f(MN\ permanent\_address) \mapsto MN\ temporary\_address$

  This function replaces the permanent address contained into the destination address field of the `IP` packet with the current temporary address of the mobile node. With respect to our mobility services and terminology, this basically corresponds to Location Lookup plus Routing. This function actually maps a node identifier to a location identifier.

[1] This class of proposals is termed *Two-Tier Addressing* and will be described later in section 4.2.2.

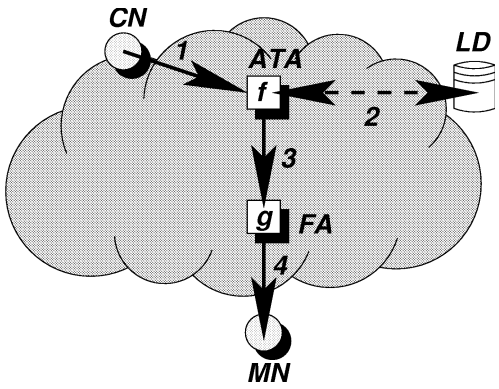

Figure 4.1: Packet Forwarding in Bhagwat's model

- function g: $g(MN\ temporary\_address) \mapsto MN\ permanent\_address$

  This function replaces the temporary address contained into the destination address field of the `IP` packet with the permanent address of the mobile node. With respect to our mobility services, this corresponds to an inverse Location Lookup plus Routing. This function actually maps a location identifier to a node identifier.

Both functions are applied on each packet. gof is an *identity mapping*. As a result, the operation is transparent to both ends.

4.1.1.2 Architecture Components

- Location Directory (LD)

  This component is a database which records the mapping between the permanent address and the temporary address.

- Address Translation Agent (ATA)

  This component performs the f function. It queries the `LD` and may cache the answer locally in order to improve processing delays.

- Forwarding Agent (FA)

  This component performs the g function.

- Location Update Protocol (LUP)

  The `LUP` is a reliable mechanism that keeps the `LD` and its cache consistent.

4.1.2 A more Detailed Abstraction Model

Our model is based on Bhagwat's model and refines it.

We define four *functions* and six *architecture components*.

4.1.2.1 Functions

The following mobility management functions may be supported:

- $update(database, node\_id, location\_id)$

  This function triggers the insertion or update in a database of a binding between a node identifier and a location identifier. This function may be performed as often as the mobile node enters a new subnetwork. The trade-off is keeping an up-to-date location identifier which optimizes routing versus minimizing signaling overhead.

- $lookup(database, node\_id) \mapsto location\_id$

  This function queries a database for a location identifier corresponding to the node identifier, provided as an input. In some cases, a location identifier may be used in place of a node identifier as an input (for instance if implemented as a chain of forwarding addresses, as justified by hierarchical schemes). This function may be performed from any place in the network between the sender and the recipient, and any number of times. It is best performed with minimum delay.

- $redirect(packet, original\_destination, new\_destination)$

  This function redirects a packet to a new destination. As a result of this function, the packet destination of the packet is modified. The packet takes a different path than the one it was originally taking (the packet is re-routed). This means that some additional operations not considered as part of the usual routing functions (like routing table lookup and output interface selection) are performed on the incoming packet. This function could be performed by means of *Encapsulation*, but is not limited to this.

- $forward(packet)$

  This function is similar to redirect but leaves the destination of the packet unchanged (the packet is not re-routed). This means that some additional operations not considered as part of the usual routing functions (like routing table lookup and output interface selection) are performed on the incoming packet. This function could be performed by means of any of the following existing mechanisms: *Decapsulation*, *Routing Extension Header* (source routing), etc, but is not limited to these.

4.1.2.2 Architecture Components

The Location Directory The Location Directory is a repository that records binding between a node identifier and its corresponding location identifier. It may be centralized, distributed or hierarchical and be subdivided into:

- Primary Location Directory (PLD): the database where is recorded the most up-to-date copy of a particular binding.
- Secondary Location Directory (SLD): a database where is recorded a less up-to-date copy of a particular binding , like a cache (i.e. a copy of the binding registered in the PLD that may not be maintained up-to-date).

Mobility Agents (MA) A Mobility Agent is an entity that performs *one* or *several* of the mobility management functions outlined in the previous section. A Mobility Agent could be any of the following ones:

- Updating Agent: a Mobility Agent that maintains a binding in the Location Directory by means of the $update$ function.
- Locating Agent: a Mobility Agent that queries the Location Directory by means of the $lookup$ function.
- Redirecting Agent: a Mobility Agent that receives a packet not intended to itself, that performs some mobility management processing on the packet ($lookup$ function), and that redirects the packet to a new destination ($redirect$ function). As a result from this, the header of the packet is modified.

- Forwarding Agent: a Mobility Agent that receives a packet not intended to itself, that mat perform some mobility management processing on the packet (*update* function), and that forwards it toward its original destination ($forward$ function). As a result from this, the header of the packet is *not* modified.

Location Update Protocol (LUP) The `LUP` is a reliable mechanism that keeps the `LD` and its cache consistent. It performs Location Update (management of the Location Directory) and Location Lookup (query of the Location Directory).

With respect to the three mobility services as outlined in section 2.4.2 and to our abstraction model, Location Update means *how to update the Location Directory* whereas Location Lookup means *how to query the Location Directory*. Location Update and Location Lookup are performed respectively by the *update* and *lookup* functions. These functions account for the signaling between the Mobility Agents and the Location Directory. In addition, Routing means *how to deliver datagrams* to the specified destination given its location identifier and is performed by means of two new functions, $redirect$ and $forward$, in addition to the usual routing functions at the routers. With this model, we are able to determine where in the network the components are located.

4.2 Mobility Support Frameworks

All studied mobility proposals make use of all or some of the components of our abstraction model, but in a different way. Our study shows that the main difference between the different proposals is the architecture of the Location Directory, its location in the topology, and the location and number of Mobility Agents. In practice, mobility components may be located anywhere in the network and may be distributed (duplicated or hierarchical) or centralized. In addition, they make use of distinct node identifier and location identifier. This allows us to identify two distinct categories and a few basic frameworks in which all schemes could be ranged. Typically, most of the studied mobility schemes could be ranged in more than one framework at the same time.

4.2.1 Network-based Category

This category includes all frameworks that rely on network mechanisms (unicast routing, broadcasting, etc) to support mobility. The mobile node retains its address and routing protocols are used to propagate its topological location in the entire network. Specific mobility management services are not necessarily required in the network. According to our abstraction model, the role of the Location Directory is in a sense plaid by the forwarding table, routers are somehow MAs, and the Location Update Protocol is performed by routing protocols.

4.2.1.1 Routing-based Framework

In the *routing-based framework*, the three mobility services depicted in section 2.4.2 (Location Update, Location Lookup and Routing) are performed by enhanced routing protocols. No specific Mobility Agents, nor Location Update Protocol are really needed. The MN retains its address which is used both as the node identifier and location identifier. As a result of the displacement of the MN, *host-specific routes* are propagated by the routing protocol and new routes are computed. The MN doesn't participate actively in this process. CNs do not need to know the topological location of the MN. Packet forwarding from a CN to the current topological location of the MN solely relies on the routing protocol and the location information recorded in the forwarding table.

This framework adapts to the mobility addressing problem (see section 2.4.1) by avoiding the address change. However, it requires the ability for the routing protocol to react quickly to topology changes and routers to keep *host-specific* entries in their forwarding table. This contrasts with the route aggregation effort of conventional routing protocols. Since the lack of routing aggregation does not scale to a large number of nodes, a solution based on this framework is clearly inadequate for Wide-Area Mobility. On the other hand, it may be adequate for Local-Area Mobility as demonstrated by `Cellular IP` [Valkó, 1999] and `HAWAII` [Ramjee et al., 1999a,b]. Both solutions define a specific routing protocol to handle Local-Area Mobility and make use of `Mobile IPv6` to handle Wide-Area Mobility. In practice, new mechanisms like *paging* are introduced as a means to keep state in routers only for active MNs. Paging offers better scalability, while it reduces battery consumption and signaling. Authors usually claim faster handoffs, and the ability to react quicker to failed links. [Carlberg, 1992] also follows this framework for Local-Area Mobility in `CNLP`. More active work in this area is effectively being proceeded at the `Internet Research Task Force` (`IRTF`) in the `Routing Research Group`. [Roberts and Loughney, 2001] investigates the limits and the issues of a routing-based solution and worth reading. A comparison between concurrent Local-Area Mobility proposals can be found in [Campbell and Gomez, 2001; Campbell et al., 2002].

4.2.1.2 Broadcast-based Framework

In this framework, a node that has a packet to send to a MN does not care about its current location and simply broadcasts the packet directly in the entire network. The MN retains its address which is used both as the node identifier and location identifier. CNs are able to communicate directly with the MN without cumbersome location maintenance and query.

No particular proposal ranges into this framework, apart *link-local mobility* support, which is out of scope of the present study. However, this framework is somewhat followed by schemes that propose *paging* as a means to route packets to *idle* MNs (`Cellular IP`). Idle MNs then switch to an active mode and another technique is used to determine the location of the MN and route packets to it.

This framework is very simple and straight forward to implement and is in a way similar to the *routing-based framework*, but does not require to update the routing tables as a result of mobility. However, flooding the network with data packets is a very important drawback. Thus, this framework could only apply to Local-Area Mobility when a small number of small packets are transmitted to a MN.

4.2.2 Two-Tier Addressing Category

Two-Tier Addressing, as defined in [Bhagwat et al., 1996] (see section 3.1.1), adapts to the addressing problem posed by mobility quite well by associating a mobile node with two addresses: a permanent address, used as the node identifier, and a temporary address, used as a routing directive (location identifier). Thanks to *Two-Tier Addressing*, the mobility management is transparent to the already deployed network, which is probably one of its main advantages. No changes are required at upper-layers either. The first issue is how to distribute the routing directive to a number of nodes or servers in the network. The second issue is how to route the packet to the current topological location of the MN. *Encapsulation* and *source routing* are the main mechanisms used to redirect packets to a new address. They do this without actually rewriting the destination address of the packets.

In practice, the node identifier could alternatively be a virtual address, a forwarding address, or a multicast address, while the location identifier could alternatively be the address of a *forwarding address*) (e.g. the address of a RA), or a chain of forwarding addresses, or a physical address (a topologically correct address owned by the mobile node on its current point of attachment). Any number of *lookup* functions may be performed along the path between the source and the destination of a packet. In this case, subsequent calls to the Location Directory return a chain of *forwarding addresses*. The MN is usually responsible to maintain an up-to-date binding between the two addresses in the Location Directory (Location Update service).

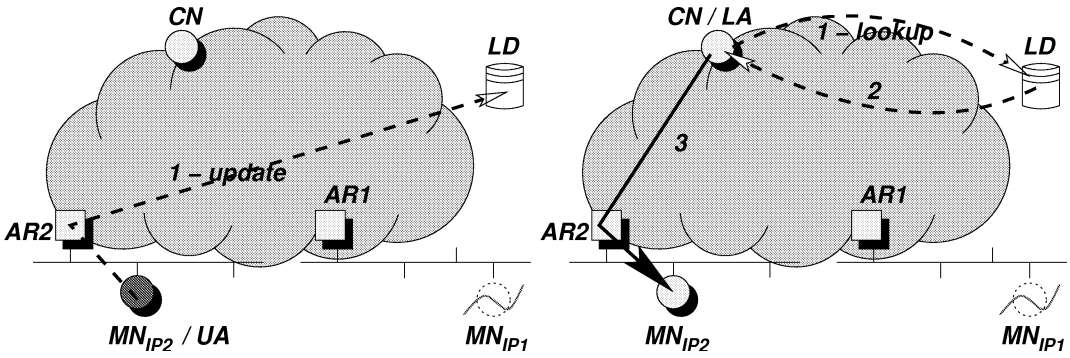

(a) Location Update

(b) Location Lookup and Routing

Figure 4.2: Location Directory Framework

With respect to our abstraction model, the Updating Agent is co-located with the MN while the other components (Locating Agent, Forwarding Agent and Redirecting Agent) are distributed differently.

4.2.2.1 Location Directory Framework (Proactive Framework)

This framework is based on the *Two-Tier Addressing Framework* and is illustrated on fig.4.2. A remote Location Directory holds bindings between the node identifier and the location identifier. The binding in the Location Directory is maintained by the MN, itself acting as a Updating Agent. The Location Directory is queried by the CN, acting as a Locating Agent, for the current location identifier of the MN before packets could actually be sent to it. No other MA is needed.

A central and real-time realization of the Location Directory is infeasible, as demonstrated in [Awerbuch and Peleg, 1991; Bhagwat et al., 1996], thus the Location Directory must be distributed. In this case, the location identifier returned by the Location Directory may be cached in a Secondary Location Directory directly at the Locating Agent. There is of course a tradeoff between querying the Location Directory before sending each packet, which leads to longer delays but provides the most up-to-date binding for the MN; and querying it from time to time which results in less processing delays, but may cause packets to be routed to a non-accurate location. Thus, a Location Update Protocol must keep bindings up-to-date at both the Primary Location Directory and Secondary Location Directory. In any case, querying the Location Directory, and maintaining consistency between the Primary Location Directory and the Secondary Location Directory incurs a considerable amount of traffic.

One of the first mobility schemes, designed as early as in the eighties by Sunshine and Postel [Sunshine and Postel, 1980] could be ranged in this framework [2]. [Snoeren and Balakrishnan, 2000] also ranges here. A similar approach based on the DNS was also considered by [Padmanabhan N. and Katz, 1998].

4.2.2.2 Third Party Framework (Reactive Framework)

This framework, as illustrated on fig 4.3, is based on the *Two-Tier Addressing Framework*. The CN does not need to query the Location Directory for the temporary address of the MN and does not care about its topological location. Thus, the Location Update Protocol is minimized. Packets are directly sent to the MN's permanent address but get intercepted by a MA (the "third party") that implements both a Redirecting Agent, and a Locating Agent. When it receives a packet intended to a MN, this MA acts as a Locating

[2] Although the term *virtual network* is used, this scheme doesn't fall into the *virtual network framework* as defined in section 4.2.2.5

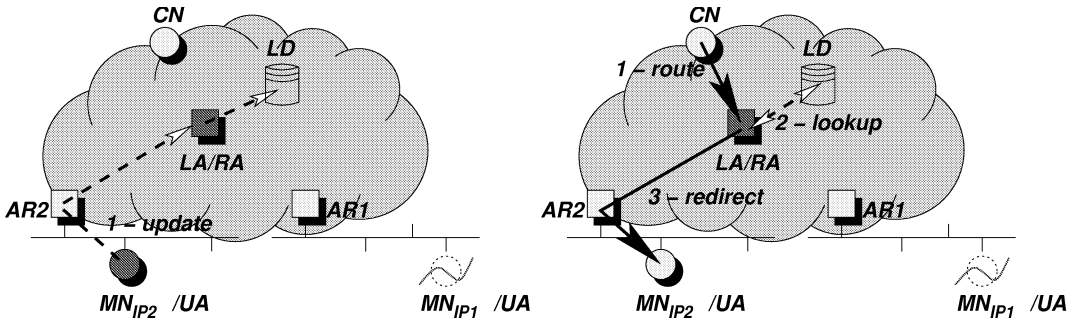

(a) Location Update (b) Routing, Location Lookup, and re-Routing

Figure 4.3: Third-Party Framework

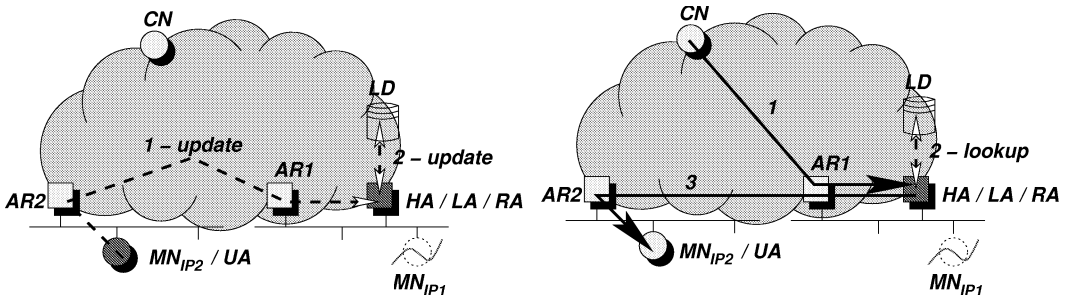

(a) Location Update (b) Routing, Location Lookup, and re-Routing

Figure 4.4: Home Agent Framework

Agent and queries the Location Directory. Then, as a Redirecting Agent, it redirects the packet towards the actual location of the MN.

This framework may be further subdivided into the *Home Agent Framework*, the *Hierarchical Framework*, the *Virtual Network Framework* and the *Multicast Framework*.

4.2.2.3 Home Agent Framework

This framework, as illustrated on fig 4.4, is a sub-case of the *Third-Party Framework* and obeys to the *owner maintains rule*. The MN has two addresses: the permanent address is the physical home address obtained on the native subnetwork (*home link*) and used as the node identifier, and the temporary address is a physical address obtained on each visited link, used as the location identifier. A dedicated router on the home link, usually termed the home agent (HA), plays the role of the *Third-Party*. The principle of this framework is to allow a MN to be always reachable at its home address. The HA acts as the Locating Agent and the Redirecting Agent and implements a local Location Directory. The MN acts as a Updating Agent and keeps the binding up-to-date at the HA. Since an individual HA only cares about MNs that have obtained their home address on its link, a HA must be deployed in every subnetwork. The Location Directory is therefore distributed amongst all the HAs in the Internet.

This framework has a number of drawbacks. First, the HA is a single point of failure and must be notified by the MN upon every displacement in the topology. Second, packets do not follow the most optimal path (triangle routing). Longer delays and data losses may result during handoffs particularly when the

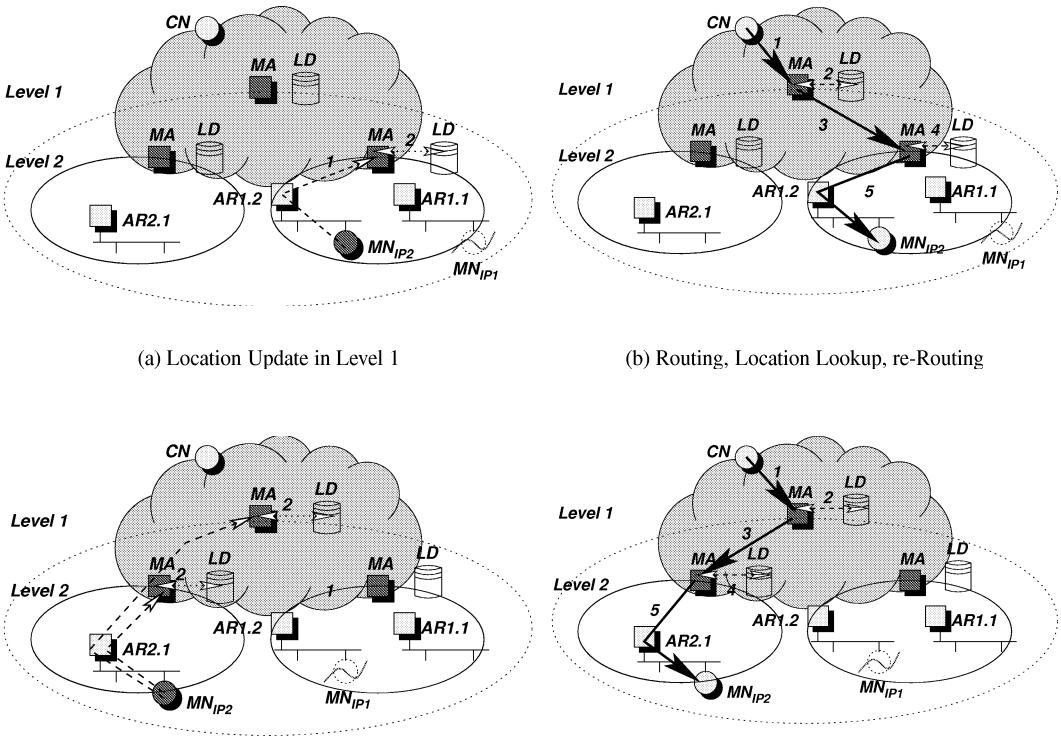

(a) Location Update in Level 1

(b) Routing, Location Lookup, re-Routing

(c) Location Update in Level 2

(d) Routing, Location Update, re-Routing

Figure 4.5: Hierarchical Framework

MN performs Wide-Area Mobility. Basically, this framework is more appropriate for mobile nodes that usually reside on their home link and occasionaly move away from it.

`Mobile IPv4` [Perkins, 1996a] is the most popular implementation of this framework. The performance of the *Home Agent Framework* could largely be enhanced with a Secondary Location Directory duplicated in the corresponding networks, usually at the CN itself. The CN would then also act as a Locating Agent. This avoids triangle routing via the HA but puts an additional burden in the network in terms of signaling load to maintain up-to-date bindings in the Secondary Location Directory. This maintenance is both the MN's responsibility and CN's responsibility. Proposals with this enhancement fall both in the *Location Directory Framework* and the *Third-Party Framework*. This is the case for `Mobile IPv4 with Routing Optimization` [Perkins and Johnson, 2000], `Mobile IPv6` [Johnson and Perkins, 2000], and other proposals ([Johnson, 1993], ...).

4.2.2.4 The Hierarchical Framework

This framework is based on the *Third-Party Framework*. As illustrated on fig. 4.5, the Internet is divided into a hierarchy of levels. A hierarchical Location Directory is distributed between each level and records a chain of forwarding addresses. The Location Directory at level m records a binding between a forwarding address used as a routing directive to route the packet up to level m and a forwarding address to route the packet up to level $m-1$. A third-party MA is co-located with the Location Directory at each level and serves all the MNs in the lower levels. The MA plays both the role of a Locating Agent, and a Forwarding Agent or a Redirecting Agent. The MN obtains a forwarding address in each level in the hierarchy of MAs. CNs are not aware of the current location of the MN, they only know the node identifier, which is the forwarding address that indicates the top-level MA. The MN registers itself with the MA in the lowest

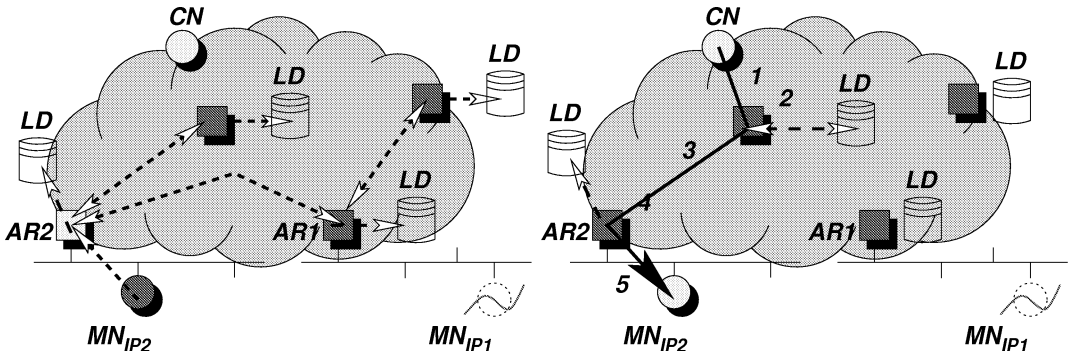

(a) Location Update

(b) Routing, Location Lookup and re-Routing

Figure 4.6: Virtual Network Framework

level. It only registers with the MA at the level above when it crosses level boundaries.

A large number of proposals could be ranged into this framework, but usually not exclusively ([Perkins, 1996b], `Hierarchical Mobile IPv6` [Soliman et al., 2001; Castelluccia, 2000], `Caceres` [Caceres and Padmanabhan, 1996],`CLNP` [Carlberg, 1992]). All proposals that distinguish Local-Area Mobility from Wide-Area Mobility could indeed range into this framework. In this case, a framework is used to manage Wide-Area Mobility, while another one is usually used to manage Local-Area Mobility.

This framework has a number of advantages when used to manage Wide-Area Mobility:

- A MN appears stationary from the point of view of the upper level. Local motion of the MN is therefore transparent to the CN. There is no impact on upper layer protocols, and this provides location privacy to the MN.
- Signaling load resulting from the local motion of the mobile is confined locally. Since most signaling is not exposed to the core network, this framework diminishes the signaling load burden and scales to a larger number of MNs.
- Compared to the *Home Agent Framework*, it permits faster handoffs and thus results in less handoffs latency and losses during the transition phase which is only performed with the closest MA.
- It could ease deployment of distinct Local-Area Mobility protocols in each administrative domain. Wide-Area Mobility between domains that run a distinct Local-Area Mobility could be achieved by means of an inter-operability protocol as the one presented in [Castelluccia and Bellier, 1999].

4.2.2.5 Virtual Network Framework

This framework is based on the *Third Party Framework*. A set of collaborative MAs form a *virtual network* which is an abstraction of the actual physical underlying internetwork. The goal of the virtual network is to hide the physical topology of the underlying network. This notion is very useful to support mobility as it hides the migration of the MN to the CNs. It indeed is similar to the concept of virtual memory where a virtual memory space hides the physical one.

The MN has two addresses: the node identifier is a permanent address taken from a virtual address space. The location identifier is a temporarily acquired physical address obtained on the visited link and is not known to the CNs. From the CN's point of view, a MN therefore appears stationary into the virtual network although it migrates within the physical network. MAs are both Redirecting Agents, Forwarding Agents,

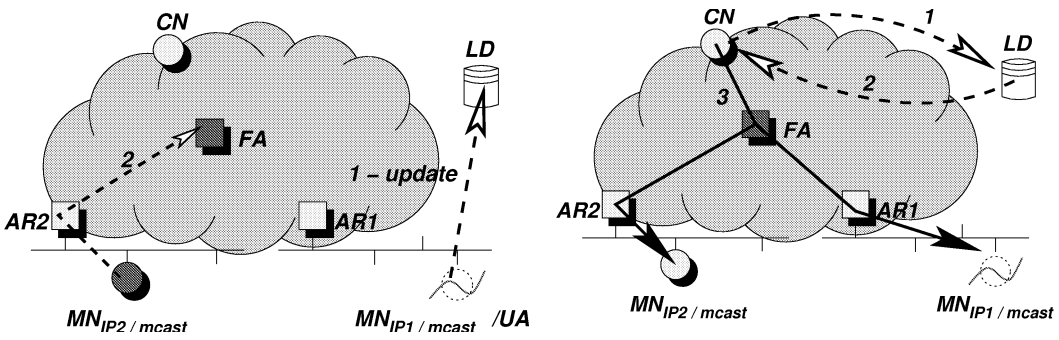

(a) Location Update

(b) Location Lookup and Routing

Figure 4.7: Multicast Framework

Locating Agents, and Updating Agents. They are also co-located with a distributed Location Directory. The MN, as the initial Updating Agent, updates its physical address with the closest MA which in turn updates this address to all or some of the other MAs.

The IPv4 proposal designed by Ioannidis et al from Columbia University [Ioannidis et al., 1991; Ioannidis and Maguire Jr, 1993] fall into this framework. Although the term *virtual network* was first introduced by Sony [Teraoka et al., 1991, 1992; Teraoka and Tokoro, 1993] in their IPv4 proposal, it does not exactly fall into this framework as it could additionally be ranged into the *Home Agent Framework*.

4.2.2.6 The Multicast Framework

This framework is based on the *Two-Tier Addressing* (figure 4.7) and also relies on network mechanisms developed for multicast routing. Each MN is associated with a multicast address that actually corresponds to a group with only one member. In practice, the MN has three addresses: a permanent IP address, used as the node identifier and recorded in the DNS, a permanent multicast address, used as a location independent and invariant location identifier, and a temporary physical address obtained on each visited link. This last address is used to join the multicast group and as a further location identifier.

The Location Directory records the binding between the permanent IP address and the multicast address. There are no particular requirements upon its location in the framework. The Location Update service is the MN's responsibility (Updating Agent). A multicast address must be obtained first, and registered permanently in the Location Directory. Following this, upon every re-location in a new subnetwork, the MN must join and leave the group with the transcient physical address obtained on the visited subnetwork. As a Locating Agent, the CN calls the Location Directory for the multicast address of the MN. As this is a permanent binding, the Location Directory need only be called once. CNs send packets directly to this multicast address and remain unaware of the physical location of the MN.

The Routing service is performed by multicast routing protocols. A multicast tree is constructed down to the transcient physical address of the MN. Routers don't have the knowledge of the topological location of the mobile, they simply forwards datagrams along the multicast tree towards the MN. Routers can be seen as Forwarding Agents, but in practice they are standard multicast routers with no additional facilities pertaining to mobility support.

Numerous proposals have investigated the use of multicast one way or another to benefit from its diffusion property. This framework is usually used as a means to route data packets from CNs down to the MN (CBTM [Castelluccia, 1998b], MSM-IP [Mysore and Bharghavan, 1997], Helmy [Helmy, 2000], Blazevic [Blazevic and Le Boudec, 1999]). Multicast is also sometimes used to deliver paging messages or

another broadcast-like packet, to provide smooth handoff between adjacent subnetworks (Deadalus [Seshan and Balakrishnan, 1995]).

The *multicast framework* has a number of advantages:

- It achieves the separation of the location identifier and node identifier by allocating an address that exhibits exactly what is needed by mobility: a location independent and invariant addressing.
- It hides local motion to CNs and thus limits signaling and provides location privacy to the MN.
- Ability to support mobility with minimal changes in the infrastructure, since the mechanism may not be deployed only for mobility. Although the multicast technology is still not mature enough, [Mysore and Bharghavan, 1997] suggests that the multicasting infrastructure could be used to solve the problem of mobility support essentially for free. For doing so, the forthcoming development of multicast should address issues that are common to multicast and mobility support. Multicasting and mobile networking indeed exhibit interesting similarities which may be well addressed together. Authors of these proposals advocate that an effective multicast protocol that meets the requirements would solve both the question of multicast routing and mobility support at the same time since it merges the effort made in the two distinct areas.
- Multicast capabilities must be supported by every IPv6 router, which means that subscription to a group and multicast forwarding are presently supported.
- The MN can receive packets on distinct interfaces at the same time if it joins the group with distinct addresses. This could facilitate hand-offs when packets are sent both on the previous visited link and the current visited link.
- Another perceived benefit of multicasting is its desirable properties in terms of resource reservation in future *Integrated Services Networks* [Mysore and Bharghavan, 1997].

As for the drawbacks, there presently doesn't exist a multicast protocol able to support all the usual requirements (scalability, low overhead, ...), particularly as far as Wide-Area Mobility is concerned.

4.3 Conclusion

All proposals without exception handle mobility at the network layer. Most papers generally depict that mobility is essentially an address translation problem [Bhagwat et al., 1996; Ioannidis et al., 1991] that comes from the dual use of the IP address at the network and transport layers. A location independent and invariant addressing scheme must be provided as a means to overcome this. Thus, transparent routing of packets to and from mobile nodes is achieved by two distinct means: most of the initial proposals fit into the *Two-Tier Addressing Category*, while some more recent proposals fit into the *Network-Based Category*.

The frameworks are summarized in tab.4.1, 4.2 and 4.3. Tab.4.4. summarizes all the studied mobility support propals. For each proposal, we list one or more frameworks in which it could be classified (we use *LD*, *Hier* and *3rd Party* for *Location Directory Framework*, *Hierarchical Framework* and *Third Party Framework* respectively). As we see, proposals can hardly be classified in a single framework. The propals where a binding is maintained at the CN belong partly to the *Location Directory Framework*. This is the case for Mobile IPv4 with Routing Optimization, Hierarchical Mobile IPv6, and MIP. Most proposals are able to handle Wide-Area Mobility, but we advocate that most of them fail to scale to a wide-area network due to the amount of signaling generated in the backbone. Lastly, the routing is nearly optimal when packets don't have to transit trough a virtual or actual home network but have to be encapsulated between MAs.

Proposals in the *Network-Based Category* do not need a Location Directory nor a Location Update Protocol. The mobile node retains its address and the network adapts to topology changes. On the other

Category	LUP	UA	LA	RA	FA	LD
Network-Based	routing protocol	routers				routing table
Two-Tier Addressing	dedicated protocol	MN	dedicated server(s)			dedicated server(s)

Table 4.1: Categories

Framework	LUP	UA	LA	RA	FA
Routing-based	none (routing protocol)	all routers	all routers	none	all routers
Broadcast-based	none	none	none	none	all routers
Location Directory	dedicated protocol	MN	CN	none	none
Third Party	dedicated protocol	MN	Third-Party		none
Home Agent	dedicated protocol	MN	HA		none
Hierarchical	dedicated protocol	MN	a hierarchy of Third-Parties		
Virtual Network	dedicated protocol	MN + 3rd-Parties	a set of Third-Parties		
Multicast	dedicated protocol	MN	CN only once	none	multicast router

Table 4.2: Taxonomy of Frameworks - Location Update Protocol

hand, proposals in the *Two-Tier Addressing Category* achieve the separation of the node identifier and the location identifier by assigning two addresses to a mobile node: a Location Directory is needed to record bindings between the node identifier and the location identifier, an a Location Update Protocol to maintain up-to-date bindings and locate the mobile node. Simply speaking, mobility is hidden to the end-nodes and supported by the network in the first category, whereas mobility is hidden to the network and supported by the end-nodes and a limited number of dedicated servers in the second category.

The first category has a few perceived advantages over the second. The first category avoids the overhead introduced by mobility management signaling, encapsulation, source routing and triangle routing, at the expense of flooding data packets (*broadcast-based framework*) or at the expense of flooding *host-specific routes* (*routing-based framework*). It avoids single points of failure since routing protocols are usually designed to adapt quicky to topology changes under link failures. Second, it facilitates *quality of service* (QoS). The second category indeed maintains state in the network transparently to the routers, which makes traffic reservation harder. Third, it is very questionable that solutions falling into the *Two-Tier Addressing Category* are able to scale to a large number of mobile nodes due to the amount of signaling.

On the other hand, the most important drawback of the first category is that there is no support for Wide-Area Mobility, which means another mechanism is needed to support Wide-Area Mobility. This second mechanism is likely based on the second category. It seems therefore interesting to separate Local-Area

Framework	LD	Binding
Routing-based	none(routing table)	none (host-specific routes)
Broadcast-based	none	
Location Directory	dedicated server + CN	node identifier $\mapsto$ temporary physical address
Third Party	Third Party	forwarding address $\mapsto$ temporary physical address
Home Agent	HA	home address $\mapsto$ temporary physical address
Hierarchical	distributed between 3rd-Parties	forwarding address $\mapsto$ forwarding address
Virtual Network	distributed between 3rd-Parties	vitual address $\mapsto$ temporary physical address
Multicast	dedicate server	permanent address $\mapsto$ multicast address

Table 4.3: Taxonomy of Frameworks - Location Directory

Proposal	Framework(s)	Wide-Area Mobility	Direct Routing
Mobile IPv4	Home Agent	not scalable	no
Mobile IPv4 with Route Optimization	Home Agent + LD	not scalable	yes
Mobile IPv6	Home Agent + LD	not scalable	yes
Hierarchical Mobile IPv6	Hierarchical + Home Agent + LD	yes	nearly
Sunshine	Location Directory	not scalable	yes
LSR	Home Agent + LD	not scalable	yes
VIP Sony	Home Agent + Virtual Network	not scalable	nearly
LINA / LIN6	Location Directory	yes	yes
Columbia	Virtual Network (+LD)	not scalable	nearly
Cellular IP	Hier: Routing-based + Home Agent	yes	no
HAWAII	Hier: Routing-based + Home Agent	yes	no
INRIA CBTM	Hier: Multicast + Home Agent + Virtua	yes	nearly
MSM-IP	Hier: Multicast + LD + Virtual	yes	nearly
INRIA HMIPv6	Hier + Home Agent + LD	yes	nearly
CLNP	Hier: Routing-based + 3rd Party	yes	yes
Hierarchical Foreign Agents	Home Agent + Hierarchical	not scalable	no
Caceres	Home Agent + Hierarchical	yes	no
Awerbuch and Peleg	Location Directory + Hierarchical	yes	yes
LAR	Home Agent + Multicast	not scalable	yes
Deadalus	Home Agent + Multicast	no	no
Helmy	Multicast	no	no
DCM	Multicast	yes	yes
Mobile Internet	Hier: Location Directory + Routing-based	yes	yes
DNS	Location Directory	not scalable	yes
Mosquitonet	Home Agent		

Table 4.4: Taxonomy of Proposals

Mobility from Wide-Area Mobility, and there are other reasons for doing so. As a study shows, there is a geographic locality in user mobility patterns [Kirby, 1995]. Furthermore, as expressed in [Caceres and Padmanabhan, 1996; Castelluccia, 1998c], when a mobile node visits a domain, correspondent nodes in other domains do not need to be aware of the mobile node's motion within its domain. The use of a hierarchical scheme to separate both types of mobility allows to confine Local-Area Mobility with the domain only and transparently to the global network.

Therefrom, a hierarchical combination of the frameworks emerges as a requirement to support mobility. Wide-Area Mobility seems only achievable by means of the *Two-Tier Addressing Category* while any other technique could be used to support Local-Area Mobility. The use of multicast is also promising, but scalable and cheap multicast is required.

Part II

Network Mobility Support

Chapter 5

Problem Statement and Requirements

The purpose of *host mobility support* is to provide continuous Internet access to mobile hosts. In contrast to *host mobility support*, *network mobility support* is concerned with situations where an entire network moves and attaches to different locations in the Internet topology. Such a network is referred to as a mobile network. *Network mobility support* must be considered separately from *host mobility support* because the addressing, locating and routing issues caused by mobility are not the same whether the mobile entity under consideration is a host, a router, a subnetwork, or a set of subnetworks.

In this chapter, we will first define a work context that describes the goal we want to achieve and limits the scope of our study (section 5.1.1). This allows us to depict a number of characteristics (section 5.1.2) and a number of issues we are faced with (section 5.2). We then investigate what constraints limit the implementation and the deployment of a potentially and ideally good solution, and what requirements solutions must comply with. Most constraints and requirements that we have listed under section 5.3 are equally applicable to *host mobility support* and *network mobility support*. Some of them have been debated in the literature as far as *host mobility support* was concerned; we have extended this list to include those related to *network mobility support* according to the scope or our study. Section 5.4 reviews the brief literature for supporting mobile networks and exhibits that current schemes do not meet our requirements. Section 5.5 discusses a few potential approaches. We conclude that a Mobile IPv6-based solution should be investigated in priority.

Most of the material presented in this chapter, together with the terminology defined in section 1.3, has been submitted to the IETF in our *internet-draft* [Ernst et al., 2001b] for the consideration of the Mobile IP Working Group. It currently serves as a basis for the definition of the problem scope and solution requirements in on-going discussions.

5.1 Scope of our Study

5.1.1 Objectives

The purpose of this study is to provide mobile network nodes with an uninterrupted Internet access and to route datagrams between correspondent nodes and mobile network nodes by the most optimal path in both directions. We will limit ourselves to mobile networks that are stub networks, i.e. the mobile network does not forward traffic not intended to itself. We are not going to consider the case where a fixed router suddenly becomes mobile, nor the case where a mobile router is detached from the mobile network it is responsive from. We are either not going to consider an entire administrative domain that changes its point of attachment or mobile networks rarely changing their points of attachment. This

present study won't either consider issues specific to multi-homing or *ad-hoc* networks.

5.1.2 Characteristics

The instances of mobile networks highlighted in section 1.3 justify the need to consider potentially large mobile networks containing hundreds of hosts and several routers. The train example highlights that the number of correspondent nodes could also be very large, and that these may be sparsely distributed in the Internet. It also justifies the need for true worldwide mobility in the Internet. A mobile network may attach to very distant parts of the Internet topology, provided it is granted access to it, therefore requiring both Local-Area Mobility support and Wide-Area Mobility support. As exhibited by our scenario, there is also a desire to achieve two levels of mobility: *node mobility* and *network mobility*, since a mobile phone carried by the passenger is mobile with respect to the aircraft or the train (i.e. a VMN or even a mobile IP-subnet) (the passenger is carrying a PAN), which is itself mobile with respect to the Internet. Moreover, since people move from a mobile network to another, and since instances of mobile networks like trains, car, aircrafts cross country boundaries, both mobile nodes and mobile networks are also most likely to cross ISP boundaries and therefore to move between topologically distant parts of the Internet.This means we must allow VMNs belonging to potentially different administrative domains to visit the mobile network. Under those circumstances, we make the following observations:

Structure of the Mobile Network A MR changing its point of attachment does not cause the associated MNNs to change their own physical point of attachment. Thus, the internal structure of a mobile network is not modified as a result of the mobile network changing its point of attachment.

Mobile Router is a Transit Point All packets sent from a CN to a MNN necessarily transit through a MR. As a result, providing the CN with the current topological location of the MR may be sufficient for optimally routing packets intended to a MNN. From a routing perspective, a mobile network could be virtually perceived as a single node (the MR) with a topologically correct address or prefix. If the mobile network is multi-homed, it could be virtually perceived as a single multi-homed node (the MR) with several topologically correct addresses or prefixes.

Size of the Mobile Network A mobile network may comprise one or more subnetworks. Its size could scale from a sole subnetwork with a few IP nodes, such as in the case of a PAN, to a collection of subnetworks with hundreds of IP nodes, such as in a train.

Large Number of CNs A mobile network may have a very large number of CNs. For instance, each passenger in a train may be considered a MNN. Each of them may be communicating with a few CNs. As a result, the total number of CNs could be several times as large as the number of MNNs and scale up to a few thousands.

Sparseness of the CNs CNs are typically sparsely distributed in the Internet and belong to distinct administrative domains.

Handoff Frequency All mobile networks may not move with the same speed and frequency. For instance, a PAN connected to the Internet via a 802.11b WLAN, or a car connected to the Internet via GSM are likely to change their point of attachment very frequently, while an aircraft or a boat may be connected to the Internet via the same satellite link for a couple of hours. In addition, mobile networks may not move at all for a large amount of time.

Routers in the Mobile Network All routers in the Internet are considered to run a number of protocols such as a routing protocol, Neighbor Discovery, ICMP, and others. This also applies to routers in the mobile network, including the mobile router.

5.2 Issues

Network mobility support is not only concerned with providing a location invariant identifier, determining the topological location and routing packets to the current location as this have exhibited for mobile nodes. Because we move a router and all nodes behind it, we are indeed faced with new routing, addressing, and security issues. Moreover, moving an entire network has an impact on a number of network-layer protocols. This section therefore lists a number of major issues faced by *network mobility support* specifically. Since this topic is really new, more issues would probably be raised in the near future once discussion becomes more active at the IETF.

5.2.1 Routing Issues

All packets sent to a mobile network node must transit through the current AR of the mobile network and the mobile router itself. As a result of mobility, the path to the mobile network is varying according to the AR to which the MR is currently attached. We have to investigate how this path could be determined in order to route packets via the most optimal path. Particularly, we need to examine if this is best solved by routing protocols or by some transient means as this is done for mobile hosts. We need to investigate:

- if there is a need for a CN to find out that the node it is corresponding with is in a mobile network.
- if there is a need to determine the actual topological location of the mobile network or the MNN.
- if there is a need to determine the AR the mobile network is currently attached to.
- if CNs should be aware of the topological location of the mobile network or the MNN or if this should this be transparent to them.
- if forwarding should be established from a former AR to a latter one.

5.2.2 Addressing Issues

Impact on the MR Following existing IPv6 specifications, any host, is in theory required to obtain a topologically correct address on the link on which it is currently attached to. We must investigate if this can alternatively be done for a single host or for a router. If yes, this means that the external interface of the mobile network is configured with the foreign prefix. For instance, we need to examine if in this situation a MR can perform address auto-configuration as any MN. We also have to investigate if the configuration of the MR's interface with a new address has an impact on MNNs.

Impact on MNNs Since MNNs don't actually change their own point of attachment, it is very questionable whether MNNs should also get a topologically correct address that actually reflects their topological and hierarchical location in the Internet. If we conclude that MNNs should get a topologically correct address, we have to determine how this could be performed internally in the mobile network, and how the impact of the address change at upper layers would be overcome. If we renumber, we have to investigate where to advertise the new addresses if we want to maintain connections; if we do not renumber, we have to investigate how to perform optimal routing between CNs and MNNs.

Topologically Correct Addresses The above paragraph questions what is the scope of a topologically correct address. Should it be topologically correct within the mobile network, within the administrative domain, or within the whole Internet. We then have to investigate the suitable IP addressing for the support of mobile networks. Particularly, we may investigate if the Aggregatable Global Unicast address format is a suitable address format or if a new address format is required to facilitate network mobility support.

5.2.3 Network Protocols Issues

As seen in section 5.1.2, all routers in a mobile network are routers like the others and have to run a number of protocols. Following the existing IPv6 specifications, they particularly should run at least an instance of a routing protocol, and other protocols like Neighbor Discovery, etc. We therefore have to investigate how the network protocols running in the mobile network must interact with the network protocols running in each subsequent visited network and how the mobile router is going to interact with the ARs it attaches to. This raises a number of issues for each network protocol, as listed in the following sections.

Impact on Neighbor Discovery One task of Neighbor Discovery is to send Router Advertisements and Router Solicitations. When the mobile router is attached to some AR in a visited network, it should receive such Router Advertisements from its current AR. We have to investigate how those Router Advertisements should be processed by the mobile router and how the mobile router should interact with this instance of the protocol running at the AR. We also have to investigate what is the impact on this protocol when the mobile network leaves its point of attachment.

Impact on the Visited Network We have to investigate if the subsequent ARs and the other routers in the visited network should be aware that the visiting mobile node is a router and not a host. In addition, we have to examine if it is necessary to let them know that there is an entire network behind the mobile router. In such a case, a network route may have to be propagated in the visited network and this raises an additional number of issues as discussed in the section about routing protocols.

Impact on Routing Protocols We have to investigate how the mobile router interacts with the routing protocols running at each of its subsequent ARs. The impact may not be the same whether the mobile network is limited to a single IP-subnet or a number of IP-subnets. Indeed, a single mobile IP-subnet may not need to run an instance of a routing protocol whereas a mobile network comprising more than one router is necessarily running one. We have to evaluate what kind of routing protocols may run in a mobile network and how it interacts with the routing protocol running at each of its subsequent ARs.

- In case the mobile network is running the same routing protocol as its ARs, it is questionable whether the mobile network should participate or not in the routing protocol running in the visited network. If it does, the mobile network can be seen as a partition of the local network. The topology computed by the routing protocol becomes more dynamic and we have to evaluate how existing protocols are able to handle this case. Moreover, mobility may cause a routing table partition.
- In case the mobile network doesn't participate in the routing protocol running in the visited network, the mobile network can be seen as a kind of Autonomous System that is running an instance of an Internal Gateway Protocol.

In both cases, we must evaluate the interaction between the routing protocol running in the mobile network and the routing protocol running in the visited network. In addition, we must determine what routing protocol is suitable within the mobile network. We also have to evaluate the impact on routing protocols when the mobile router is multi-homed, when the mobile network comprises more than one mobile router, and when the mobile network itself receives a mobile network.

5.2.4 Security Issues

All security concerns that usually apply to *host mobility support* also apply to *network mobility support*. In addition, *network mobility support* faces a number of additional ones that complement the addressing issues, network protocols issues, and routing issues depicted in the previous sections. In order to assist *nested mobility*, we face administrative and access control concerns for VMNs and potential *mobile IP-subnets*.

5.3 Design Requirements

Solutions for *network mobility support* must or should fulfill a number of requirements, whereas a number of constraints may limit the deployment of a potentially good solution. The purpose of this section is to detail an initial list of such requirements and constraints that we propose; once more people are joining us to work on *network mobility support*, more will likely be added. Basically, most requirements discussed for *host mobility support*, like for instance in [Bhagwat et al., 1996], [Myles and Skellern, 1993b], [Castelluccia, 1997] or [Caceres and Padmanabhan, 1996], are equally applicable to *network mobility support*. However, they must be refined to comply with the specific characteristics and issues that apply to mobile networks. In addition to this refinement, we have also extended this list with a number of new requirements peculiar to mobile networks.

5.3.1 Wide-Area Mobility

Permanent and uninterrupted world-wide mobility requires the ability to move between heterogeneous networks, i.e. Wide-Area Mobility. Nothing but administrative and security policies should prevent a mobile network to attach anywhere in the `Internet` topology. In practice, a given mobile network must be able to roam between administratively distinct access networks and via any available access technology (`802.11b WLAN`, `Bluetooth`, satellite link, `GSM`, ...). In addition, network mobility support must also accommodate correspondent nodes deployed in distinct administrative domains. This requires a unified mobility support scheme. We must avoid a situation where distinct *network mobility support* schemes deployed in distinct access networks are unable to inter-operate with each other. This lack of standardization clearly happens in cellular telephony where Japaneses or Americans cannot use their mobile phones in Europe because their phones do not speak the `GSM` protocol. This also happens when the mobile phone user is unable to connect its phone to distinct operators that compete in the same country. Then, not only standard between countries and organizations is required, but also between networks in which different policies may apply. For doing so, we need a network mobility support scheme that fits well into the existing standards, that is easy to deploy and that does not require too many additional services in the network.

5.3.2 Optimal Routing

Non-optimal routing increases bandwidth consumption and transmission delays. The amount of traffic intended for the mobile network is understandably more significant than in the case of a single mobile node, then non-optimal routing is an even more important requirement that it was already for *host mobility support*.

5.3.3 Minimum Signaling Overload

Routing packets efficiently from a CN to the current location of the mobile network is usually performed at the cost of control traffic. The cost of this control traffic has to be balanced against the expected gain

of optimal routing. Minimizing the amount of control traffic has always been an important concern for host mobility support. Due to a potentially large number of CNs, this becomes an even more important requirement for network mobility support.

5.3.4 Scalability

Scalability has always been an important concern in the design of new protocols. As far as host mobility is concerned, mobility support has to take into consideration a growing number of mobile nodes and should even assume that a major part of the nodes composing the Internet are mobile in the near future. This means that signaling load and memory consumption should scale to an important number of mobile nodes. Network mobility support has to address scalability differently, and in three ways:

- Scaling to a large number of mobile networks.
- Scaling to a large number of correspondent nodes.
- Scaling to the size of large mobile networks comprising several subnetworks and hundreds of MNNs.

Network mobility support must thus be able to support large mobile networks containing hundreds of nodes like a train and corresponding with thousands of CNs, and a very large number of small mobile networks such as PANs comprising a few IP nodes. Scaling to a large number of CNs indeed deserves more consideration than scaling to a large number of mobile networks. Moreover, the number of CNs is somewhat a function of the size of the mobile network; the more MNNs we have, the more CNs we are likely to have. Scaling to a large number of CNs is a requirement that has never been considered in *host mobility support*, to the contrary of the ability to support a large number of mobile nodes.

5.3.5 Transparency

Mobility management transparency for MNNs We have seen in section 5.1.2 that MNNs appear to move from the point of view of their CNs although they don't change their own point of attachment within the mobile network. Mobility management of a mobile network is therefore better seen as the mobile router's responsibility and should be transparent to MNNs. MNNs should better have no responsibility in network mobility management, particularly LFNs and LMNs. They should only be concerned about managing their own mobility if they themselves appear to change their point of attachment. However, MNNs may encounter variable delays of transmission and loss with their respective CNs as the network is moving,

Migration Transparency Mobility support must maintain continuous access to the Internet without disruption of service for MNNs regardless of the location of the mobile network. This means that all MNNs must always be reachable regardless of the point of attachment of the mobile network. In addition, there shouldn't be an abrupt interruption of the IP sessions. In practice, mechanisms should be added to forward packets in transit to the current location of the mobile network.

Operational Transparency The application or the user must not perform any action to remain connected to the Internet as a result of network mobility. This means that the network layer is solely responsible to support Internet access from and to the mobile network in an absolute transparent manner to the application or the user.

Performance Transparency (Seamless Mobility) Network mobility support should exhibit low latency, incur little or no data loss, minimum delays, minimum signaling load, and minimum bandwidth consumption for datagrams delivery. The solution is termed "efficient" provided network mobility is supported without performance degradation of the Internet. Loss and delays should indeed range into those experimented for communication flows between two fixed nodes. Moreover, both Local-Area Mobility and Wide-Area Mobility need to be handled as efficiently. At last, the addition of network mobility support shouldn't impact the performance of upper layers. In order to limit losses during hand-offs and to avoid degradation of performance at the upper layers, it may be necessary to perform seamless handovers.

Layers Independence Handover of IP sessions is performed at the network layer. With respect to the layer separation of the Internet protocol suite, handover must be managed at the network layer only and transparently to upper layers, despite of the migration of the mobile network in the network topology. Therefore, a change of topological location must not have an impact on layers above the network layer other than a transient loss of performance, as depicted in the above paragraph. If this is respected, compatibility with existing transport and application layers is maintained. In practice, the node identifiers used at the transport layer should be independent from the physical IP addresses used at the network layer for routing. If upper layer protocols require a location independent and invariant identifier, the network layer must provide it with an identifier irrespective of the actual topological location.

5.3.6 Nested Mobility

Network mobility support must allow VMNs to enter the mobile network and LMNs to leave it. Both mobile networks and mobile nodes should therefore be considered at the same time. This means that communications with VMNs and LMNs should be handled as efficiently as with LFNs. *Network mobility support* should also allow a mobile network to visit another mobile network (this is the example of a PAN in a train). This latter case may however not allow for recursive nested mobility and may be limited to simple instances of mobile networks.

5.3.7 Mobile CN

Network mobility support must be optimized to handle the case where the CN is itself a mobile node or located in a mobile network(particularly if it is a visiting mobile node). It must perform efficiently in both cases.

5.3.8 Backward Compatibility

Network mobility support is constrained by backward compatibility with existing and forthcoming IPv6 standards. As such, it should not prevent MNNs from operating any standardized protocol. To ensure backward compatibility, the solution for supporting mobile networks must provide the necessary extensions to existing protocols if needed. This, among other, includes Mobile IPv6 and IGMP:

- Mobile IPv6: host mobility support in IPv6 is achieved by Mobile IPv6 which must be supported by any IPv6 implementation. Thus, mobile nodes located in a mobile network (i.e. LMNs and VMNs) must still be able to operate Mobile IPv6 once they move in, within, or away from the mobile network.
- IGMP: any IPv6 router is supposed to allow hosts on its attached subnetworks to participate in multicast sessions. Group membership is gathered by the IGMP protocol.

5.3.9 Minimum Impact on Existing Protocols and Infrastructure

Minimum impact on the already deployed infrastructure was an important issue as far as IPv4 was concerned since it was not possible to require all hosts to implement new features[1]. On the other hand, the emergence of IPv6 is an opportunity for making changes, if necessary. An important number of specifications has already been defined, but there is still scope for adding new capabilities if needed since IPv6 deployment is only dawning. However, in order to provide a quickly deployable solution, *network mobility support* should better make use of the existing protocols whenever possible and impose minimum changes or extensions on the existing ones. It is also desirable to minimize infrastructure installation costs and complexity.

5.3.10 Security

Network mobility support must comply with usual IPv6 security policies and standardized protocols. In addition, and unlike fixed nodes, mobile network nodes are more exposed to security threats, particularly identity usurpation. Network mobility support must provide MNNs and their CNs with at least as good security as for fixed nodes and mobile hosts. It particularly shouldn't leave more room for intruders to usurp an identify nor to perpetrate any kind of attack against the MNNs nor the CNs. In practice, all control messages required by network mobility support must be exchanged in a secure manner and must ensure the following:

- Confidentiality All control messages transmitted in the network must ensure MNNs'confidentiality. Only the recipient of the control message may be able to decrypt the content of the datagram.
- Authentication All control messages must be authenticated by recipients in order to prevent intruders to usurp the identity of a MNN.
- Authorization The recipient of a control message must ensure that the sender is effectively authorized to perform the mobility support operation indicated in the control message.
- Location Privacy Network mobility support must provide means for MNNs to hide their location to any third party. It shouldn't be possible to determine the succession of topological locations of the mobile network or a particular MNN by monitoring the exchange of control messages. In practice, MNNs may wish to hide their location to some or all of their CNs, or anyone else but the CNs. It would also be desirable to hide the location of the entire mobile network to all CNs without discrimination between MNNs.

5.3.11 Addressing Constraints

The IP address is used for routing and to identify the subnetwork where an interface is attached to. Every interface must therefore be configured with the network prefix of the current subnetwork.

5.4 Network Mobility Support in the Literature

The study of mobility support proposals in chapter 3 showed that mobile networks are only mentioned in two proposals, namely Mobile IPv4 (section 3.1.1.3) and Hierarchical Mobile IPv6 (section 3.1.2). There is also some interesting work turning around *ad-hoc* networking that worth mentioning. None

[1] Interesting to note, some details of the IPv4 specification are not fully implemented. For instance, Loose Source Routing was specified as a requirement in RFC1122 [Braden, 1989], but only a few routers have efficiently implemented it)

of the two proposals solve the issues we raised in section 5.2 nor they meet the requirements defined in section 5.3.

5.4.1 Mobile IP and Mobile Networks

Mobile IPv4 [Perkins, 1996a] and Mobile IPv6 [Johnson and Perkins, 2000] have introduced *host mobility support* for IPv4 and IPv6 respectively. Although is it claimed that it could support mobile networks equally as mobile nodes ([Perkins, 1996a] section 4.5, [Perkins, 1998] section 5.12, [Solomon, 1998] section 11.2), we argue that they cannot be supported efficiently. A means to support LFNs and VMNs is suggested, but not really detailed. Indeed, the specification more provides hints rather than a workable solution. The use of routing protocols to provide connectivity is mentioned, but this introduces the same issues as mentioned in section 5.2. One of the most obvious drawbacks is the lack of routing optimization. Routing optimization between correspondent nodes and mobile network nodes is the feature that raises the more issues and this is probably why mobile networks are not considered explicitly in Mobile IPv4 with Routing Optimization nor Mobile IPv6 anymore.

5.4.2 Hierarchical Mobile IPv6

Hierarchical Mobile IPv6 Extended Mode (see section 3.1.2) proposes a solution to support for mobile networks. This was first claimed in the second version of this work in progress and following our first intervention on *network mobility support* at the IETF in August 2000 (i.e. our solution that is discussed latter in section 7.1). However, this proposition is only targeted to VMNs and does not address the whole issue: it does not handle redirection of packets from the MR's HA, nor optimal routing for LFNs, not other issues discussed latter in this document.

5.4.3 MIPMANET

[Jonsson et al., 2000] studies how to connect an *ad-hoc* network to the Internet in order to allow a mobile node to enter and leave the *ad-hoc* network while performing Mobile IPv4. It does however not study the issue of moving an entire *ad-hoc* network and attaching it to distinct points of attachment, which would raise similar issues than for a mobile network.

5.5 Potential Approaches

In this section, we investigate how we could achieve *network mobility support*. During the course of this study, we came across a few possible approaches to tackle with the above listed issues and requirements. Most approaches we had in mind derive one way or another from our study of the *host mobility support* proposals (chapter 3), and from our classification (chapter 4). The first approach we came through is based on the *Location Directory Framework* and we investigated the ability of the DNS to record the current topological location of the mobile network. The second approach is based on the IPv6 *renumbering* mechanism and could be ranged into a new category of the *Network-Based Framework*. In the third, we investigated how the *Routing-Based Framework* could be applied to *network mobility support*, and lastly, we investigated a *Two-Tier Addressing Framework*.

Since distinct address formats could coexist simultaneously in IPv6, we have also investigated how a new address format could be specifically designed to support network mobility. As such, we have considered a geographical address format that would reflect the actual position of the mobile network on Earth. However, we have decided to left it out of this present study because it seemed hardly achievable.

5.5.1 DNS-based Approach

We first considered a solution based on the DNS to record directly in this database mappings between the *domain name* of the mobile network and its topological location. This approach indeed falls into the *Location Directory Framework* that we have defined in section 4.2.2.1. The topological location could be a forwarding address (the address of a *Third-Party*). Upon query for the topological location of a MNN, the DNS would return its IPv6 address (used as the node identifier) and the forwarding address for the mobile network (used as the location identifier). This mapping would be updated by the MR itself and propagated by the DNS up to the CNs.

However, the DNS is not adapted to frequent updates, thus it is not applicable to mobile networks that frequently change their point of attachment. It indeed takes some time before the former prefix advertised in the DNS is depreciated and effectively replaced by the new one in all existing caches. The lifetime of the prefix and addresses could actually be set to a very low value, but still, the DNS is not designed to propagate frequent address changes. Last, there is no mechanisms embedded in the DNS to deliver updated mappings to the caches at the CNs nor to force them to check for new mappings, other than relying on the lifetime of the record. Thus, the new address cannot be propagated fast enough unless the DNS is enhanced for doing so. Such enhancements would surely impact its design philosophy.

While a solution based on the DNS is clearly inappropriate for frequent moves, thus for Local-Area Mobility, it still may be adequate to support Wide-Area Mobility. This advocates for a hierarchical scheme where another means is used for Local-Area Mobility.

5.5.2 Renumbering-based Approach

We then considered a solution based on the existing IPv6 renumbering techniques. IPv6 renumbering allows an entire network to change its address automatically. With this mechanism, the MR could obtain a new topologically correct mobile network prefix on the current visited link and advertises it in the entire mobile network. All MNNs would renumber to a new address that complies this new prefix. The mobile network prefix would then be used as a routing directive since it identifies the topological location of the mobile network.

This approach faces a number of issues. First, the prefix advertised on the visited link must be large enough to comply with the size of the mobile network and its architecture. This questions how a mobile network would negotiate a suitable prefix. Second, renumbering is not a mechanism suitable for networks that renumber frequently. From the discussion on the IPv6 mailing list, we note that networks shouldn't renumber more than about once a week. The reason is that the new prefix must be propagated in the DNS and this takes time before all the depreciated addresses are removed from caches, as said in section 5.5.1. Third, as a result of an address change, we fall into the same problem as encountered for mobile hosts and described in section 2.4.1: how on-going communication could be maintained between CNs and MNNs.

To conclude, renumbering alone does not solve the question of address change. We still lack a permanent node identifier. Additional mechanisms are therefore needed, particularly a Location Update Protocol.

5.5.3 Routing-based Approach

Our third approach considered a solution that relies on routing protocols to advertise in the Internet a route to the mobile network via the current AR. As debated in section 4.3, this avoids the burden of the change of address. In the case of a mobile network, a *network-specific route* would be propagated in the network instead of a *host-specific route*.

Existing routing protocols deployed in the Internet are optimized to quickly adapt to topologies changes in

order to avoid network partition and packet losses. However, this quick process is usually performed at the cost of signaling. Topology changes are propagated in the topology and cause a signaling burst. Indeed, routing protocols are optimized to perform well on topologies with minimum changes and not for highly dynamic topologies such as topologies comprising mobile nodes. In order to support mobility, the routing protocol need to be optimized for highly dynamic network topologies, to handle temporary loss of routes during handover. Mobility management by means of routing protocol indeed incur more requirements on the routing protocols like smooth handover, packet loss, fast handover. It is questionable whether this should be the role of the routing protocol to take care about this. Furthermore, these routing protocols do not take into account the characteristics of the wireless media. Wireless links have limited bandwidth, experiment longer delays and are more lossy than wired media. In addition, handover may also cause more packet losses, which would in turn cause more retransmissions.

Host mobility proposals that fall into the *Routing-Based Framework* are clearly developed for mobile *hosts* only: they introduce *paging* as a means to reduce topology updates. The network searches for the topological location of the mobile node only when there is actual traffic. In the mobile network case, a mobile network could only be considered as *idle* if no MNN engage in any communication outside the mobile network (see section 1.3).

Relying on routing protocols to achieve Local-Area Mobility of a mobile network seems achievable, as this discussed in section 4.3. This would first require to evaluate `Internal Gateway Protocols`'s ability to support frequent topology changes, and probably to design new protocols better adapted to dynamic topologies. On the other hand, in order to achieve Wide-Area Mobility, the mobile network prefix would need to be propagated in the entire Internet, or at least in all routing tables of the `External Gateway Protocol` (likely `BGP`) running in the Internet backbone. This obviously wouldn't scale to a large number of mobile networks since each instance of the routing table would contain one entry per mobile network. Moreover, the Internet is hierarchical. All networks with the same prefix are aggregated in a single entry. This allows routing aggregation. Propagating a route to the mobile network prefix would break this routing aggregation. Thus, Wide-Area Mobility could only effectively be achieved by means of additional mobility support features. *Network mobility support* can therefore not only rely on routing protocols.

5.5.4 Two-Tier Approach

Our fourth approach considered how the *Two-Tier Addressing* concept could be applied to *network mobility support*, as a means to overcome the generic addressing problem. In this approach, each MNN would be associated with a permanent address, used as the node identifier, and with a temporary address, used as the location identifier. This location identifier would then be advertised to a number of nodes in the network.

The most advanced solution for *host mobility support* that falls into the *Two-Tier Addressing Framework* is `Mobile IPv6`. Moreover, it should become a standard very soon, i.e. a base protocol that every `IPv6` stack must implement. As such, the suitability of this protocol for *network mobility support* should be considered and evaluated in priority. At first sight, `Mobile IPv6` doesn't fully meet all the design requirements described in 5.3. One of the main drawback of `Mobile IPv6` is its scalability issue. It provides for optimal routing between any two nodes at the expense of an important signaling cost throughout the entire network. Every home agent and correspondent node must constantly be notified the primary careof address of the mobile node. Even if the mobile node doesn't move, the registration must be refreshed periodically. In this situation, `Mobile IPv6` hardly scales to a very large number of mobile nodes communicating with a large number of correspondent nodes, whether mobile nodes are frequently changing or not their point of attachment. As shown in [Castelluccia, 1998c], most of the available bandwidth may be solely consumed by `Mobile IPv6` signaling if a significant part of the end-nodes become mobile. This scaling issue deserves even more consideration if we want to support a large number of mobile networks with a large number of correspondent nodes. As an attempt to overcome this scalability issue for *host mobility support*, `Hierarchical Mobile IPv6` [Soliman et al., 2001] (see section 3.1.2) proposes to confine locally signaling pertaining to Local-Area Mobility. Then, a combination of the two protocols could also be evaluated for *network mobility support*.

5.5.5 Conclusion

Most of the approaches described in this section somehow follow the tracks already investigated in the literature for *host mobility support*. Mobile IPv4 resulted from these past discussions as an easy solution for an immediate need. However, as for *network mobility support*, the above sections exhibit that we don't have anything near at hand. A more detailed analysis is required. Basically, none of the potential approaches is straight forward.

Despite this, a Mobile IPv6-based approach seems promising because it surely has less impact on the currently deployed set of protocols than the other envisioned approaches. Thus, we think we should pay more attention to Mobile IPv6, and we should therefore study its ability and shortcomings more carefully, all the more a section in the Mobile IPv4 specification suggests a means to support VMNs (see section 3.1.1.3).

From the above discussion and section 4.3, we conclude that the ideal solution is probably a mix of some of these approaches, where one approach would rather be used for Local-Area Mobility and another one for Wide-Area Mobility. For instance, we may think about the following combinations: a Mobile IPv6-based solution for Local-Area Mobility combined with Hierarchical Mobile IPv6 for Wide-Area Mobility, or a Mobile IPv6-based solution for Local-Area Mobility combined with renumbering or a DNS-based solution for Wide-Area Mobility, As we see, this raises potentially interesting research subjects, and we do not pretend to investigate in depth all of them. We would probably better try to use what already exists if we want to avoid to re-invent the wheel. Thus, we should try to base our work on an existing solution, provided it is not proven inadequate. Mobile IPv6's suitability must therefore be investigated first.

5.6 Conclusion

In this chapter, we have seen that *network mobility support* faces a number of issues and that solutions need to meet a number of requirements. *Network mobility support* has deserved very little consideration in the literature up to this date and none of the potential solutions address the whole issue nor meet a significant part of the requirements. We therefore conclude that this question needs to be tackle explicitly. We have therefore considered a few approaches, and we have concluded that none of the them offers a sufficiently straight forward solution. However, a solution based on Mobile IPv6 seems promising as long as this protocol is not proven inadequate. We have therefore decided to investigate a Mobile IPv6-based approach in the rest of this study.

Chapter 6

Mobile IPv6 Shortcomings

Following our conclusion in the precedent chapter, this chapter examines how mobile networks could be supported by Mobile IPv6. While there was an attempt to support mobile networks in Mobile IPv4 (section 3.1.1.3), they are not mentioned anymore in Mobile IPv4 with Routing Optimization [Perkins and Johnson, 2000] nor Mobile IPv6. We therefore study Mobile IPv6's ability for supporting mobile networks and we conclude that the current Mobile IPv6 specification has no provisions to handle them. We first discuss the lack of explicit support for mobile networks. Particularly, we show that the HA is unable to redirect packets intended to mobile network nodes and that there is no way optimal routing could be performed between correspondent nodes and mobile network nodes. Minor adjustments would indeed allow the HA to redirect all packets intended to mobile network nodes. However, those adjustments do not meet most of the design requirements as discussed in section 5.3. As hundreds of correspondent nodes may be communicating simultaneously with mobile network nodes located in the same mobile network, the questions of locating, optimal routing and signaling overload are significantly more critical issues for mobile networks than for mobile nodes. Slightly enhancing the Mobile IPv6 specification for supporting mobile networks without taking into account their specific characteristics (section 5.1.2) and requirements (section 5.3) would not provide optimal routing without overloading the network with Mobile IPv6 signaling. As it will be demonstrated in the following sections, we also face some scalability and security concerns. Extensions specific to mobile networks are therefore needed in order to support mobile networks efficiently.

6.1 Mobile IPv6 and Mobile Networks

As said in 5.4, *network mobility support* is discussed in Mobile IPv4, but not in Mobile IPv6. In theory, Mobile IPv6 could support mobile networks similarly as in Mobile IPv4. Despite this, the Mobile IPv6 specification doesn't mention them anymore. The specification merely states that a mobile *node* may either be a mobile *host* or a mobile *router*. As a result of this, Mobile IPv6 works effectively independently for a host or for a router as long as the final destination of the packet is the mobile router itself, but nothing in the specification tells what should be done for packets intended to mobile network nodes. The following sections examine what happens if the mobile router is operating Mobile IPv6 as is. In fig. 6.1, we have refined our terminology in order to comply with Mobile IPv6.

6.1.1 What the Mobile IPv6 Specification Says

Obtaining a care-of address Obtaining a careof address can be done independently for a host of for a router. The mobile router MR has a permanent home address MR_{ip} on its home link and gets a new

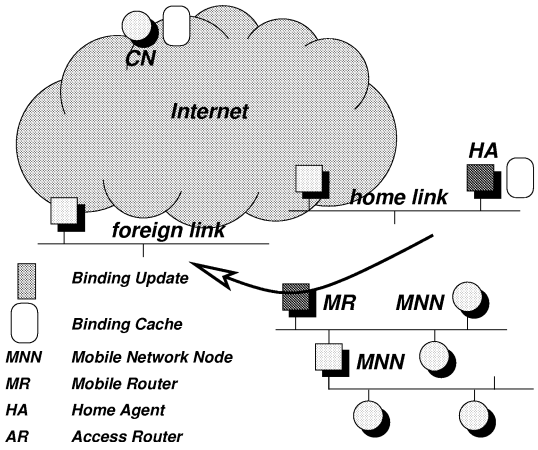

Figure 6.1: Terminology for Mobile Networks Applied to Mobile IPv6

careof address MR_{coa} on each subsequent foreign link it visits as any mobile node would. The prefix of the newly obtained careof address is the foreign prefix.

Registration with HA and CNs As for a standard mobile node, MR constructs a BU containing a mapping between its home address and its primary careof address. It then registers the binding between the MR_{ip} and the MR_{coa} with a HA on its home link. It may also send a BU to each of its *own* correspondent nodes. The recipient of the BU then authenticates the sender. As a result of a successful authentication, which is the case, the home agent and the correspondent nodes of the mobile router add a *host-specific route* for the MR_{ip} to the MR_{coa} as this has already been described in section 3.1.1.5. The home address is then used as the key for searching the Binding Cache ([Johnson and Perkins, 2000] section 4.6). On the other hand, we note that no BUs are sent to correspondent nodes of the mobile network nodes.

6.1.2 Experiment

We have conducted two experiments to highlight the lack of `Mobile IPv6` for network mobility support. Our experiments have been conducted in our lab using the `FreeBSD Mobile IPv6` implementation. The testbed is presented on fig. 6.2. The mobile router MR has two interfaces. The first one is the internal interface and is configured with the home prefix $3ffe:306:1130:100::/64$ as the home address. The second one is the external interface and is configured with the mobile network prefix $3ffe:306:1130:200::/64$. When the mobile network moves and attaches to the foreign link, the internal interface is configured with the foreign prefix $3ffe:306:5555:7777::/64$ as the careof address. In a first experiment, a correspondent node CN in the fixed Internet sends a a packet to MR directly. In a second experiment, CN sends a packet to MNN.

Packets intended to the MR itself In this experiment, CN sends a packet to MR's home address $3ffe:306:1130:100::eui64$. The first packet get routed to MR's home link. When the packet enters the home network, AR checks its routing table for a route to $3ffe:306:1130:100::eui64$. The next hop toward this address is MR. AR sends `NDP` messages on the home link to discover MR's MAC address. HA answers with its address on behalf of MR since it has a host-specific route for MR in its Binding Cache. Packets are therefore correctly intercepted by HA. Packets are then tunneled to the $MR_{coa}3ffe:306:5555:7777::eui64$. Receiving an encapsulated packet is an indication for the MR that the sender does not have a binding for its careof address. MR may decide to send a BU to CN in order to optimize routing. Following packets are then directly sent from CN to $3ffe:306:5555:7777::eui64$ using a `Routing Extension Header` containing MR's home address.

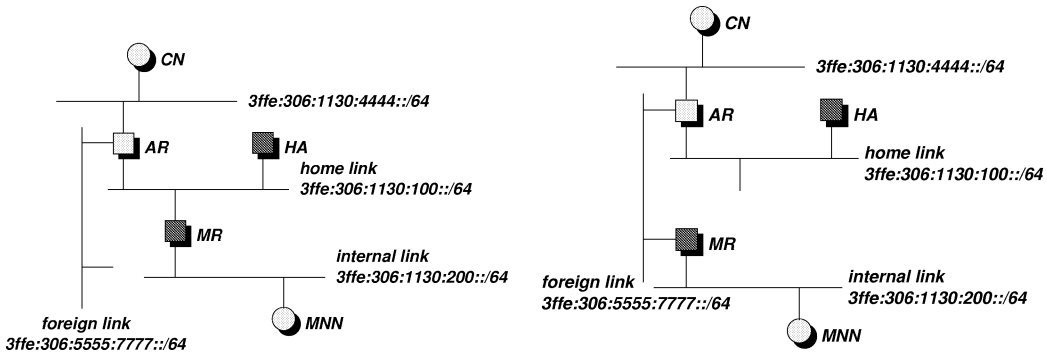

(a) Mobile Network on the home link

(b) Mobile Network on the foreign link

Figure 6.2: Testbed

Packets intended to a MNN behind the MR In this experiment, CN sends a packet to MNN's address $3ffe:306:1130:200::eui64$. The packet get routed to the MR's home link. When the packet enters the home network, AR checks its routing table for a route to $3ffe:306:1130:200::eui64$. The next hop toward this address is MR. AR sends NDP messages on the home link to discover MR's MAC address. HA answers with its address on behalf of MR since it has a host-specific route for the MR in its Binding Cache. Packets are therefore correctly intercepted by HA. The tricky thing follows: although HA is also able to intercept packet intended to MNN, it is unable to encapsulate it to MR's careof address because it does not have a route to MNN. As a result of the mobile router's home registration, the home agent has only recorded in its routing table a *host-specific route* from the MR's home addressto MR's careof address. HA does not know what to do with the packet; it is sent back to the default route (i.e. AR) and the packet enters a routing loop, until the `TTL` expires. A subsequent result is that no communication at all is possible between CN and MNN since its actual topological location is unknown to CN.

6.1.3 Conclusion

From the experiment, we conclude that `Mobile IPv6` is able to support a mobile router but not the mobile network nodes behind the mobile router. The HA is unable to redirect packets intended to the mobile network nodes because HA only adds a *host-specific route* from the MR's home address to the MR's careof address as a result of MR's registration. Indeed, the Binding Cache entry registered by MR worths only for the MR_{ip}. Basically, the *routing loop* demonstrated in our experiment only seems to arise in some implementations, the ones which are in strict conformance with the `Mobile IPv6` specification. Thus, some implementations may allow the HA to encapsulate packets intended to the MNNs, but this means that the `Mobile IPv6` specification is extrapolated. Discussions on the `IETF Mobile IP` mailing list came to the conclusion that the `Mobile IPv6` specification needs to clarify this. We argue that clarifying `Mobile IPv6` is not sufficient to claim that it is able to support mobile networks. Even if the home agent would be able to redirect packets intended to mobile network nodes, we note that there wouldn't be optimal routing between correspondent nodes and mobile network nodes. As this has been discussed in chapter 5, routing optimization is a necessary feature which can not be left aside. Basically, routing optimization prevents `Mobile IPv6` for supporting mobile networks as easily as in `Mobile IPv4`. Lastly, only mobile IP-subnets have been considered so far while larger mobile networks exhibit a number of routing issues, as discussed in section 5.2. So far, these issues have never been discussed. We conclude that the current specification of `Mobile IPv6` is unable to support mobile networks efficiently and needs specific extensions.

6.2 Mobile IPv6 Issues

From the above discussion, we can now conclude that network mobility support using Mobile IPv6 and its existing mechanisms is not that straight forward. In order to keep network mobility support as close as possible to the Mobile IPv6 specification, we need to deliver a careof address to both the home agent and all the correspondent nodes. This raises a number of issues. We fisrt discuss what careof address shall be registered with the home agent and correspondent nodes. We conclude that we should only make use of the mobile router's careof address. We then discusses how the mobile router's careof address could be delivered to the home agent and all correspondent nodes. We conclude that the registration should be made by the sole mobile router in a way for the recipients to understand that this careof address must be used not only to communicate with the mobile router, but also with all mobile network nodes. A number of remaing issues is also listed in the last section.

6.2.1 Security Considerations

Authentication According to Mobile IPv6, each packet that includes a BU option must also include either an AH or *ESP* header providing sender authentication, data integrity protection and replay protection [Johnson and Perkins, 2000]. Authentication of BUs is a necessary task performed by the BU recipient in order to ensure that it has not been sent by an usurper. Without successful authentication, CNs should never send packets directly to the careof address provided in the BU.

Authorization Mobile IPv6 states that BUs cannot be sent by a node on behalf of a mobile node. This raises an issue concerning the identity of the sender of BUs relative to the mobility management of a mobile network and the proper authentication of the sender by recipients of those BUs.

6.2.2 Obtaining a careof address

We have identified two alternatives for obtaining a careof address. In the first one, each mobile network node obtains a new one; in the second one, a unique careof address is obtained by the mobile router.

Individual careof address for each MNN MNNs are not effectively changing their point of attachment and may even not notice that the network in which they are located is mobile. Since this is usually done by monitoring Router Advertisements, this would simply lead to renumbering, which is not exactly what we are looking for. There is no existing means to inform the mobile network nodes that they should obtain a new careof address and register it with their respective home agent and correspondent nodes. Not only this would break mobility transparency (5.3.5) and impose them to take active part in mobility management, but it would break some of the features of the Mobile IPv6 specification.

Single careof address shared by for all MNNs As for an alternative, it makes more sense to use a single careof address for the entire mobile network. The careof address obtained by the mobile router makes a good candidate since the Mobile IPv6 specification effectively allows a mobile router to obtain one as any standard mobile node. Since the mobile router all packets necessarily have to transit via the mobile router, it is sufficient to advertise the mobile router'scareof address in order to route packets to the current topological location of the mobile network.

6.2.3 Registration of the careof address

We have identified three possible alternatives for registering the mobile router's careof address with the home agent and all correspondent nodes. In the first one, the registration is made individually by mobile network nodes; in the second one, the registration is performed by the mobile router on behalf of the MNNs whereas in the third, the registration is still made by the mobile router, but worth for the entire mobile network.

BU processing by individual MNNs As a first trial, we could consider that each MNN is in charge of sending BUs to its respective CNs and HA, i.e. a binding between MNN_{ip} and MR_{coa}. This solution first requires a mechanism to distribute the MR_{coa} to all MNNs. Consequently, MNNs would take part in mobility management, which is not desirable as stated earlier. Moreover, the distribution of the careof address to MNNs also requires two new `Mobile IPv6` entities in order to distinguish between the standard MN Operation, the MR Operation, and the MNN Operation. No changes in the CN Operation are required. Then, all MNNs would send a BU to both their HA and CNs. As a result of these registrations, the HA would defend each individual home address on the home link by means of a gratuitous `Neighbor Advertisement` message on behalf of each MNN. Not to say, for large mobile networks, processing all these individual registrations would impose an unnecessary burden to the HA. On the other hand, this approach would be quite advantageous since the process of sending and authenticating BUs would be left unchanged. BUs could also alternatively be piggybacked or sent alone.

BU processing by the MR on behalf of the MNNs It makes more sense if the node which is assigned the careof address also sends BUs. The mobile router would consequently send its careof address to all correspondent nodes of the mobile network, on behalf of MNNs, i.e a binding between MNN_{ip} and MR_{coa}. This alternative has the advantage of avoiding changes in the CN Operation. However, `Mobile IPv6` states that BUs can not be sent on behalf of an other node [Johnson and Perkins, 2000, section 10.8]. The CN would indeed be misled by the sender of the BU. Before the binding carried in the BU could be recorded in the Binding Cache, the MR has to be authenticated by the CN as the node that owns the address MNN_{ip} carried in the Home Address Option. This means that the MR must use the same security association as its MNN. An additional mechanism is required to exchange security associations between MNNs and the MR. Although the MR and its MNNs are likely to trust each other and to adopt the same administrative policy, it is not desirable to mislead the recipients, then we don't recommend this solution. We also see that MNNs the mobility management of the mobile network is not kept transparent to MNNs. Another drawback is that CN communicating with several MNNs from the same mobile network would redundantly record the same careof address for each MNN. Not only bandwidth consumption is wasted, but also memory allocation in the Binding Cache. A CN communicating with two MNNs from the same mobile network may receive a duplicate BU containing the primary MR_{coa}. This increases BU processing at the MR and at the CN and consumes unnecessary memory at the CN also wastes bandwidth. Another drawback is that CN communicating with several MNNs from the same mobile network would redundantly record the same careof address for each MNN. Not only bandwidth consumption is wasted, but also memory allocation in the Binding Cache. A CN communicating with two MNNs from the same mobile network may receive a duplicate BU containing the primary MR_{coa}. This increases BU processing at the MR and at the CN and consumes unnecessary memory at the CN also wastes bandwidth.

BU processing by the MR The more promising alternative is to establish a binding between the entire mobile network and the careof address obtained by the MR. However, nothing in the specification currently allows the CN to understand that all packets intended to a mobile network should be sent via the MR_{coa}. The BU only tells the CN that datagrams sent to the MR_{ip} should be sent via the MR_{coa}. Then, the format of the BU should be enhanced, and the operation of the MR be defined as a new entity in the `Mobile IPv6` specification. This can only be done provided the CN Operation is changed. This alternative has the benefit of hiding mobility of the network to the MNNs and frees them from any mobility management. Doing so while keeping MNNs out of any mobility management requires that the mobile router is

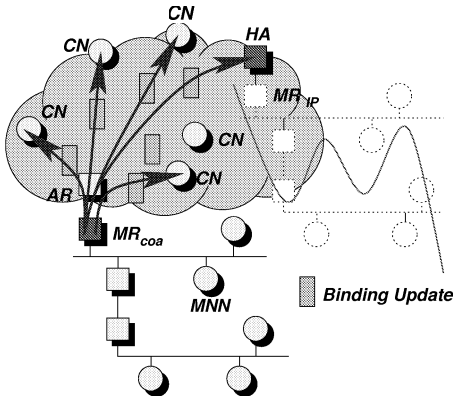

Figure 6.3: Binding Update Explosion

able to track the list of CNs communicating with MNNs and to send them BUs. This could be done by recording the source address of all encapsulated datagrams. The reception of an encapsulated datagrams is an indication that the sender does not have the primary MR_{coa} recorded in its binding cache. However, due to the possible high number of CNs of the mobile network, tracking and recording the identity of all CNs may add an important burden to the MR. Another drawback is that the MR takes the decision to advertise the MR_{coa} instead of the MNN which may prefer to keep its location secret. This could be negotiated between MR and MNNs, but requires new code in the specification. As for drawbacks, piggybacking can not be done by the MR without rewriting the `AH or ESP header` which may be present. This does not comply with `IPv6` recommendations. No headers but the `Routing extension header`, under some particular conditions, can be rewritten by routers along the path. Alternatively, encapsulation could be used, but this mechanism is more expensive.

6.2.4 Binding Update Explosion

The purpose of BUs is to offer optimal routing between the correspondent nodes and the mobile network nodes. Since optimal routing cannot be left aside for mobile networks, we note that periodic BUs would be sent individually to each CN, whatever the alternatives above discussed. Typically, 5 consecutive BUs should be sent to each CN every time MR obtains a new MR_{coa}, and periodically to refresh the Binding Caches, whether or not there is actual traffic between the CN and the MNN. As we have seen in section 5.1.2, the total number of CNs may be very large. In this case, the periodic emission of BUs would waste bandwidth resources and may cause congestion on the nearest links, principally on the wireless link and in the visited network. Links close to the MR would be periodically overloaded by a bulk of BUs. It may also waste processing power and overload capabilities of the MR. We term this Binding Update Explosion, and this is illustrated on fig.6.3. This issue is outlined by our simulations in section 9.2.1 and 9.2.2.

6.2.5 Other Issues

Routing Protocol Issues As a router, MR should process `Router Advertisements` sent on its home link. Is this advisable when MR is attached to a foreign link? Moreover, routing protocols use link-local addresses in their messaging. Link-local addresses are not handled by `Mobile IPv6` and are not tunneled between HA and MN.

Privacy In order to preserve privacy of their location, Mobile IPv6 offers MNs the ability to discriminate between the correspondent nodes that will be sent the primary careof address. If they wish so, datagrams from correspondent nodes to the mobile node transit through the HA and mobile nodes agree to suffer from higher transmission delays due to a non-optimal path. Similarly, it is advisable to offer mobile network nodes the ability to discriminate between the CNs that will be sent the MR_{coa}. Furthermore, it is also advisable that a VMN has this ability too.

6.3 Conclusion

In this chapter, we have studied the ability and shortcomings of Mobile IPv6, the proposed IETF standard for *host mobility support*, for supporting networks in motion. We have outlined that the intrinsic mechanisms offering optimal routing are not able to support mobile networks as efficiently as mobile nodes and would potentially cause an important waste of bandwidth resources and processing power. We have also outlined security concerns about the authentication of BUs and stressed the need for optimal routing and for mobility transparency to MNNs. Hence, a solution based on Mobile IPv6 to support mobile networks must provide optimal routing in a way that minimizes signaling, that limits memory and processing overhead, that complies with security requirements and that perform transparently from the MNNs.

Before proposing such changes to an existing protocol, it is usually required to balance the expected benefit over the deployment cost. As far as mobility support in IPv6 is concerned, there is no actual standardization since Mobile IPv6 is still a work in progress. This leaves us room to deliver new propositions which comply with IPv6 and not necessarily with Mobile IPv6.

Chapter 7

Proposed Mobile IPv6 Extensions

In this chapter, we propose *network mobility support* extensions to Mobile IPv6. Our main interest is to offer permanent and un-interrupted Internet connectivity and optimal routing to all mobile network nodes, while scaling to a large number of correspondent nodes and a large number of mobile networks.

We first propose Prefix Scope Binding Update as a straight forward extension of the Mobile IPv6 specification which addresses the two issues discussed in section 6.1, namely the redirection at the HA of packets intended to mobile network nodes, and optimal routing between correspondent nodes and mobile network nodes.

Although Mobile IPv6 is a straight forward and adequate solution for supporting mobile networks with few correspondent nodes, the mechanism used to distribute the careof address to correspondent nodes does not scale to a large number of correspondent nodes as individual BUs are sent to each correspondent node. We therefore propose to use multicast to deliver Prefix Scope Binding Updates, where correspondent nodes are the group members. We indeed propose two multicast techniques. The first one is more targeted to a large number of correspondent nodes and makes use of standard multicast protocols that build a multicast tree between the mobile router and the correspondent nodes. The second is more targeted to a smaller number of correspondent nodes or to distinct groups of correspondent nodes split into smaller groups and makes use of a new technique, List-Based Multicast, where the IP address of the destinations is actually recorded in the Prefix Scope Binding Update itself.

As a further proposition, we envision a few framework architectures that combine our multicast propositions with some of the frameworks described in section 4.2, like the *Hierarchical Framework* and the *Virtual Network Framework*.

7.1 Prefix Scope Binding Updates

Prefix Scope Binding Update is an original work that we have presented in the Mobile IP Working Group at the 47th IETF meeting (Pittsburgh, August 2000) for the first time. Details are in the third revision of our internet-draft [Ernst et al., 2000a]). It attempts to provide an immediate solution to an immediate need. Our proposal suggests Mobile IPv6 extensions in order to advertise the mobile network prefix to the home agent and to all correspondent nodes communicating with mobile network nodes.

Basically, Prefix Scope Binding Update is a special format of a BU which registers an entire mobile network in one step as opposed to Mobile IPv6's BUs which only register a sole end-node. This is performed by making use of most existing Mobile IPv6 features, with minor extensions to the specifica-

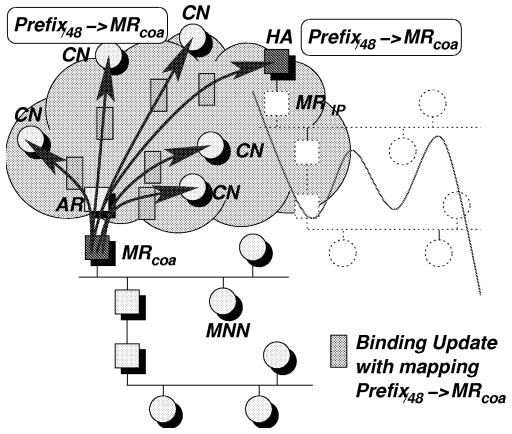

Figure 7.1: Prefix Scope Binding Updates

tion. A mobile router operates `Mobile IPv6` and obtains a careof address at each subsequent point of attachment. This careof address is sent to all correspondent nodes of the mobile network using Prefix Scope Binding Updates. Prefix Scope Binding Updates carry bindings that associate the mobile router's careof address with the mobile network prefix instead of the full 128-bit `IPv6` home address as defined in the existing `Mobile IPv6` specification[1]. As a result of the reception of the Prefix Scope Binding Update, the record in the Binding Cache can be seen as a *network-specific route* for the mobile network prefix.

Prefix Scope Binding Updates ensure optimal routing between CNs and MNNs. The concept of Prefix Scope Binding Update reduces the overhead of mobility management because the number of BUs is reduced by regarding the MNNs of a mobile network collectively. By aggregating the mobility management of an entire network in a single Prefix Scope Binding Update, individual registrations of mobile network nodes are avoided and mobile network nodes are preserved from the mobility management of their network. All the burden is entirely on the mobile router's side, the node that effectively changes its point of attachment. It limits bandwidth consumption between the mobile network and the home agent and saves memory and CPU. It also avoids a correspondent node to receive duplicate BUs and to register a duplicate entry in its Binding Cache. Prefix Scope Binding Updates scale to the size of the mobile network. *Authentication* of Prefix Scope Binding Updates is performed as for standard BUs: the mobile router is the sender of the Prefix Scope Binding Update and is authenticated as such. Therefrom, recipients of Prefix Scope Binding Updates are not misled by the identity of the sender since the mobile router does not send Prefix Scope Binding Updates on behalf of its mobile network nodes. However, this faces *authorization* concerns, as debated in section 7.1.3.3. Our proposal does also not help to scale to a large number of correspondent nodes and doesn't address the *Binding Update Explosion* issue.

7.1.1 Implementation

As detailed below, this proposition makes use of most of the existing `Mobile IPv6` features with minor extensions. It defines a new `Mobile IPv6` entity, the MR, which performs most of the existing MN Operation and it redefines the CN Operation. The format of the BU Option is slightly modified and we define a new `Destination Sub-Option` (Mobile Network Prefix Sub-Option).

The Prefix Scope Binding Update instructs CNs to add a binding between the mobile network prefix and the mobile router's careof address, i.e. a network route in their Binding Cache. Before sending

[1] Accustomed readers may have noted that the MR's prefix on the home link (the home prefix) is not the same as the one on the internal link (the mobile network prefix). In addition to the binding between mobile router's home address and mobile router's careof address, a binding between the mobile network prefix and the mobile router's careof address is needed.

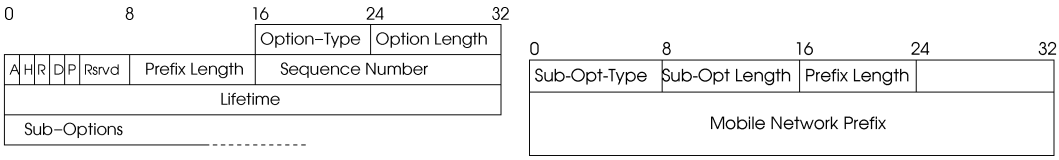

(a) Modified Binding Update Option Format

(b) Mobile Network Prefix Sub-Option Format

Figure 7.2: Prefix Scope Binding Update Format

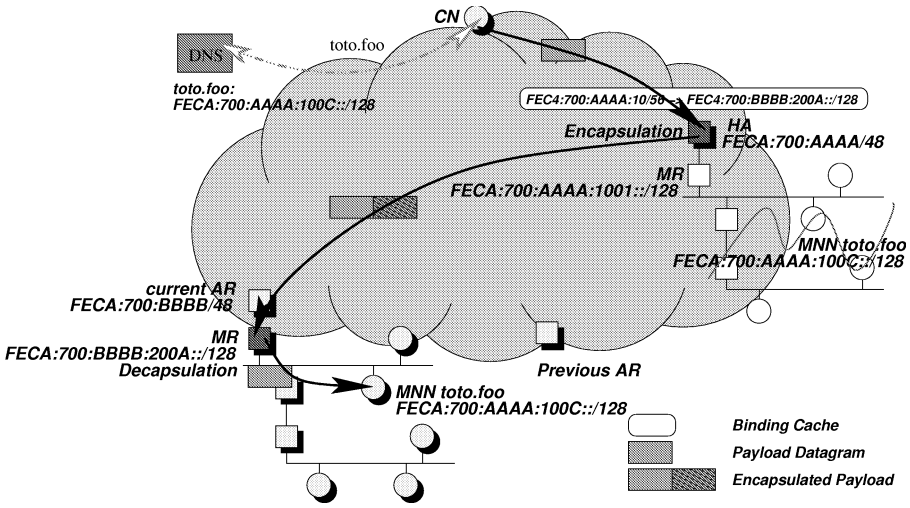

Figure 7.3: First datagrams transit through the home agent HA

a datagram, the CN checks if the prefix of the destination address matches the mobile network prefix recorded in the binding cache. If so, datagrams are sent via the MR's careof address using a `Routing Extension Header`. Fig. 7.1 illustrates the resulting record in the CN's Binding Cache.

Modified Binding Update Option format A new bit is taken from the reserved set of bits in order to indicate that the registration worth for the prefix as indicated in the Mobile Network Prefix Sub-Option, and not for a single 128 bits `IPv6` address. The format is described in figure 7.2 (a).

Mobile Network Prefix Sub-Option The Mobile Network Prefix Sub-Option is a new BU Option Sub-Option that we propose for carrying the mobile network prefix. It contains the mobile network prefix used as a netmask. This Sub-Option is inserted in the BU. The packet format is described in figure 7.2 (b).

The Mobile Router, a new Mobile IPv6 entity The MR is a new `Mobile IPv6` entity that embeds all the MN Operation with some extensions. Thus, the mobile router is a mobile node enhanced with the ability to send Prefix Scope Binding Updates to the list of correspondent nodes corresponding with the mobile network. Unlike the standard `Mobile IPv6`, the Prefix Scope Binding Update includes the Mobile Network Prefix Sub-Option. No changes are required to the standard MN Operation.

Extended Correspondent Node and Home Agent The CN Operation is extended to process the Mobile Network Prefix Sub-Option and to transmit via the MR's careof address all packets bearing a destination

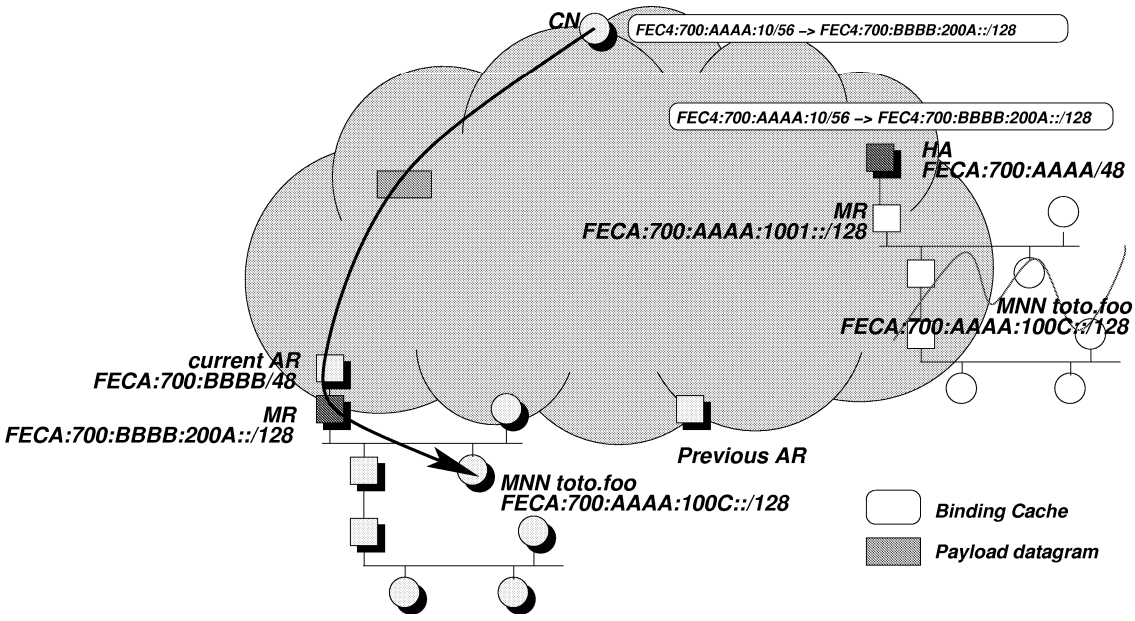

Figure 7.4: Optimal routing between CN and MNN

address that matches the mobile network prefix. The HA Operation is enhanced in the same way in order to process the Mobile Network Prefix Sub-Option. Our solution only requires that the HA is able to redirect packets intended to all nodes that have a common prefix to the MR's careof address, and not only to the MR.

7.1.2 Protocol Operation

Initialization In order to perform mobility management using Prefix Scope Binding Updates, the mobile network must be configured with a unique network prefix. This could be configured manually or automatically. The mobile network prefix could be of any length depending on the size and the organization of the mobile network. For instance, a large mobile network may be allocated a $/48$, whereas a single mobile IP-subnet may be allocated a $/64$.

MR Operation The MR obtains a new careof address MR_{coa} at each of its subsequent points of attachment using either stateless or stateful `DHCPv6` address autoconfiguration as a standard MN. Following this, it sends a Prefix Scope Binding Update to its HA and to all the CNs recorded in its Binding List. The Prefix Scope Binding Update contains a BU Option with the Mobile Network Prefix Sub-Option, a Home Address Option, and an *AH or ESP header* for authentication. The Prefix Scope Binding Update is sent at periodic time intervals to ensure that CNs do not delete the binding because its lifetime has expired. The Home Address Option is the same as for the existing `Mobile IPv6` specification. It contains the mobile router's home address used as an identifier. As a result of this registration, the MR will receive packets intended for the MNNs. Those packets are either tunneled from the HA in which case these are de-tunneled, or they are directly sent to the mobile router's careof address in which case the mobile router's MR_{coa} is switched with the MNN_{ip} contained in the `Routing Extension Header`. In both cases, the packet is forwarded to the relevant MNN. The MR monitors encapsulated packets and adds the source address of the CN in its Binding List. Following this, the MR sends a Prefix Scope Binding Update to the CN.

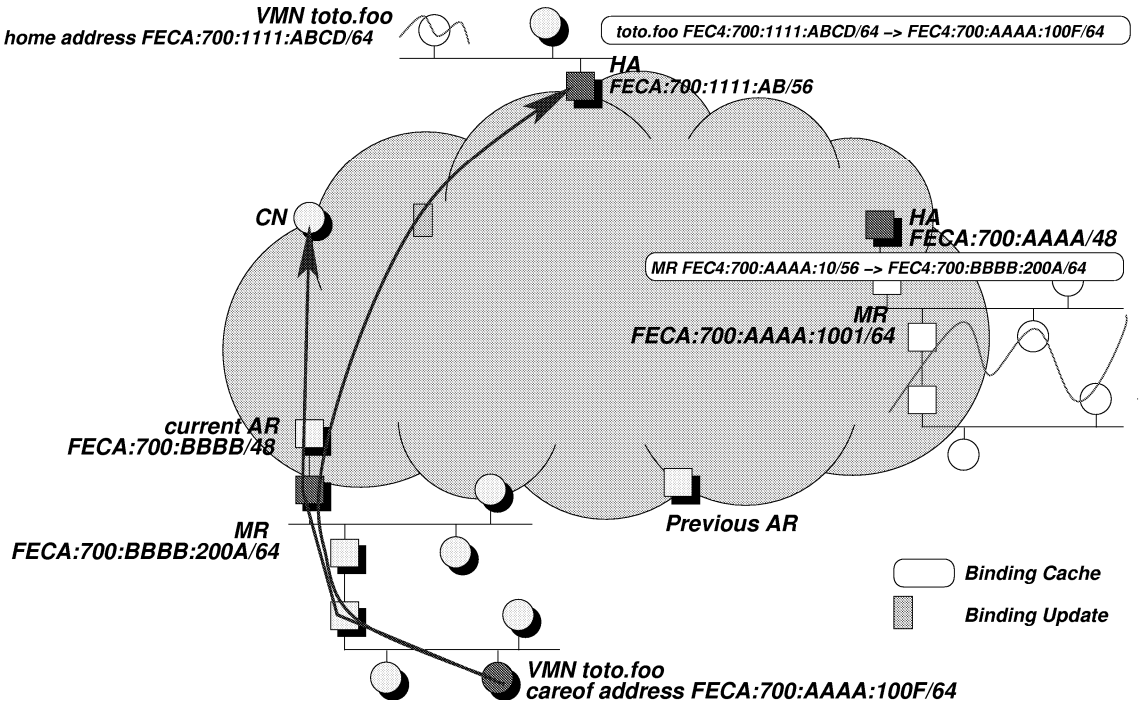

Figure 7.5: VMN Registration

CN Operation Prior to communication establishment with a MNN, the CN calls the `DNS` for the IP address corresponding to the domain name of that node. The `DNS` supplies the IP address MNN_{ip}. First packets are therefore sent to the MNN_{ip} and get routed up to the home link of the MR. Packets are intercepted by the home agent on the home link and tunneled to MR_{coa} (fig. 7.3). Upon reception of a valid Prefix Scope Binding Update, the CN authenticates the sender, registers the binding between the mobile network prefix and the MR_{coa} in its Binding Cache and sets the expiration timer. When it has a pending packet to send, the Binding Cache is searched. If the prefix of the pending packet's destination address matches the mobile network prefix, the mobile router's careof address MR_{coa} is returned. The pending packet therefore gets routed to the destination address MNN_{ip} via MR_{coa} using an `IPv6 Routing Extension Header` (fig. 7.4).

7.1.3 Discussion and Open Issues

7.1.3.1 Nested Mobility

In order to provide optimal routing to VMNs, two registrations are required: the first one is performed by the VMN by means of the existing `Mobile IPv6` and tells VMN's HA and CNs that the VMN is currently in the mobile network. A second is needed to tell them where is the mobile network.

A VMN that performs standard `Mobile IPv6` has a permanent home address VMN_{ip} from its home link and obtains a careof address VMN_{coa} in the mobile network. The prefix of this careof address is the mobile network prefix. The VMN registers its current VMN_{coa} with its own HA and its CNs, as shown in fig. 7.5. As a result of this first registration, packets sent by CN and intended to VMN are first routed to VMN_{coa}. Since the prefix of the MN_{coa} is the mobile network prefix, packet get routed to the MR's home agent. If the mobile network is not on the home link, packets are intercepted by the MR'home agent and encapsulated to MR_{coa}. MR decapsulates and forwards them to the VMN where the `Routing Extension Header` is processed. Any CN of the VMN, including its home agent, is regarded as a CN of the mobile network and is added in the MR's BU List by monitoring incoming packets

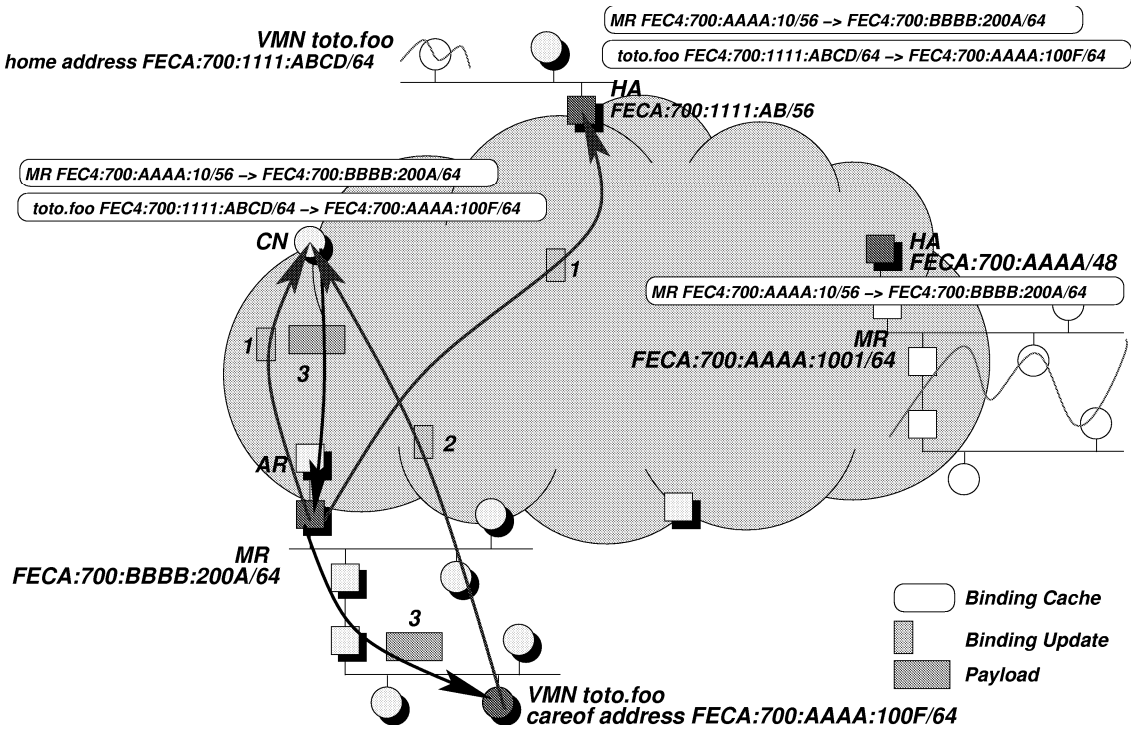

Figure 7.6: MR Registration for a VMN

intended to the VMN. Thus, Prefix Scope Binding Update could also be sent to VMN's CNs. In order to insure optimal routing, the Routing Extension Header must be filled with two addresses. This requires extensions in the CN Operation to record two addresses in the Binding Cache and to return both addresses when there is a pending packet for the VMN. Then, the Routing Extension Header must be filled in the right order so that packets are first routed the MR_{coa}, then to VMN_{coa}, and then delivered to the VMN_{ip}. With this solution, the mobility of the visited network is kept transparent to VMN which keeps sending BUs containing the VMN_{coa} to its own HA and CNs.

7.1.3.2 Piggybacking

In order to limit BU emission, `Mobile IPv6` provides a mechanism called *piggybacking* to include BUs directly in payload datagrams when it is actively communicating with it. In order not to conflict with the `IPv6` security requirements, this could be done by the MR by means of encapsulation.

7.1.3.3 Security Issues

The mobile router is clearly authenticated as the sender of Prefix Scope Binding Updates using the same mechanisms as for a standard mobile node. However, the most debatable issue is the authorization for the mobile router to register its careof address for a prefix (mobile network prefix) instead of its home address. There are currently hot discussions in the `IETF Mobile IP Working Group` mailing list about authorization. Without a means to check if the MR is authorized to send a Prefix Scope Binding Update for this prefix, there is a security hole that would allow a MR to pretend the ownership for a shorter prefix than the one it actually owns. This would allow a malicious MR to capture all traffic intended for a whole portion of the Internet. We are therefore paying much attention to this issue. We are currently working on it.

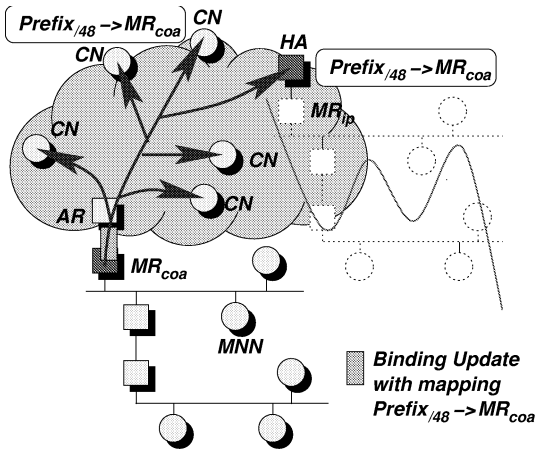

Figure 7.7: Multicast Delivery of Prefix Scope Binding Updates

7.2 Standard Multicast Delivery of Binding Updates

Prefix Scope Binding Update fits small mobile networks but is clearly inefficient in face of a large number of CNs since it does not prevent from the periodic *Binding Update Explosion*. When the careof address is sent by means of Prefix Scope Binding Update, we observe that each CN receives Prefix Scope Binding Updates carrying identical information: the MR's careof address and the mobile network prefix. Therefrom, for a large number of CNs, it may be wise to multicat Prefix Scope Binding Updates instead of sending individual copies.

We therefore propose to send Prefix Scope Binding Updates by means of a multicast routing protocol. Each mobile network is allocated a permanent multicast address that idenfies the BU Recipient Group, i.e. the group composed by all CNs. CNs join the BU Recipient Group using standard `IPv6` multicast group membership mechanisms. The MR sends periodic Prefix Scope Binding Updates to the multicast address and the Prefix Scope Binding Update is delivered to all subscribed CNs. As described in section 7.1, the Prefix Scope Binding Update instructs CNs to add an entry in their Binding Cache. Before sending a datagram, the CN checks if the prefix of the destination address matches the mobile network prefix recorded in the Binding Cache. If so, packets are sent via the MR's careof address using a `Routing Extension Header`.

The multicast delivery of Prefix Scope Binding Updates is best designed to reduce signalling load both in the entire network and on the wireless link specifically, and to scale to a large number of CNs, particularly when the emission rate of Prefix Scope Binding Updates is significant. This is demonstrated by our simulation results in section 9.3. Sending Prefix Scope Binding Updates to a multicast address alleviates the need for the MR to track individual CN since it is unnecessary to know the group members to send to a multicast group. This saves memory and processing power at the MR.

7.2.1 Implementation

Our proposition first assumes that an efficient and cheap *inter-domain multicast routing protocol* is available to mobile sources. Given this assumption, we make use of all the existing `Mobile IPv6` features and the Prefix Scope Binding Updates extensions described in section 7.1. Minor modifications are brought to the CN Operation and the MR Operation in order to process multicast packets. A new message is provided to inform CNs that they shall join the multicast group advertised in the message. No other specific extensions are required. As an optimization, which we are not going to detail, the multicast address may be recorded in the `DNS` and returned directly to the CNs by means of a new `DNS Resource Record`.

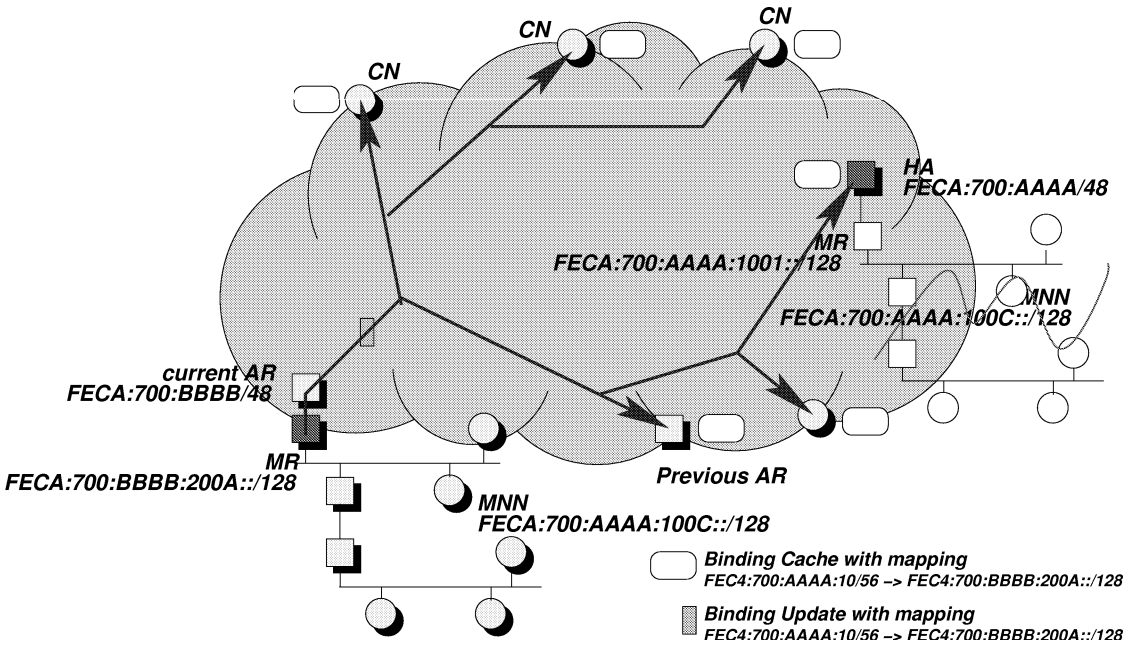

Figure 7.8: Binding Update Distribution from the MR

Extended Correspondent Node The CN Operation is extended for processing the Mobile Network Prefix Sub-Option as described in section 7.1. It is also extended to process a warning message containing the multicast address and sent by the MR. Alternatively, the multicast address could be returned by the DNS.

Extended Home Agent The HA Operation is extended for processing the Mobile Network Prefix Sub-Option and to redirect the packets with a destination address corresponding to the mobile network prefix as described in section 7.1.

Extended Mobile Router The MR is an IPv6 multicast-enabled router (with the ability to obtain and register a multicast address, and to send packets to a multicast group). The MR Operation is similar to the MR Operation highlighted in section 7.1. In addition, the MR Operation is extended with the ability to send the multicast address to the CNs.

Multicast Address The multicast address is constructed from the IPv6 address format described in section 2.2.2.2 (fig. 2.2 (b)). We propose to use the mobile network prefix as the suffix used in the multicast address that identifies the BU Recipient Group. A bit from the *flags* field could be taken to indicate that the multicast address worth for a mobile network.

7.2.2 Protocol Operation

Initialization Prior to its first movement, the mobile router forms a multicast address from the mobile network prefix. This multicast address identifies the BU Recipient Group for this mobile network. The multicast address is then registered with the multicast authority.

MR Operation MR obtains a new careof address on each subsequent visited link using either stateless or stateful DHCPv6 address autoconfiguration. Following this, it sends a Prefix

Scope Binding Update to its HA and to the multicast address identifying the BU Recipient Group. The BU contains a BU Option with the Mobile Network Prefix Sub-Option, a Home Address Option, plus the AH or ESP header. The BU is sent at periodic time intervals to ensure that CNs do not delete the binding because its lifetime has expired. The Home Address Option is the same as for the existing Mobile IPv6 specification. It contains the mobile router's home address used as an identifier. Fig. 7.8 shows Prefix Scope Binding Update sent to the multicast group of subscribed CNs. As a result of this registration, the MR receives packets intended for the MNNs. Those packets are either tunneled from the HA in which case these are de-tunneled, or they are directly sent to the mobile router's careof address in which case the mobile router's careof address is switched with the MNN_{ip} contained in the Routing Extension Header. In both cases, the packet is forwarded to the relevant MNN. Receiving a packet encapsulated by the HA is an indication that the original sender does not have the MR_{coa}. In this case, the MR shall warn the CN with a message containing the multicast address.

CN Operation: Prior to communication establishment with a node MNN, the CN calls the DNS for the IP address corresponding to the domain name of that node. First packets are therefore sent to the MNN, intercepted by the home agent on MR's home link, and tunneled to the MR's careof address (fig. 7.3). When appropriate, the CN joins the BU Recipient Group with the multicast address provided by the MR. This could be performed by means of Multicast Listener Discovery [Deering et al., 1999], the protocol used for routers to discover neighboring hosts interested in getting multicast datagrams. Following BUs sent by the MR to the multicast address are forwarded up to the CN. The CN authenticates the sender of the received BU, registers the binding between the mobile network prefix and the mobile router's careof address in its Binding Cache and sets the expiration timer. When it has a pending packet to send, the Binding Cache is searched. If the prefix of the pending packet's destination address matches the mobile network prefix, the mobile router's careof address is returned. The pending packet therefore gets routed to the MNN via the mobile router's careof address by using an IPv6 Routing Extension Header (fig. 7.4). The CN may leave the multicast group. hen communication is over.

7.2.3 Open Issues

7.2.3.1 Multicast Issues

The candidate multicast routing protocol must take into account the specific characteristics of dynamic multicast groups with a unique and mobile source. It should scale to a large number of mobile networks and sparsely distributed CNs. In addition, it should not be limited to intra-domain multicasting.

Inter-Domain Multicast Routing Protocol As highlighted in section 2.3.2.2, there is presently no standardized protocol to support inter-domain multicast.

Cost of Multicat The performance of our proposal depends on the underlying multicast routing protocol. Any multicast technique such as Shortest-Path Tree (SPT) or Core-Based Tree (CBT) (see section 2.3.2) may be used to build the multicast delivery tree leading to the CNs. However, the gain of multicasting Prefix Scope Binding Updates must be balanced against the multicast cost. This cost includes the computation and maintenance cost of the multicast tree, the group membership cost. It varies with the density of the group members, and the mobility frequency of the source. A protocol like DVMRP is clearly inapropriate for a mobile source because the tree must be recomputed upon every relocation of the source and packets are flooded to the entire network.

Sparse Group As observed in section 5.1.2, CNs are likely sparsely distributed in the Internet. In this situation, a Sparse-Mode multicast protocol seems more appropriate (see section 2.3.2). Dense-Mode

protocols are clearly inefficient since data packets are occasionally sent over many links that do not lead to receivers of the group [Deering et al., 1994]. If a SPT-based protocol rooted at the MR is used, the source of the multicast group is mobile and the delivery tree has to be updated upon every new point of attachment of the mobile network. If a CBT-based protocol is used, the position of the core and the delivery tree may not be optimal. This question will be studied when we present our simulation results (see section 9.3).

7.2.3.2 Security Issues

The distribution of a *Security Association* between a multicast source and the group members is necessary for authentication. This is a current hot topic and was first debated during San Diego 49th IETF in the MSEC BOF. This BOF came up with a bulk of work already developed at Internet Research Task Force (IRTF) in the SMUG Working Group and it is very likely that standards are available soon. We are therefore waiting for the output of those working groups.

7.2.3.3 Privacy Concerns

The MR has no control over the group members and is unable to discriminate between subscribed CNs since multicast BUs are sent to all or none group members. As a result, the MR can not hide its location to some of the CNs while not to the other. Location privacy of MNNs can therefore not be ensured to the MNNs which may wish so. To overcome this, we may further enhance our proposal to add, for instance, encryption of the multicast address, encryption of the BU, or limit subscription to the BU Recipient Group.

7.2.4 Related Work

Previous papers addressing *host mobility support* have already suggested the use of multicast as this was seen in chapter 4, but not for the delivery of BUs. For instance, in [Mysore and Bharghavan, 1997], mobile nodes are identified by a unique multicast address, which is location independent and invariant. More recently, [Helmy, 2000; Blazevic and Le Boudec, 1999] has also proposed the use multicast to deliver packets to the MN. In these proposals, the multicast infrastructure is always used to route payload packets down to the mobile node, not to deliver BUs to CNs, as we propose.

7.3 List-Based Multicast Delivery of Binding Updates

In this section, it is provided another multicast extension of Mobile IPv6 to deliver Prefix Scope Binding Updates. In contrast to the previous section which proposes the use of traditional multicast routing protocol, this section proposes a more innovative multicast technique that records the list of destinations in the packet header itself (see section 2.3.3).

We therefore propose to send Prefix Scope Binding Update to a list of CNs self-contained in the packet itself. CNs do not need to join in order to receive Prefix Scope Binding Updates. When the MR has a pending BU to send, it inserts all the list of CNs or a subset of the CNs recorded in its BUList in the Prefix Scope Binding Update.

This technique is more appropriate for a small number of CNs since the length of the packet grows with the number of destinations. In practice, Prefix Scope Binding Updates could also be sent to distinct sets of CNs, thus this solution is still adequate to address the *Binding Update Explosion* problem for a large number of CNs. The performance of this solution is evaluated in section 9.4.

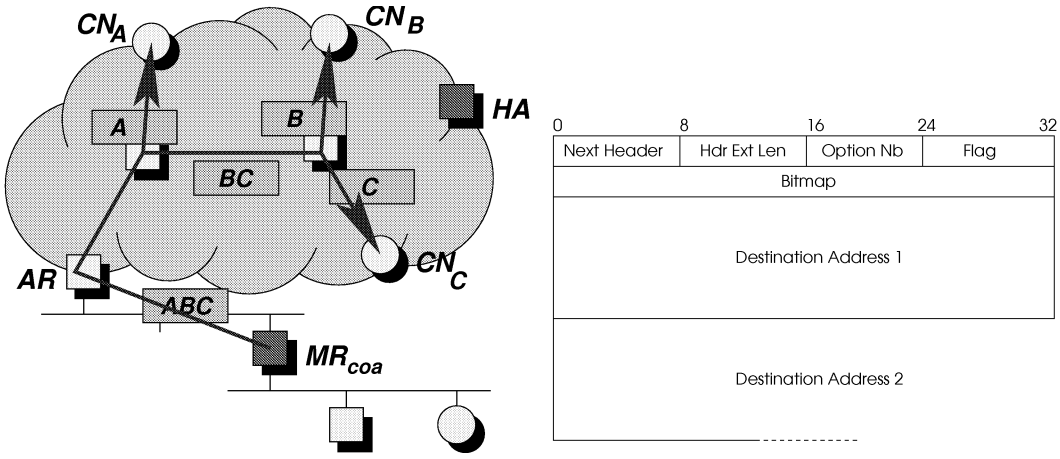

(a) Mobility Support using List Based Multicast

(b) LBM Hop-by-Hop Extension Header

Figure 7.9: List Based Multicast

As a first advantage, there is no need for a multicast routing protocol, nor for a group membership protocol. This greatly simplifies the question of *inter-domain multicast routing* since there is no difference whether the CNs are in the same routing domain or not. CNs shall be split into disctinct lists, thus the MR can discriminate between CNs: it may send Prefix Scope Binding Updates containing distinct careof addresses or distinct lifetime to distinct groups.

7.3.1 Implementation

As detailed below, our proposition makes use of most of the existing `Mobile IPv6` features, and requires the Prefix Scope Binding Update extensions detailed in section 7.1.2. Only other minor extensions are needed. The MR performs most of the existing Prefix Scope Binding Update MR Operation, and we slightly redefine the CN Operation. To carry the list of CNs, we define a new `Hop-By-Hop Options Extension Header`. This new header is added in the Prefix Scope Binding Update.

LBM Extension Header The new header, as illustrated on fig.7.9 is a `Hop-by-Hop Options Extension Header`. As such, it should be processed by all routers that understand the option. The *header length* is the total length of this header and varies according to the actual number of destinations. The *option number* is to distinguish this option from other already defined options and is assigned by a global Internet numbering authority. Each *address* field corresponds to a destination of this packet. Any number of addresses may be specified; this number is only limited for performance considerations. *Flags* are options that tell the `LBM`-enabled router how it should process this option. Currently, we have only defined a *remove flag*. If set, the router should remove the destination addresses not reachable from the interface where it will duplicate the Prefix Scope Binding Update. If not set, the list remains in the Prefix Scope Binding Update. In the former case, a *bitmap* field must be filled and updated by each router. It is 32-bits wide. Each rank corresponds to a destination recorded in the *address* field. When set, it indicates that the destination remains still undelivered. If more than 32 destinations are required, another *bitmap* field must be added. In both cases, the total number of destinations is derived from the *header length*.

LBM Address Format The address specified in the destination field of the `IPv6` header could alternatively be the address of one of the CN, or preferably a specific address format reserved for `List-Based Multicast`, or the address of a `LBM`-enabled router.

LBM Header Processing For each bit set to 1 in the bitmap, the List-Based Multicast-enabled router reads the corresponding address and then interrogates its routing table for ascertaining the next hop towards this address. If the next hop differs from the one specified in the destination address field in the IPv6 header, it duplicates the packet. This results in a number of duplicate packets. There are as many duplicate packets as there are distinct next hops towards the destinations from the given router. In order to prevent loops and avoid the processing of destinations to which a duplicate packet was already transmitted, the router sets the destination address of each duplicated packet to one of the undelivered destination and also sets the corresponding bits in the bitmap. A packet with all bits in the bitmap set to 0 means that no duplication is required anymore.

Extended MR The MR embeds the extensions described in section 7.1. In addition, the MR is able to include the list of CNs in a LBM Extension Header and to fill the corresponding fields. It is also provided a basic decision mechanism to decide how to send BUs: sending copies of a Prefix Scope Binding Update to individual CNs, sending a single Prefix Scope Binding Update embedding a LBM Extension Header, or sending several Prefix Scope Binding Updates to several set of CNs.

Extended CN The CN Operation must be extended to redirect the Prefix Scope Binding Update to other CNs still recorded in the LBM Extension Header.

7.3.2 Protocol Operation

MR Operation The MR, on receiving an encapsulated packet, includes the address of the CN in the Binding List. When it has a pending Prefix Scope Binding Update to send, it fills a LBM Extension Header with the list of CNs and send a Prefix Scope Binding Update message that also comprises the LBM Extension Header. It may also decide whether CNs should be removed by LBM-enabled routers when the packet is duplicated on distinct interfaces by setting the *remove flag*.

CN Operation On receiving a Prefix Scope Binding Update containing a LBM Extension Header, the CN checks if there is still destinations to which the packet remains undelivered. In such a case, the CN duplicates the packet and sends it to the next undelivered destination.

7.3.3 Open Issues

Security The IPv6 security policies prevent routers to rewrite the information contained in the header and extensions headers. In practice, this means that List-Based Multicast-enabled routers are not allowed to rewrite the destination field, nor to remove the list of destinations that are not on the same next hop than the one where the packet is transmitted. There are several means to address this issue which definitely deserves an in-depth study along as all other security aspects in general. A first means to overcome this is encapsulating the Prefix Scope Binding Update to the next undelivered CN or to the next List-Based Multicast-enabled router on the way to the CN, which adds a 40-bytes overhead. Concerning the list of destinations, it could be specified that this field be set to 0 when the security check is applied on the packet as this is done for the *hop limit* and the *flow label* field of the IPv6 header. There is no perceived risk that the list of destinations is forged by a malicious router, since this would only result in the non-delivery of the Prefix Scope Binding Update.

Processing Overhead As for the drawbacks, recording the list of CNs in the packet itself requires more processing of the packet at intermediate routers, and would result in overloading routers and delaying the BU. Hence, processing the list of CNs at each intermediate router is not feasible. A good compromise is to

process the option only at some well-located and dedicated routers. To avoid processing of the Hop-by-Hop Extension Header, the packet may be encapsulated between two LBM speakers.

Deployment In order to ensure the delivery of BUs to all CNs, at least all CNs listed in the packet must be LBM-enabled. If only CNs are able to process the option, the number of hops taken by the packet would depend on the order of CNs that get it. Hence, the addresses in the list may be ordered by the MR for better efficiency. For sure, the packet will not take an optimal path, thus it is desirable that at least some well located routers are also able to process this option. We envision to deploy the processing of this option at transit routers.

Privacy By recording the list of recipients in the packet we face some confidentiality and privacy threats. Every router that implements the option processing is able to determine the complete list of CNs for a mobile network.

Packet Overhead This technique is inevitably restricted to a bounded number of CNs since the more CNs in the packet, the larger the packet length. To overcome this, the list of addresses could be compacted to minimize the packet size. Although List-Based Multicast is only appropriate for a small number of CNs, it could also be used to send BUs to a large number of CNs if CNs are split in distinct independent groups.

Large Number of CNs In this case, the complete list of CNs may actually be split into several shorter lists. This ability is envisioned in the following instances: BUs with distinct lifetime, confidentiality with respect to the other CNs, BUs with distinct careof addresss, reducing packet overhead, etc.

7.4 Prospective Framework Architecture

In this section we highlight that both multicast techniques are complementary one to the other. This means that a good combination of the unicast and multicast delivery of BUs, together with other techniques as described in section 4.2, could offer a large panel of solutions for the problem posed by mobility in general, and for the *Binding Update Explosion* in particular. By combining the two techniques, we are able to scale LBM to a much larger number of correspondent nodes and to lessen the cost of the standard multicast technique.

The framework architectures presented in this section are purely prospective but show a direction where network mobility support could lead in future work in this area. In practice, we have envisioned the following combinations of mechanisms:

Overlay Network of LBM Speakers In order to ease the deployment of List-Based Multicast and to limit the processing of the LBM Extension Header, we envision a network of collaborative LBM-speakers. Prefix Scope Binding Update containing the LBM Extension Header would be encapsulated from the MR to a well known LBM-enabled router, possibly located in the closest BR, for instance. This router would duplicate the packet on the relevant interfaces and in turn forward it to the next LBM-speaker. Therefrom, individual copies of the BU could be sent encapsulated to the remaining CNs.

Combination of LBM and standard multicast In order to scale List-Based Multicast to a more important number of correspondent nodes, we envision another embodiment of this technique, combined with standard multicast. In each domain where there exists a number of correspondent nodes for the same

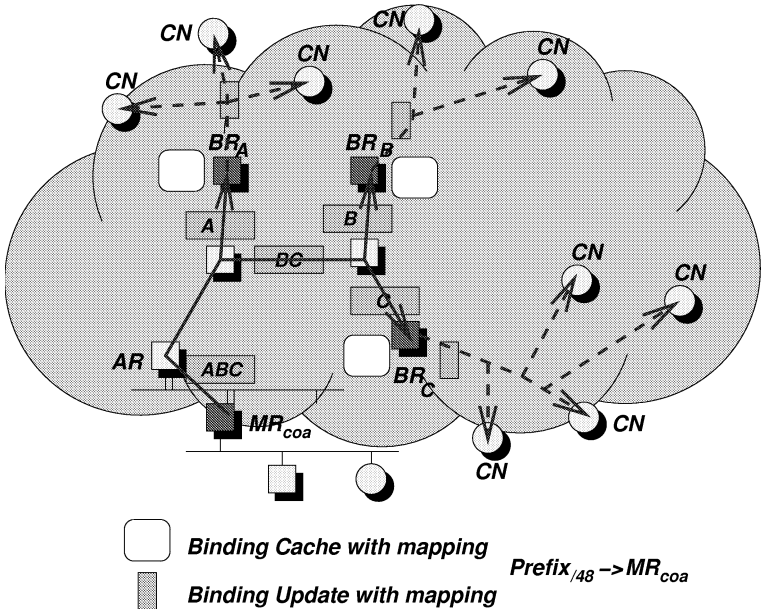

Figure 7.10: Combination of LBM with standard multicast

mobile network, correspondent nodes join a multicast group bounded to their domain. The multicast tree is rooted at the BR of the domain down to the correspondent nodes in the same domain and is built using an `intra-domain multicast protocol` like `PIM-SM`. The LBM Extension Header of the Prefix Scope Binding Update sent by the MR contains the multicast addresses of each intra-domain multicast group instead of the list of correspondent nodes. This packet is relayed by LBM-speakers well located in the backbone down to each domain, where it is forwarded on the intra-domain multicast tree.

Overlay Network of Mobility Servers Instead of sending Prefix Scope Binding Updates up to alll the correspondent nodes of the mobile network, we envision to send Prefix Scope Binding Updates to a number of Mobility Servers located near the correspondent nodes instead of the actual correspondent nodes. Mobility Servers serve as repositories. They are queried by correspondent nodes that enquire for the routing address of MNNs. As illustrated on fig.7.11, Prefix Scope Binding Updates are sent by the MR to a nearby Mobility Server. A multicast protocol delivers the Prefix Scope Binding Updates to the group of subscribed Mobility Servers.

In this embodiment, the establihmennt of the multicast tree is easier since the group of Mobility Servers is much less dynamic than the group of CNs. For better efficiency, the multicast tree could be shared by all mobility management signaling pertaining to mobile networks and therefrom be set permanently, by means of tunnels between adjacent Mobility Servers. A multicast tree shared by more sources of Prefix Scope Binding Updates insures a larger ratio of packet sent over the delivery tree compared to the multicast tree maintenance cost. The tree maintenance cost then becomes negligible compared to the amount of data transmitted on the tree.

7.5 Conclusion

In this chapter, we have proposed `Mobile IPv6` extensions to support mobile networks. All our enhancements allow to keep mobile network nodes away from any mobility management pertaining to the displacement of the mobile router. The first proposition is a straightforward solution that answers to an immediate need. It establishes a binding between the mobile network prefix and the current careof address of the mobile router. This binding is sent to correspondent nodes by means of Prefix Scope Binding

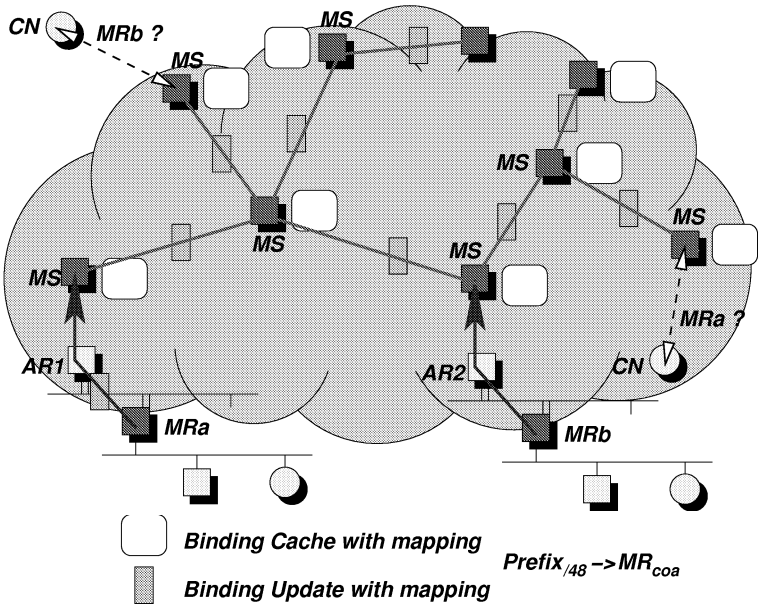

Figure 7.11: Overlay Network of Mobility Servers

Update. In order to reduce the signalling load incurred by the distribution of the careof address to the multiplicity of correspondent nodes we propose two mechanisms based on multicast. The common distinguishing feature of these mechanisms is the multicast technique used to build the multicast distribution tree. The first proposition is a generic multicast architecture that can be used to support mobility effectively. It uses traditional multicast. We assume that scalable and cheap multicast capabilities are deployed in the entire Internet in order to make use of them to provide ubiquitous mobility of the mobile network, though the state-of-the-art of inter-domain multicasting is still inadequate for supporting multicast group with members sparsely distributed in the Internet. A detailed discussion of the candidate multicast protocol is indeed beyond the scope of this dissertation. The second proposition is to record the list of correspondent nodes in the BU itself. This is a new multicast technique which is best envisioned for small groups since the length of the packets is growing with the number of addresses.

Part III

Performance Evaluation

Chapter 8

Simulation Process

In this chapter, we describe the method we use to evaluate the performance of our propositions, by means of simulation. The first section describes our needs, the second section describes our simulation tool and the extensions we brought to it. Then, we detail our simulation scenarios to configure the network topology and to determine the mobility pattern of the MN. The last section defines the metrics we are using for the analysis of the simulation results.

8.1 Needs

The purpose of the simulation is to evaluate the performance of Mobile IPv6 and to compare the multicast delivery of BUs over the standard unicast delivery. For doing so, we don't need to simulate an actual mobile network, a single mobile node, playing the role of a MR, suffices for evaluating the amount of signaling. This mobile node must operate Mobile IPv6 and our multicast extensions. We therefore need models simulating Mobile IPv6, and also List-Based Multicast, a Core-Based Tree multicast protocol and a Shortest-Path Tree multicast protocol. In order to obtain meaningful results, our simulation must involve a sufficiently large number of CNs. In addition, we need a low ratio of CNs over the total number of nodes in the topology. This means we need to manipulate large topologies, at least hundreds of nodes. These topologies must be good representations of the Internet topology. As it would be too cumbersome to determine by hand a good representation of the Internet, a topology model is needed. We also need a mobility model that preferably exhibits Wide-Area Mobility, i.e. displacements between topologically distant ARs (cf section 1.4).

8.2 Simulation Tool

For our simulations, we make use of NS-2 (version NS-2.1b6) [Fall and Varadhan, 2000], a discrete time event simulator developed by the VINT project, a join effort between the University of California at Berkeley, USC/ISI, Xerox PARC, LBNL. In order to conduct our study, we had to bring an important number of extensions. The following sections describe our simulation platform and the enhancements relevant to the results presented in this PhD dissertation. More details about our NS-2 enhancements can be found in our technical report [Ernst, 2001b].

NS-2 modules simulating Mobile IPv4, wireless communications (MAC 802.11, and ad-hoc routing protocols have already been contributed by SUN and Carnegie Mellon University (CMU). They were first developed for NS-2.1b2 and were then incorporated in the NS-2.1b6 distribution. How-

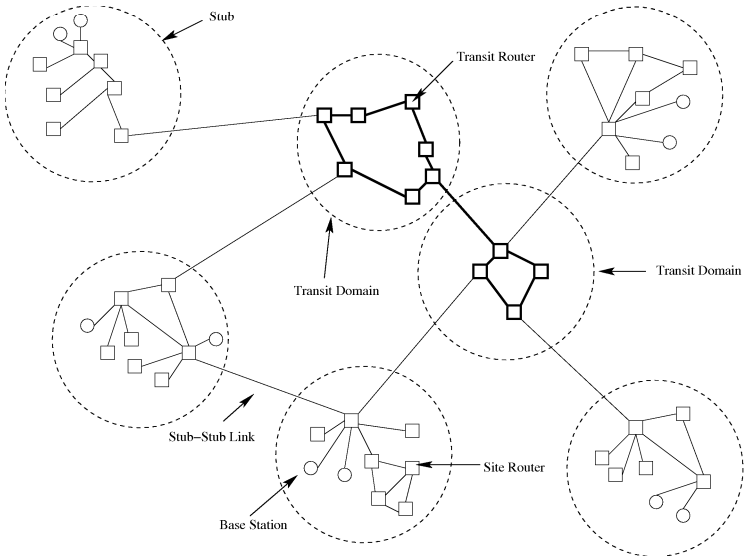

Figure 8.1: Internet Topology Generated by GT-ITM (Transit-Stub Model)

ever, those modules do not fulfill our needs and particularly not the movement of nodes in a `Wide-Area Network`. First, the actual models simulate movement of nodes within a bounded geographical grid (i.e. geographical movement) whereas we are more concerned by topological movements. Second, `Mobile IPv6` is not supported and multicast protocols do not work in coordination with the mobility modules. And last, there is no easy means to configure and manipulate large topologies.

8.2.1 Network Topology Model

Our network topologies are created with `GT-ITM` [Calvert and Zeggura, 1996; Calvert et al., 1997], a topology generation tool which uses the *Transit-Stub model* [Zegura et al., 1996] and aims to produce graphs that reflect the locality and hierarchy present in the Internet. The word *stub* basically corresponds to what we call a *site*. The model produces topologies according to parameters that indicate the number of domains, the number of sites connected to each router in the domain, and an average number of nodes in each site, and probabilities that a node is connected to the others in the same level. A probability indicating that two sites are directly connected can also be specified.

In addition to topologies produced by `GT-ITM`, we add a number of additional nodes (access routers, or ARs) where MNs are allowed to attach. An AR is an actual `NS-2` implementation of a base station, with `MAC 802.11` wireless capabilities. ARs are attached to a site router in the specified site. An instance of a topologies generated by `GT-ITM` with the addition of ARs is illustrated in fig. 8.1.

The instance topology shown on fig.8.2 is a transit-stub graph of 1050 routers, with 10 transit-domain (backbones) and 100 stub-domains (sites) comprising an average of 10 nodes. This topology is generated using the following `GT-ITM` parameters: `ts 1 21 ; 2 0 0 ; 10 20 3 0.5 1.0 ; 5 20 3 0.5 1.0 ; 10 10 3 0.7 0.8`. This topology is not very visible due to the large number of nodes, but anyway it gives a rough idea of its shape.

8.2.2 Manipulation of Large Topologies

A tool is needed to translate `GT-ITM` format into a format comprehensible by `NS-2`. There exists one, but the position of the node in the topology hierarchy is lost during the process which does not facilitate

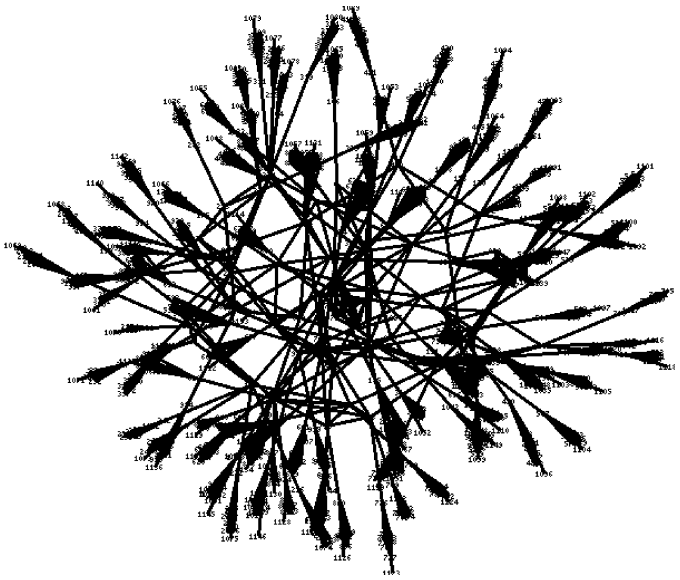

Figure 8.2: Instance of a Topology used for our Simulations

the simulation configuration of large topologies. We have therefore created a new tool that translates the GT-ITM output into a NS-2 format. The tool classifies nodes according to their position and function in the hierarchy: transit router, border router, site router, access router, mobile node. The tool also includes a library of procedures that allows automatic configuration of the topology according to the desired simulation scenario. For instance, the library is used to render all border routers LBM-enabled in one simulation, and to determine a number of access routers that will be visited by the mobile node.

8.2.3 Wide-Area Mobility Extensions

There is no actual means in NS-2 to simulate displacements in the topology other than geographical mobility between ARs with a 2-dimensional geographical co-ordinates and embedded wireless capabilities similar to a base station. Nodes with wireless capabilities have 2-dimensional geographical co-ordinates within the grid and communicate one with another by means of a radio channel, with a propagation model that determines if two nodes are able to listen to one another given their co-ordinates. The NS-2.1b6 distribution includes a number of geographical mobility models that have been designed to evaluate the performance of various ad-hoc routing protocols. They simulate geographical displacements within a bounded geographical grid. Indeed, the NS-2 micro-mobility models generate (frequent) geographical displacements (*micro-movements*). They are not suitable to simulate Wide-Area Mobility: displacements are too frequent and it would be too cumbersome to determine the geographical co-ordinates of all ARs with respect to both their location in the topology hierarchy and to their neighbors.

We therefore extended NS-2 in order to perform *macro-movements* between topologically distant ARs, i.e. ARs that belong to distinct sites, while *micro-movements* between ARs that belong to the same site is still possible using the existing geographical mobility models. A macro-movement is simulated by moving the MN from one site to another. In our simulations, we therefore define Local-Area Mobility as mobility within a single site and Wide-Area Mobility as mobility between two sites. We associate each site with a distinct radio channel and the same geographical grid. All nodes in a site with wireless capabilities have 2-dimensional geographical co-ordinates within the grid and communicate by means of this channel. Only nodes in the same site are able to listen to the same radio channel, then two nodes in distinct sites with the same 2-dimensional co-ordinates are unable to listen to one another. The MN begins the simulation in a site. The wireless interface of the MN is attached to the channel of that site. At a given time, the wireless interface of the MN is detached from the channel in the previous site and is attached to the channel in the

	Topology 1	Topology 2
Number of Sites	100	100
Number of ARs per site	1	1
Number of Wired Node	450	1050
Number of Mobile Nodes	1	1
Total Number of Nodes	551	1151
Total Number of Wired Links	1064	

Table 8.1: Simulation Topologies

new site. Then, the MN simply starts listening to Routing Advertisements sent by ARs in the new site.

8.2.4 Mobile IPv6 Extensions

The NS-2 distribution does not provide a model to simulate Mobile IPv6. Only the Mobile IPv4 protocol (without Routing Optimization) is included. We have founded easier to rewrite the code from scratch for a number of implementation and evolution reasons. We have implemented Mobile IPv6 as a set of NS-2 Agents that inherit from a base class MIPv6Agent. Following this, we have extended this code with the ability to send BUs to a group by means of standard multicast protocols and by means of List-Based Multicast. As for IPv6, we also add the Routing Extension Header processing to the structure of the NS-2 node

8.2.5 List Based Multicast and Multicast

Support for standard multicast already exists in the NS-2 distribution. Some extensions were nevertheless necessary to run simulations that simultaneously need multicast capabilities and mobility capabilities, particularly to allow a NS-2 Class MobileNode to send to a multicast group. As for List-Based Multicast, we have added a new packet format, and we have also modified the structure of all nodes with the ability to duplicate packets based on the list of destinations recorded in a header of the packet.

8.3 Simulation Scenario

8.3.1 Simulation Configuration

Topology Model The topologies are generated with GT-ITM. The parameters are illustrated in table 8.1. A single access router is added in each site and attached to a random router in the site. We only have one MN.

Communication Model A sample of n CNs is chosen randomly among all the site routers in the topology, by means of a random seed number generator (*CN seed*). At start up, the MN has only one CN in its BU List. A new CN arrives with an interval of time $cn\_delta\_arrival\_time$. When effective communication between CNs and the MN is needed, we generate a basic *Constant Bit Rate (CBR)* traffic between the CN and the MN, with a given payload size (50, 500, 1000). However, CNs do not actually need to send traffic toward the MN in order to receive a BU. They are inserted statically in the BU List of the MN.

Parameters	Set 1	Set 2
Number of Visited Sites	2	8
Inter-Site Handoff Rate	0.01	0.02
Period of time in each site	$\frac{1}{0.01} = 100sec$	$\frac{1}{0.02} = 50sec$
Maximum number of CNs	100	32
CN Delta Arrival Time	100	$8 \times 50 = 400$
Stop Time	10000	12800
Payload Packet Length		50 | 500 | 1000 bytes
Traffic (CBR)		0.02 packet per sec
Number of CN Samples	1	5
Number of Site Samples	1	5
Total Number of Simulations	1	5×5

Table 8.2: Simulation Parameters

The random and uniform selection of CNs complies with the observation we made in section 5.1.2 under *sparseness of the CNs*.

Mobility Model As we are more concerned about Wide-Area Mobility, we need a model to simulate topologically distant displacements, i.e. *macro-movements*. Our mobility scenario exhibits Wide-Area Mobility for a growing number of CNs at various positions of the MN. The *Inter-Site Handoff Rate* determines the movement frequency of the MN. A sample of s visited sites is randomly selected and uniformly distributed between the total number of sites, by means of a random seed number generator (*site seed)*. The MN begins the simulation in the first site. The MN then moves from one site to another site according to the macro-movement scenario characterized by the number s of sites and a fixed period of time $mt = \frac{1}{inter-sitehandoffrate}$ spent in each site. After the period of time mt, the MNs moves to the second site. This is repeated for the s selected sites. When all s sites have been visited, the MN moves again to the first site, then to the second site, etc. Once it has visited all s sites, a new CN is added in the BU List and the MN visits the same selection of sites again, in the same order and for the same period of time. The sequence of visited sites is repeated for the duration of the simulation, i.e. for each number of CNs. The interval of time between the arrival of two CNs is therefore given by $cn\_delta\_arrival\_time = mt \times s$ and the simulation lasts for $mt \times s \times n$ seconds

Although this mobility scenario does not appear to be realistic at first sight, the repetition of the same sequence of visited sites allows us to draw our results for a growing number of CNs. We only made this choice for convenience in order to limit the number of simulations. Concerning the random sequence of sites itself, Wide-Area Mobility between two `ISPs` that cover the same geographical area does not preclude that the two `ISPs` are topologically close in the Internet hierarchy. The remoteness of two sites in the Internet topology is independent from the geographical remoteness of those two sites. We therefore argue that a random selection of sites where the MN moves to diametrically opposed sites makes some sense.

8.3.2 Mobile IPv6 Settings

The MN listens to Router Advertisements sent by ARs. Following the reception of a Router Advertisement issued from a AR not recorded in its AR List, details of this AR is recorded for a period of time that does not exceed the lifetime of the Router Advertisement and the MN configures a new careof address with the prefix advertised by the AR. After obtaining a new careof address, the `Mobile IPv6` specification recommends to send a sequence of five BUs with an interval of one second in order to prevent the risk of loss. Then, periodic BUs must be sent to refresh the corresponding Binding Cache entries at the CNs

before their expiration. A 10-seconds interval was chosen for our simulations, but a less frequent and more optimal value could easily be chosen in order to limit bandwidth consumption. The HA is one of the ARs and is located at the same AR for all our simulations with the same topology.

8.3.3 Number of Simulations

The seed numbers allow us to run several simulations alternatively with the same CNs for distinct samples of sites, or the same visited sites for distinct samples of CNs. The simulation *Topology 1* and *Set 2* was re-run a total of 25 times for each protocol: 5 times with distinct samples of sites, each with 5 distinct samples of CNs. 9 "protocols" were simulated: standard `Mobile IPv6`, 4 distinct `List-Based Multicast` (all routers `LBM`-enabled, CNs only, CNs and border routers, CNs and transit nodes), 4 distinct standard multicast (`SPT`, `CBT` with the core at the HA, `CBT` with the core at the border router in the home site, and `CBT` with the core in the backbone network). All simulation parameters are summarized in tab.8.2.

8.3.4 Results Exploitation

Our measurements are made on a per packet basis. We aggregate the measurements for different periods of time: results are aggregated every mt and every $CN\_delta\_arrival\_time$.

8.4 Performance Metrics

In this section, we define the metrics we use to evaluate the performance of `Mobile IPv6` and our multicast extensions. Some of our metrics are derived from the usual terminology found in ad-hoc networks [Hu and Johnson, 2000] and multicast routing [Wei and Estrin, 1995] performance papers. The use of a multicast terminology is justified for two reasons: first by the analogy between a multicast group and a group of correspondent nodes and second by the comparison of the unicast and multicast distribution of BUs.

First let $G(N, L)$ be the graph representing the topology defined by N the total number of nodes and L the total number of links within the network. Then, let $g(n, l, s, t)$ be the set of correspondent nodes for a given mobile node at position s and time t and defined by n the set of correspondent nodes and l the total number of links on the path from s to all members of g.

Distance The number of hops between a source and a given destination at position d: $dist(s, d)$

We distinguish between:

- *Optimal Path*: The distance of the shortest route: $dist_{opath}$
- *Effective Path*: The distance of the path where a packet was effectively routed: $dist_{epath}$

Mean Distance The average distance from the MN at position s to any member of g is given by:

$$\overline{dist(s, g)} = \frac{\sum_{i=1}^{n} dist(s, i)}{n}$$

The mean distance from any position of the MN to any member of g is given by:

$$\overline{dist(g)} = \frac{\sum_{i=1}^{m}\sum_{j=1}^{n} dist(i,j)}{m \times n}$$

on-tree links Number of links used to send an instance of a BU to all the CNs from a position s over the total number of links in the network. For a multicast delivery of BUs, *on-tree* links is the number of links used by a given BU sent from the MN at position s to all members of g. For a unicast delivery of BUs, we define *on-tree* links as the total number of links used to send a single copy of a BU from the MN at position s to each member of g.

$$on - tree_{unicast}(s,g) = \sum_{i=0}^{n} dist_{epath}(s,i)$$

Tree Size The average distance between the MN at position s to a correspondent node of the set g is given by:

$$tree\_size(s,g) = \frac{\sum_{i=1}^{n} dist_{epath}(s,i)}{n}$$

Path Optimality

$$Optimality(s,g) = \overline{dist_{epath(s,g)}} - \overline{dist_{opath}(s,g)}$$

Density We call *density of g with respect to s* the following:

$$Density(s,g) = \frac{l}{L} = \frac{on - tree(s,g)}{L}$$

Mean Density The average density of the group g taking into account all positions of the mobile node is defined as:

$$Density(g) = \frac{\sum_{i=1}^{m} l_i}{m \times L} = \frac{\sum_{i=1}^{m} on - tree(i,g)}{m \times L}$$

where m is the number of positions of the mobile node and l_i the number of links from position $i (1 \leq i \leq m)$.

Multicast Overhead For the group g, let $M(t,p,l,g)$ be the total packet length of multicast control messages sent on link l from time $t-p$ to time t. Multicast control messages includes tree maintenance messages and group membership management messages. The total bandwidth consumed by multicast control messages in the entire network during interval p at time time t is given by:

$$Cost_{multicast}(t,p,g) = \sum_{l=1}^{L} M(t,p,l,g)$$

Mobility Management Overhead For the group g, let $N(t,p,l,g)$ be the total packet length of mobility control messages sent on link l from time $t-p$ to time t. Mobility control messages include all BUs, Binding Acknowlegments, BRs. The total bandwidth consumed in the network by mobility control messages during interval p at time time t is given by:

$$Cost_{mobility}(t,p,g) = \sum_{l=1}^{L} N(t,p,l,g)$$

Transmission Overhead For the group g, let $T(t,p,l,g)$ be the total packet length of data packets sent on link l from time $t-p$ to time t. The total bandwidth consumed in the network by data packets during interval p at time time t is given by:

$$Cost_{transmission}(t,p,g) = \sum_{l=1}^{L} T(t,p,l,g)$$

Optimal Overhead For the group g, let $O(t,p,l,g)$ be the total packet length of payload packets with no `IP` overhead[1] sent on link l from time $t-p$ to time t when packet is transmitted over the most *optimal path*. The total bandwidth consumed in the network by payload packets transmitted over the most optimal path during interval p at time time t is given by:

$$Cost_{optimal}(t,p,g) = \sum_{l=1}^{L} O(t,p,l,g)$$

Routing Overhead The routing overhead is the difference between the transmission cost of packets over the *effective path* and the optimal cost of packets over the optimal path. It is given by:

$$Cost_{routing}(t,p,g) = Cost_{transmission} - Cost_{optimal}$$

Control Overhead

$$Cost_{control}(t,p,g) = Cost_{mobility} + Cost_{multicast}$$

Total Overhead

$$Cost_{total} = Cost_{control} + Cost_{routing}$$

Bandwidth For each of the above costs, the bandwidth consumption is expressed in a number of bits per second during an interval p and is given by:

$$Bandwidth_{cost}(t,p,g) = \frac{nbbytes \times 8}{time} = \frac{Cost(t,p,g) \times period}{time}$$

[1]The first 40-bytes `IPv6` header is not considered overhead. Any additional header (`Extension Headers` and `Encapsulation` is considered overhead

End-To-End Delay The delay, expressed in a number of seconds, from when a successfully delivered packet is sent by a source node at position s until it is received by the corresponding destination node d: $delay(s, d)$

Mean End-To-End Delay The average transmission delay from the MN at position s to any member of g:

$$\overline{delay(s, g)} = \frac{\sum_{i=1}^{n} delay(s, i)}{n}$$

Chapter 9

Simulation Analysis

All our simulations are grouped in this chapter for convenience and are based on the scenario described in section 8.3 whereas metrics are those described in section 8.4. We focus on Wide-Area Mobility. The MN does not perform Local-Area Mobility displacements (no *micro-movements* within the site). This means that all our simulation results equally apply to both a MN that would operate `Hierarchical Mobile IPv6` (no BUs are sent as a result of a *micro-movement*) or to a MN that would operate `Mobile IPv6` without changing its point of attachment in the visited site (the MN does not perform any *micro-movements*).

Each of the following sections describes and analyzes a distinct experiment. The first experiment highlights the impact of mobility in a `WAN` according to our mobility pattern. The second is a performance analysis of `Mobile IPv6` on the wireless link with an important number of CNs, and in the wired Internet. Then, we evaluate the standard multicast delivery of the BUs against unicast for a `Core-Based Tree (CBT)` with several positions of the core and for a `Shortest-Path Tree (SPT)`. The last experiment evaluates the performance of our `List-Based Multicast` technique when all or only some of the routers do actually implement this technique. And finally, we conclude the analysis of all our results.

9.1 Mobility Pattern Analysis

This experiment shows the impact of Wide-Area Mobility on the *distance* from a MN to its CNs and HA. Our results show that the mean distance from the MN to all the CNs is independent of the topological location of the MN. The simulations were conducted according to scenario *Topology 1, Set 2*.

Distance to CNs The purpose of fig.9.1 (a) is to highlight our mobility scenario. It shows the *distance* between the MN and two distinct CNs. As expected, the distance to a given CN varies quite a lot according to the visited site. At time 0, the second CN is only one hop away from the MN because it is located in the first visited site.

Distance from the HA Fig.9.1 (b) shows the *distance* between the MN and the HA for two different samples of sites. This distance varies less significantly than the distance to the CNs because the HA is located quite close to the BR.

Mean Distance to CNs The first curve on fig.9.2 shows the *mean distance* $dist(s, g)$ from 8 distinct sites visited by the MN to the group of CNs for a growing number of CNs. The first curves highlight the

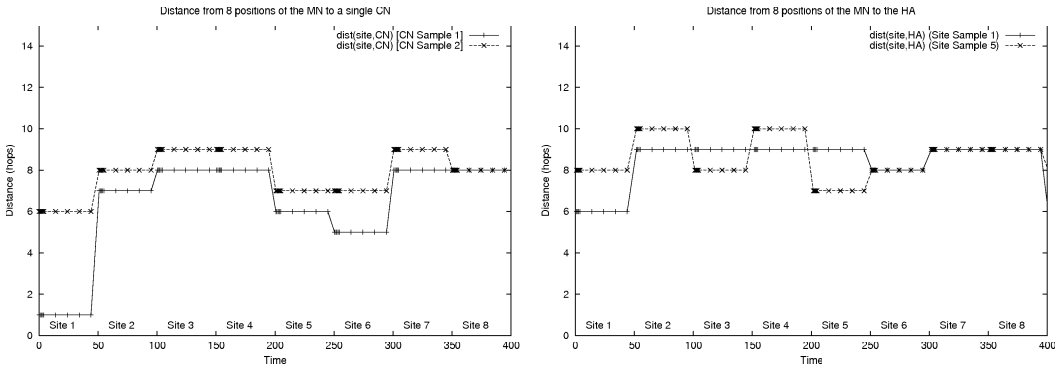

(a) Distance from MN to a single CN

(b) Distance from MN to the HA

Figure 9.1: Distance from the MN to a single node

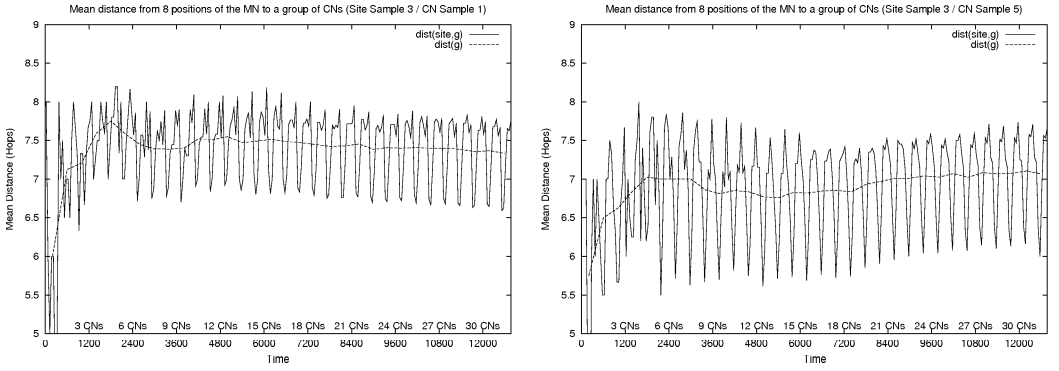

(a) First sample of CNs

(b) Second sample of CNs

Figure 9.2: Mean Distance from the MN to a growing number of CNs

oscillation of the *mean distance* in function of the visited site. Oscillations on all our simulations are easily explained by our mobility scenario. The *mean distance* $dist(s, g)$ from the MN at position s to a given number of CNs g varies according to the site s visited by the MN. The second curve is the mean distance $dist(g)$ averaged over the sample of visited sites and is therefore independent on the position of the MN. As we observe, the *mean distance* tends to a value between 7 and 8 hops on this topology. As the number of CNs grows, the *mean distance* tends to diverge less and less from this value. This behavior is observed for all samples of CNs. Fig 9.3 (a) shows this for three samples of CNs. Finally, in fig.9.3 (b), we have averaged all samples of CNs for the first sample of sites, and all the samples of sites for the first sample of CN. The third curve, which nearly looks like a straight line, is $dist(g)$ averaged over all the 25 simulations.

Conclusion We conclude that for a large number of CNs uniformly distributed in the topology, the *mean distance* $dist(g)$ from the MN to the CNs does not depend on the position of the MN in the topology. This observation is particularly interesting for the kind of mobile networks that we are willing to support, since, as noted in section 5.1.2, CNs of a train or an aircraft are likely sparsely distributed in the Internet topology.

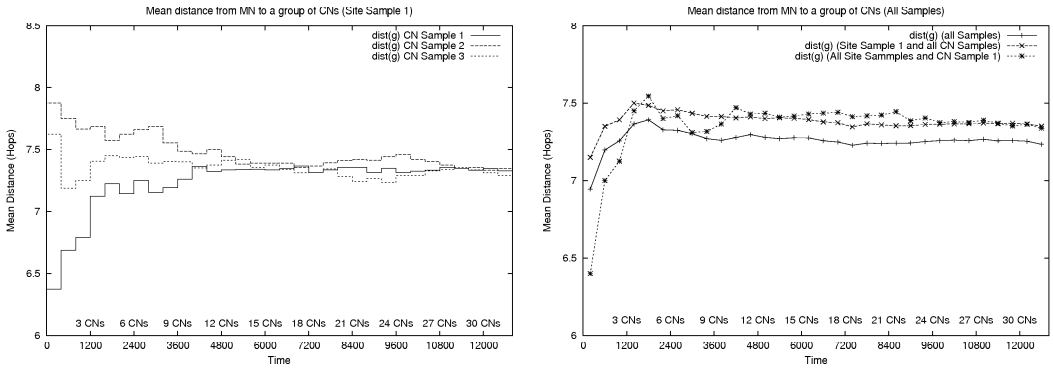

(a) Average of $dist(g)$ for all Sites Samples

(b) Average of $dist(g)$ for all CN and Sites Samples

Figure 9.3: Mean Distance from MN to CNs

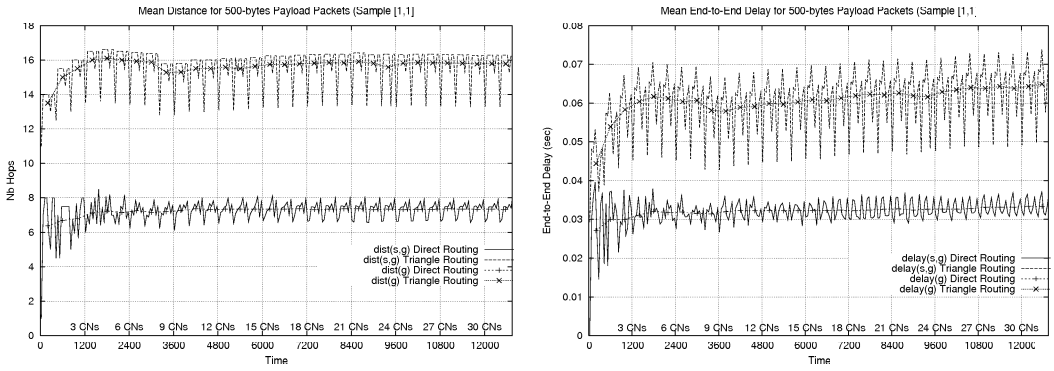

(a) Mean Distance

(b) Mean End-to-End Delay

Figure 9.4: Direct Routing vs Triangle Routing (Payload Packets)

9.2 Mobile IPv6

The performance evaluation of `Mobile IPv6` presented in this section is made for the purpose of demonstrating requirements and shortcomings for network mobility support. This exhibit `Mobile IPv6` limitations in face of a large number of CNs. In the first section we illustrate why *optimal routing* must be supported. Then, the second and third sections illustrate why the number of BUs must be limited, both in the fixed network and on the wireless link.

9.2.1 Optimal Routing

The simulation presented in this section was conducted according to scenario *Topology 1, Set 2* under tab. 8.1 and 8.2. We compare *direct routing* from the CNs to the MN and *triangle routing* from the CNs to the MN via the HA. We demonstrate that routing optimization is a feature that cannot be left aside if we want to avoid bandwidth waste, which becomes particularly significant for a large number of CNs.

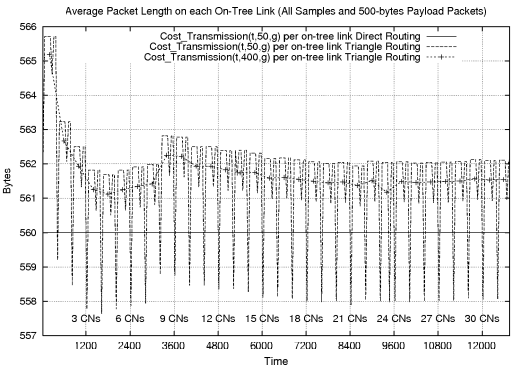

Figure 9.5: Packet Length on each *on-tree* Link

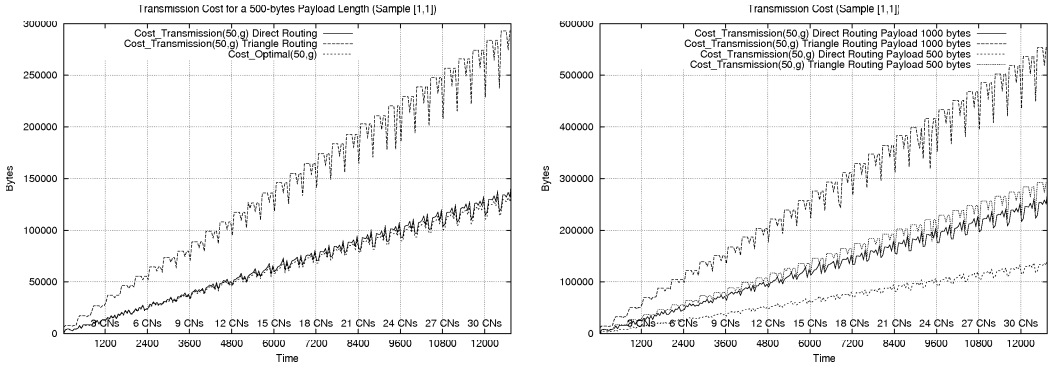

(a) Transmission Cost vs Optimal Cost

(b) Direct Routing vs Triangle Routing

Figure 9.6: Transmission Overhead for Payload Packets

Path Optimality Fig.9.4 (a) shows the *mean distance* between the group of CNs and the MN when routing optimization is used, and when it's not[1]. We observe that the route via the HA is on average more than two times the direct route, whatever the position of the MN in the network. As a result of a longer path, the *mean end-to-end delay* also increases, as shown on the second graph (e.g. 9.4 b). Both graphs also show that the mean distance and the mean end-to-end delay vary much more on the triangle route, then this further exhibits the need for routing optimization.

Packet Length Fig.9.5 shows the average packet length on each *on-tree link*, and for all simulation samples, with a 500-bytes payload. As we see, packets on the direct route have always a fixed length, whereas the average length on the triangle route varies according to the remoteness of the MN, due to encapsulation. The average packet length indeed depends on the ratio of the number of hops between the CNs and the HA over the number of hops between the HA and the MN (not explicitly shown on the graph). If measured independently of the location of the MN, the packet length on the direct route is on average shorter than on the triangle route (curve $cost_{transmission}(t, 400, g)$.). Then, on a per-link basis, the `Routing Extension Header` accounts for less overhead than encapsulation.

Transmission Overhead Fig.9.6 (a) shows the *transmission cost* $cost_{transmission}$ for payload packets in both conditions. In addition, we shows the *Optimal Cost* as for a comparison. As we see, the curves of

[1]The metric $dist(s, g)$ as expressed on the graph does not exactly follow our definition (section 8.4). Parameters s and g are taken the other way round and $dist(s, g)$ is indeed the distance from the group g of CNs to the MN

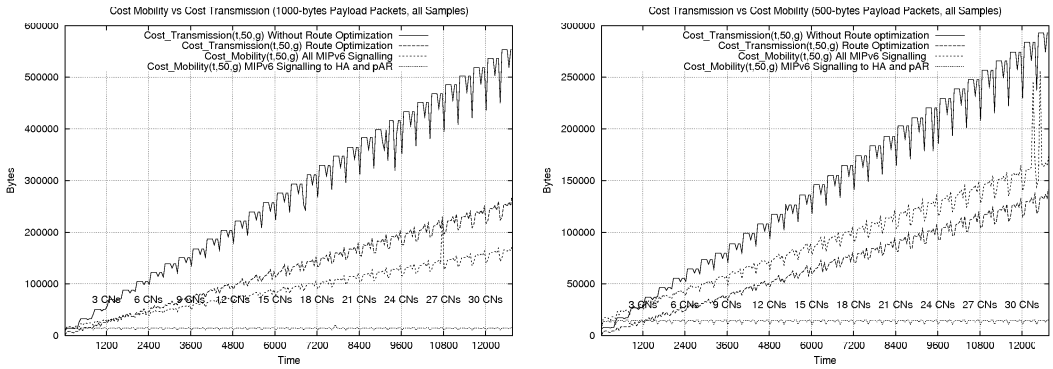

(a) 1000-bytes Payload Packet Length

(b) 500-bytes Payload Packet Length

Figure 9.7: Mobility Cost vs Transmission Cost

the transmission cost when routing optimization is used and the optimal cost are juxtaposed. This means that the `Routing Extension Header` only accounts for a minor overhead (20-bytes per packet). For large packets (1000 bytes), we see that the transmission cost over the triangle path is more than twice the transmission cost over the direct path (fig.9.6 (b)). The ratio between the two curves complies with the ratio between the curves on fig.9.4 and the average packet length (fig.9.5). However, for very small packets, the 20 additional bytes of the `Routing Extension Header` would account for a significant overhead compared to the optimal cost, but would still perform better than the transmission cost on the triangle route. This concludes that for large payload packets, the transmission cost when optimal routing is performed indeed tends to increase twice less than the *transmission cost* with no optimal routing.

Signaling Overhead On fig.9.7, the *mobility cost* of `Mobile IPv6` control messages is computed and compared against the *transmission cost* of data packets, over the direct path and the triangle path, for two distinct payload lengths (500 and 1000 bytes). The lower curve shows the mobility cost when no BUs are sent to CNs. This accounts for the BUs sent to the HA, its acknowledgments, and the BU sent to the previous AR. It is obvious from the graph that the mobility cost when BUs are sent to CNs, including BUs sent to the HA and AR increases linearly with the number of CNs. If we compare the mobility cost against the transmission cost, we see that for a large payload packet, even for a very low traffic rate (1000-bytes payload packets transmitted only every 50 seconds), it is still beneficial to perform routing optimization, while it's not for a half shorter payload packet with the same emission rate.

9.2.2 Overhead on the Wireless Link

For this set of measurements, we compute the `Mobile IPv6` overhead on a `IEEE MAC 802.11` wireless link. Our simulations are based on scenario *Topology 1, Set 1*. The packet length accounts for an additional 52 bytes that includes both the MAC and Ethernet header for each packet sent on the wireless link. In addition, the MN and the AR may exchange an optional `Request To Send (RTS)` and a `Clear To Send (CTS)` before the payload packet could actually be sent on the wireless medium. Although optional, particularly for short packets, we have set this feature on in our simulations.

We measured the signaling data rate on the wireless link during distinct interval of time (0.01, 1, 10, and 100 seconds). We made this count for `Mobile IPv6` signaling only (BUs sent to the HA, CNs, and previous AR, `Router Advertisements` only, and all the aggregated signaling, which includes `Mobile IPv6`, `Router Advertisements`, `MAC RTS / CTS`, and `ARP` (*all-signaling* curve).

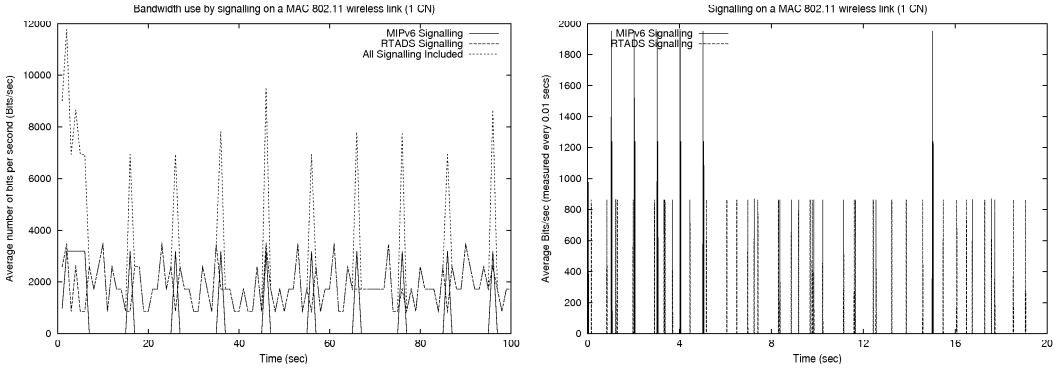

(a) 1 second interval (b) 0.01 second interval

Figure 9.8: Signaling on the wireless link (for 1 CN)

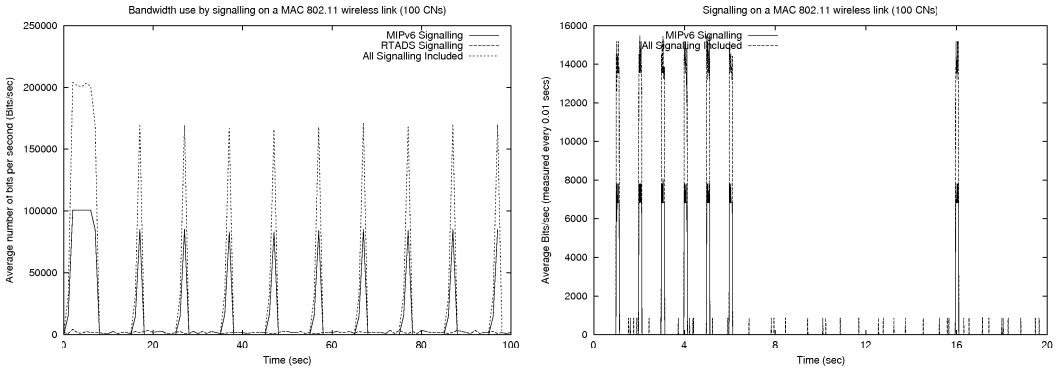

(a) 1 second interval (b) 0.01 second interval

Figure 9.9: Signaling on the wireless link (for 100 CNs)

Mobile IPv6 Overhead on the Wireless Link Figure 9.8 shows the data rate for the 3 different kinds of control packets with just one CN. Fig.(a) shows the total number of bits sent over a 1-second interval, whereas fig.(b) shows it over a 0.01-second interval. We can see that the total signaling is more than twice as much as `Mobile IPv6` on the average. We also note that the *all signaling* curve is parallel to the `Mobile IPv6` curve. The difference between the total signaling and `Mobile IPv6` signaling mainly accounts for the `MAC` signaling (not explicitly shown on the figures). Indeed, a `MAC RTS and CTS` exchange is sent prior to each `Mobile IPv6` packet. The ratio between the two curves is almost 2 because a BU on the link-layer is 122-bytes large whereas a `MAC RTS` is 44-bytes large and a `MAC CTS` is 38-bytes large, whatever the size of the data packet to transmit. This shows that, because of the link-layer overhead, the number of BUs transmitted over a wireless link is an issue more important to consider than the length of BUs itself. For a large number of CNs, it is thus preferable not to send the new careof address in separate BUs to avoid link-layer overhead. However, it is very unlikely `MAC RTS / CTS` is set for such small packets like BUs. The `MAC RTS / CTS` threshold should be fixed to a value over 122 bytes. We also note that `Router Advertisements` account for a constant bandwidth consumption about 2000 bits per second on the average[2]. When we compare fig.9.8 (1 CN) with fig.9.9 (100 CNs), we see that `Router Advertisements` do not account for a significant bandwidth consumption as the emission rate remains constant.

[2]Router Advertisements are sent every $random(1)$ sec

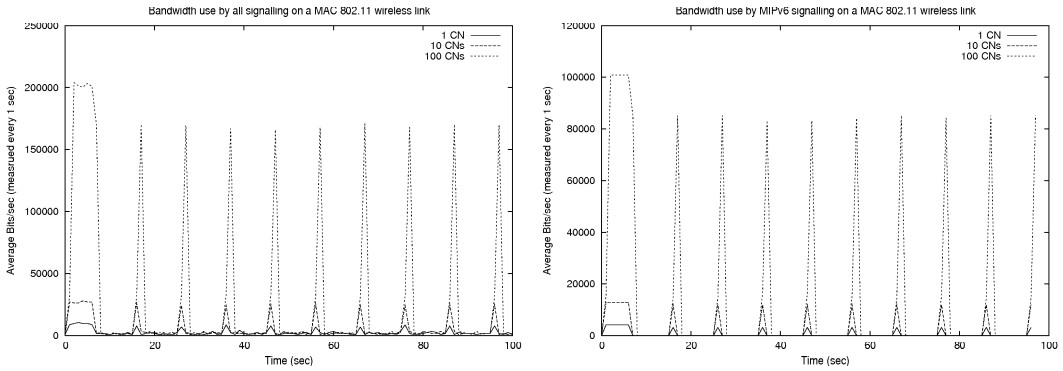

(a) MAC, MIPv6, and Routing Advertisements Signaling all included

(b) MIPv6 Signaling only

Figure 9.10: Periodic bursts on the wireless link (for 1, 10, 100 CNs)

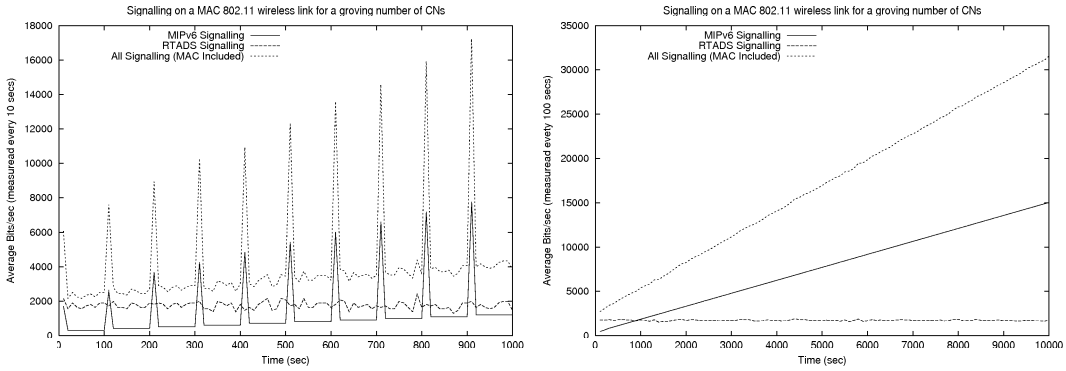

(a) Measurement every 10 secs (from 1 to 10 CNs)

(b) Measurement every 100 secs (from 1 to 100 CNs)

Figure 9.11: Signaling on the wireless link (for a growing number of CNs)

Periodic Burst As we can see from fig.9.8 and 9.9, there is a periodic burst, which corresponds to the interval between two BUs. Figures on the left show the total number of bits sent over a 1-second interval of time, whereas figures on the right show the total number of bits sent over a shorter period of time (0.01 second). Figures on the right show that a first set of BUs is sent with a 1-second interval, whereas the following BUs are sent with a 10-seconds interval, as exhibited on figures on the left. The first burst on fig.9.8 (a), 9.9 (a) and fig.9.10 is larger than the following ones since it accounts for 5 consecutive BUs. The first set corresponds to BUs sent as a consequence of obtaining a new careof address. The `Mobile IPv6` specification recommends to send 5 consecutive BUs with a 1-second interval in order to prevent the risk of loss (see section 3.1.1.5). This set of bursts cannot be avoided and is repeated every time the MN changes its point of attachment. On the other hand, 10-second interval periodic bursts correspond to BUs sent to refresh the corresponding Binding Cache entries at the CNs before their expiration. Of course, the periodicity of this second type of bursts can be limited if a more appropriate lifetime is set in BUs.

Binding Update Explosion on the Wireless Link This burst linearly increases with the number of CNs as this is shown in fig.9.11. One curve shows the bandwidth consumption when there are 100 CNs, whereas the second one shows it with 1 CN. On fig.(a), the bandwidth is averaged over a 10-second interval. As a consequence of this, bursts relative to the periodic BUs sent to refresh Binding Cache entries cannot be

perceived. This leave us with the interesting part: the bursts we see correspond to the 5 consecutive BUs sent after a change of careof address. As the number of CNs increases, this movement-related burst tends to consume a very significant part of the available bandwidth each time a MN changes its AR. This concludes that Mobile IPv6 signaling consumes periodically a significant part of the available bandwidth on the wireless link.

9.2.3 Conclusion

All our simulation results advocate a need to perform routing optimization. While routing optimization reduces the distance, it also reduces end-to-end delay, and the aggregated bandwidth consumption. Moreover, the more CNs we have, the more vital this becomes. Thus, we conclude that network mobility support *must* provide for optimal routing[3]. Graphs that followed have exhibited that a MN operating Mobile IPv6 to advertise its current careof address is a *bursty source*. It sends short packets bursts, separated by silent periods on the order of several seconds. As the number of CNs grows, this periodic burst consumes a fairly large amount of the available bandwidth, is propagated to the entire network, and causes a *Binding Update Explosion*.

From this set of experiments, we conclude that Mobile IPv6, while it provides for optimal routing, unfortunately accounts for quite an important consumption of the bandwidth. As the number of CNs grows, the periodic emission of BUs, whether it is due to a change of address or due to the periodic refreshment, is clearly inefficient and does not scale to a large number of CNs.

9.3 Standard Multicast Delivery of Binding Updates

In this section we compare the unicast delivery of Mobile IPv6 BUs against the standard multicast delivery of BUs. This comparison is made against the two most usual multicast delivery trees. Our aim is to determine when the multicast delivery is beneficial over the unicast delivery, and to determine which multicast distribution tree between a CBT and a SPT is the most appropriate.

Simulations were conducted according to scenario *Topology 1, Set 2* under tab. 8.2 and 8.1. All routers are multicast-enabled. We assume that an optimal SPT is always rooted from the current point of attachment of the MN to all CNs. Under CBT, we assume the existence of a delivery tree between the core and the CNs recorded in the MN's Binding List. BUs are encapsulated from the MN to the core, where it is decapsulated and transmitted on the tree. The CBT is compared with several locations of the core: at the HA, at the BR that connects the home site to the domain's backbone, and to a router in the backbone.

9.3.1 Unicast vs Multicast

In this section, the unicast delivery of Mobile IPv6 BUs is compared against the multicast delivery of BUs in terms of *on-tree links* and *mobility cost*. Curves show results for a unicast delivery of BUs, for a multicast delivery by means of SPT, and by means of CBT with the core located at the HA.

On-Tree Links Overhead Fig.9.12 compares the number of on-tree links in each situations. As we observe, the number of links with the unicast delivery increases linearly with the number of CNs while all

[3]This may probably be common sense to many readers, but we want to stress this because the need for routing optimization is criticized by many folks, probably the same who pretend that there is no need for mobility support at all. These critics are based on the assumption that many mobile hosts initiate the communication, which is probably true as far as today's communication are concerned, but probably not anymore in the future

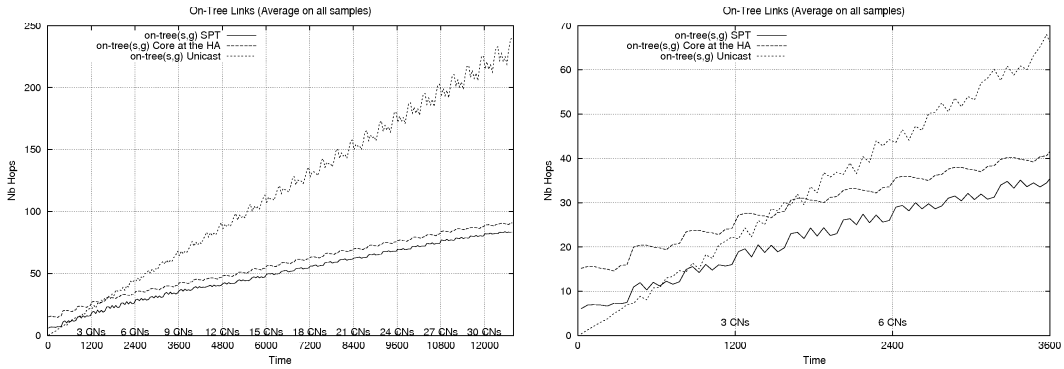

(a) On-Tree Links: All Samples

(b) On-Tree Links: Zoom All Samples

Figure 9.12: Unicast vs Multicast: On-Tree Links

multicast curves are $log(\sharp CN)$. In term of number of links consumed, any multicast technique is advantageous from a very small number of CNs. The average of all 25 simulations shows that the threshold is 2 for a SPT, and 4 for a CBT with the core at the HA.

Mobility Management Overhead The *mobility cost* is illustrated on fig.9.13. It is proportional to the number of on-tree links since all BUs have the same length whether they are distributed by unicast of by multicast. In (a) and (b), we have computed all BUs, including BUs sent to the HA and the previous AR; the right-hand side shows a zoom and a restricted number of CNs. The mobility cost in (c) and (d) accounts only for BUs sent to CNs in (c) and (d), for the same simulation sample. To conclude, (e) and (f) show these results averaged over all simulation samples. This shows that the unicast delivery of BUs does not perform well against multicast. For this topology, the crossover point is around 4 CNs. In addition, CBT and SPT perform nearly equivalently if compared to the unicast delivery.

Conclusion Those results show that it is beneficial to use multicast over unicast from a very small number of CNs, provided that the multicast signaling cost is negligible. It also shows that delivering BUs by means of multicast performs better whatever multicast delivery tree is used. This also means that the candidate multicast tree may be chosen alternatively between a CBT-based multicast delivery tree and a SPT-based multicast delivery tree. From these results, it is straight forward to note that the multicast delivery avoids the periodic *Binding Update Explosion* on both the wireless link and in the network. This is a particularly good advantage on the wireless link, due to its limited bandwidth.

9.3.2 Core-Based Tree vs Shortest-Path Tree

In this section, CBT and a SPT are compared in term of *tree size*, *on-tree links* and maximum *end-to-end latency*. Curves show results for a multicast delivery by means of SPT, and by means of CBT with the three positions of the core.

On-Tree Links Fig 9.14 (a) and (b) show the number of *on-tree* links measured in one simulation (left-hand side) and the average of all simulations (right-hand side). As expected, the number of *on-tree* links grows linearly with the number of CNs and SPT consumes less links than its CBT counterpart. This is further highlighted on the zoom fig 9.14 (c). The oscillations on the curve exhibit that the number of on-tree links varies more according to the location of the MN under a SPT tree than a CBT tree. The reason is that

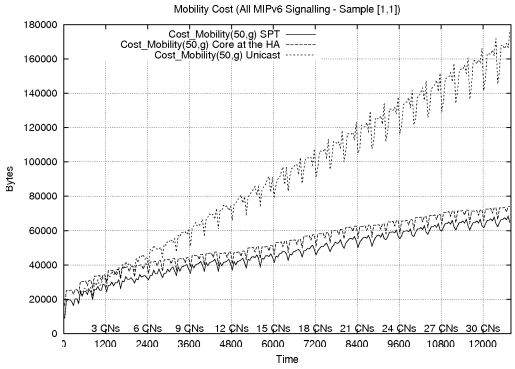

(a) Mobility Cost: BUs to HA and pAR included

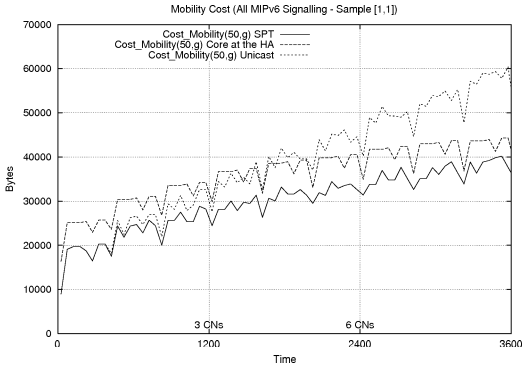

(b) Mobility Cost: BU to HA and pAR included (zoom)

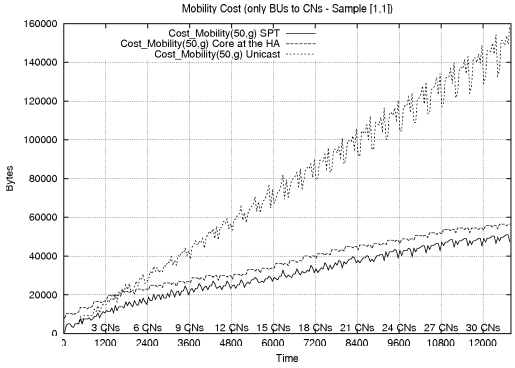

(c) Mobility Cost: BUs to CNs only (one sample)

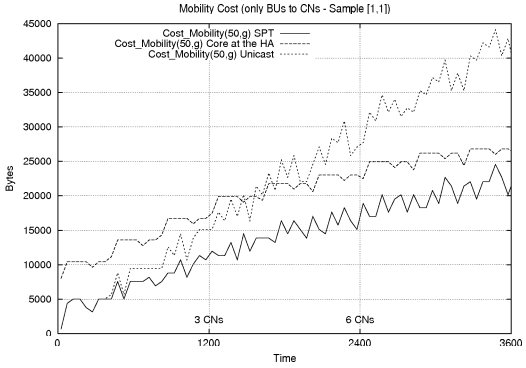

(d) Mobility Cost: BU to CNs only (zoom, one sample)

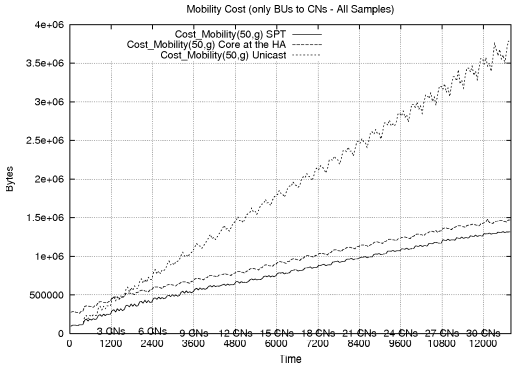

(e) Mobility Cost: BU to CNs only (all samples)

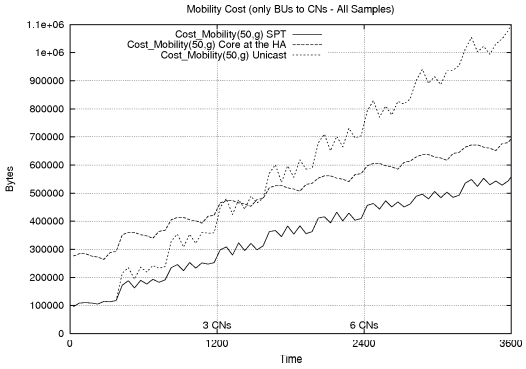

(f) Mobility Cost: BU to CNs only (zoom, all samples)

Figure 9.13: Unicast vs Multicast: Mobility Cost

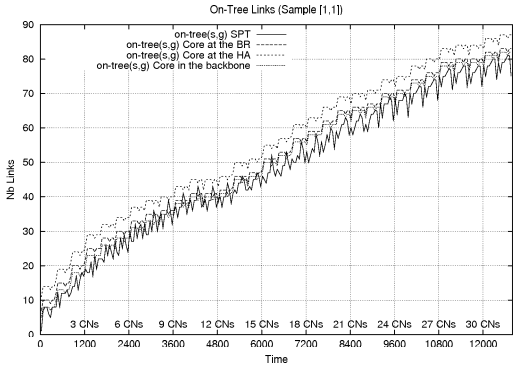

(a) On-Tree Links (Simulation Sample 1,1)

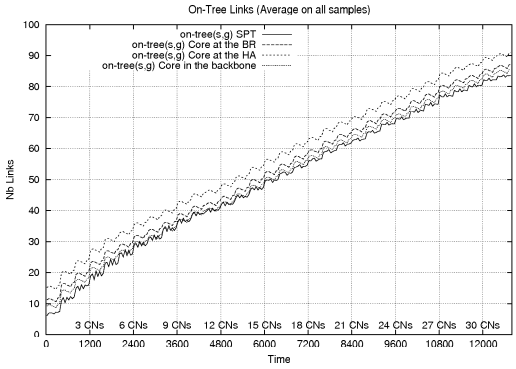

(b) On-Tree Links Average on all Simulation Samples

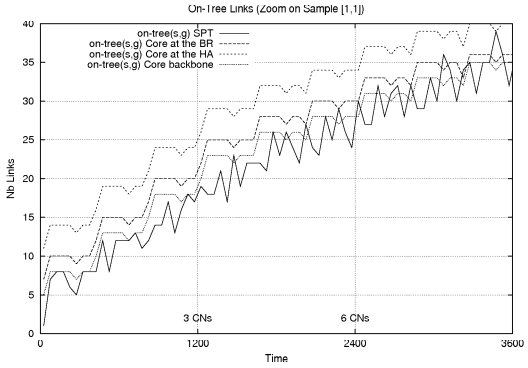

(c) On-Tree Links (zoom)

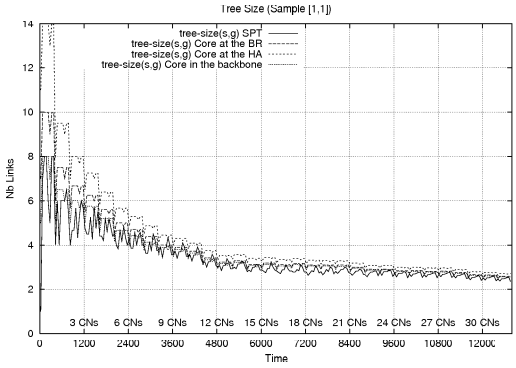

(d) Tree Size

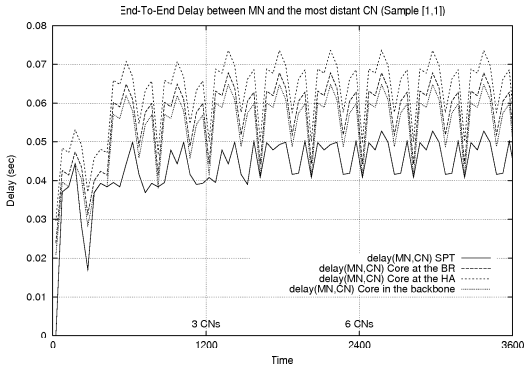

(e) End-To-End Delay (zoom on sample (1,1)

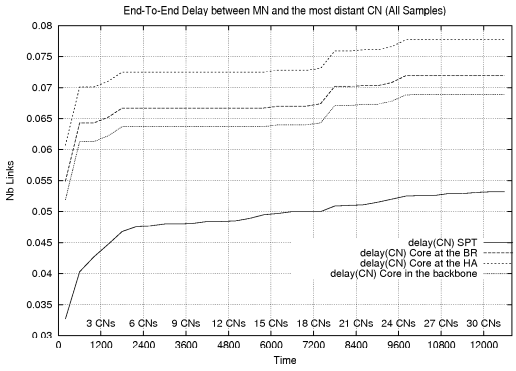

(f) End-To-End Delay Average on all Samples

Figure 9.14: CBT vs SPT

the SPT tree is updated upon every displacement of the MN, while the CBT is not subject to any tree update as a result of the displacement since packets are always encapsulated to the core. We also show the *tree-size* on fig.9.14 (d). The tree size tends to a value between 3 and 4 in this topology, whatever the number of CNs. As the number of CNs grows, the impact of the displacement of the MN decreases over both the number of on-tree links and the tree-size. This complies with the observation we made in section 9.1.

End-To-End Delay Fig.9.14 (e) and (f) show the maximum latency experienced by a member of the group to receive a BU. SPT performs better than the CBT, for all positions of the core. Under SPT, the delay is shorter and varies less according to the location of the MN.

Conclusion Not surprisingly, SPT always performs a bit better than CBT in terms of delay and number of on-tree links, whatever the number of CNs, their location, and the location of the MN, since the delivery tree is always optimal. However, the curves also evidence the fact that the SPT is more subject to tree updates than the CBT. Since updating the tree upon every displacement typically has an important cost, this would account to prefer a CBT over a SPT in such a situation. We can also conclude that a SPT technique is not suitable for frequent displacements, i.e. Local-Area Mobility. On the other hand, we notice that the performance is nearly the same whatever the position of the core for the CBT. Without surprise, the core located at the HA exhibits the worst results amongst the simulated location. The HA is the most distant core from the MN, unless the MN visits its home site. In average, it is also the most distant from the CNs. The core better located in the backbone of course performs a little better.

The evaluation of several positions of the core also shows that the position of the core is a less important issue for a mobile source performing Wide-Area Mobility displacements, when CNs are uniformly distributed, because the distance from the core is averaged over the visited sites. Indeed, any position of the core could be chosen if the visited sites are also uniformly distributed. Due to the fact that the MN may also spend a significant amount of time in its home site, and that none of the chosen core exhibit significant worse results than another, a transient conclusion is that the Mobile IPv6 HA, or the BR, would be good candidates for the core. A means to send BUs directly from the HA on behalf of the MN should be also investigated; this could both reduce the path and delay of BUs and facilitate the multicast delivery.

The evaluation of a suitable candidate protocol is left for future study. The performance of a multicast protocol highly depends on implementation choices. This would need to balance the benefit of using multicast against the signaling cost inherent to the multicast protocol for group membership management and multicast tree maintenance. Up to this date, no inter-domain multicast protocol has been standardized. This is still an open research area, as mentioned in section 2.3.2.2 [Ramalho, 2000; Almeroth, 2000]. Thus, it would probably be too immature to conduct this study now.

9.4 List Based Multicast Delivery of Binding Updates

In this section, we evaluate the performance of the List-Based Multicast delivery of Mobile IPv6 BUs and we compare it against the standard Mobile IPv6 unicast delivery. Simulations were conducted according to scenario *Topology 1, Set 2* under tab. 8.2 and 8.1.

9.4.1 Which routers should be LBM-enabled

The performance of List-Based Multicast is evaluated under four situations: when all routers are LBM-enabled, when only CNs are LBM-enabled, when both CNs and border routers are enabled, and when both CNs and transit router are enabled. In addition, we set the flag to 1 in the LBM Header to tell the LBM-enabled routers to remove the destination from the list. The LBM feature is always set at the correspondent nodes to guarantee the delivery of BU to all destinations specified in the packet.

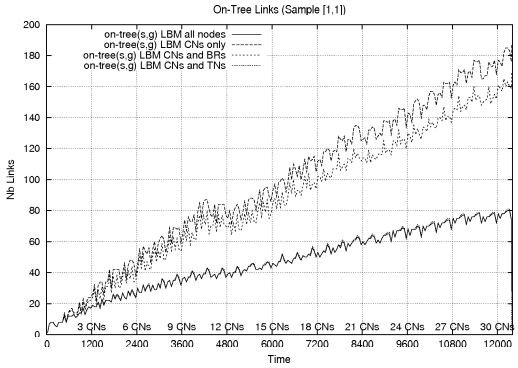

(a) On-Tree Links (All Samples)

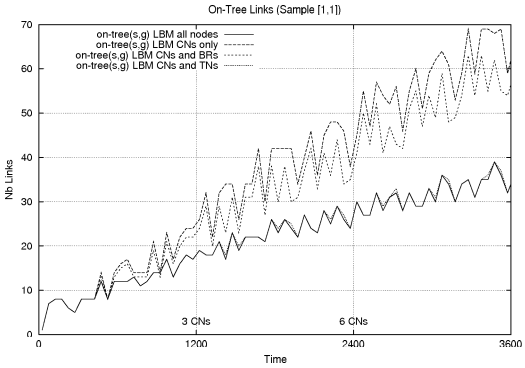

(b) On-Tree Links (Zoom All Samples)

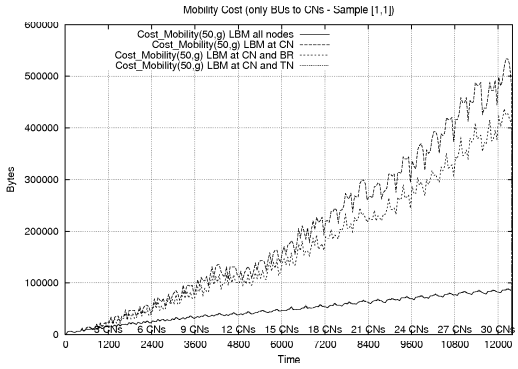

(c) Mobility Cost: BUs to CNs only (All Samples)

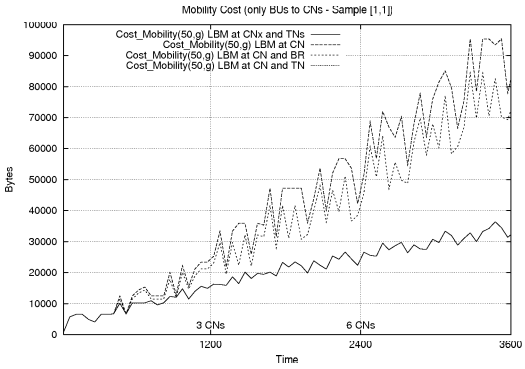

(d) Mobility Cost: BUs to CNs only (Zoom All Samples)

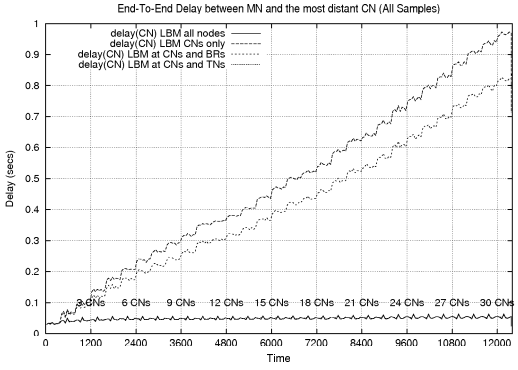

(e) End-To-End Delay (All Samples)

Figure 9.15: List Based Multicast (destinations removed from the list)

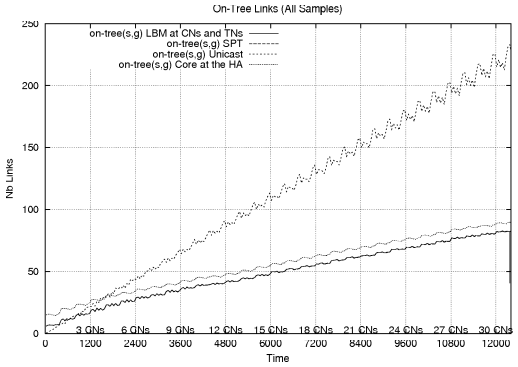

(a) On-Tree Links (All Samples)

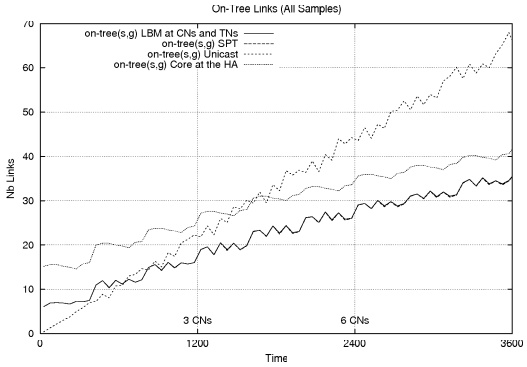

(b) On-Tree Links (zoom all samples)

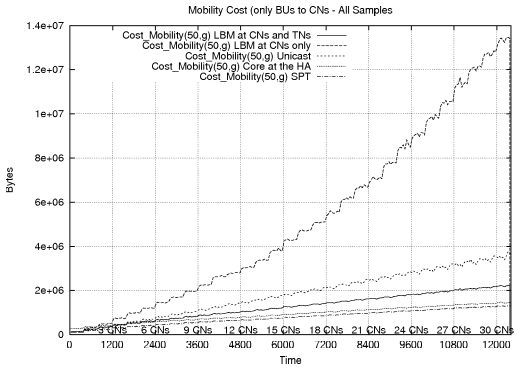

(c) Mobility Cost: BUs to CNs only (All Samples)

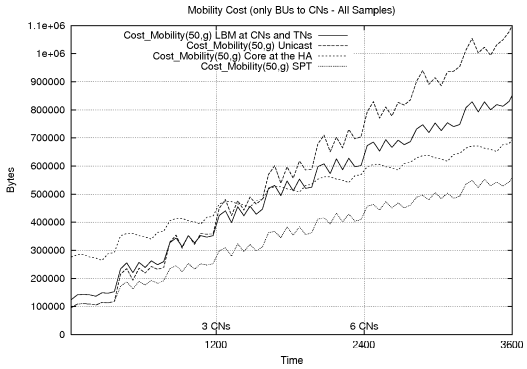

(d) Mobility Cost: BUs to CNs only (All Samples)

Figure 9.16: LBM vs SPT vs CBT vs Unicast

On-Tree Links Fig.9.15 (a) and (b) shows the number of *on-tree links* for one simulation sample. The number of on-tree links when all routers are `LBM`-enabled is nearly equal to the situation where only transit routers and CNs are `LBM`-enabled (two lower curves). Similarly, the number of on-tree links is equivalent when only CNs are `LBM`-enabled and when both BRs and CNs are `LBM`-enabled (two upper curves). The reason is that CNs are uniformly distributed in the topology (the probability that two CNs are in the same site is low) and 1 to 3 hops away from BRs. We also see that the slope of the upper curves is higher than the lower curves. The former two increase linearily whereas the latter two look like $log(\sharp CN)$.

Mobility Management Overhead The *mobility cost* is shown on fig.9.15 (c) and (d). This time, it is not proportional to the number of on-tree links, since the `LBM` packet length is increasing with the number of CNs. `List-Based Multicast` performs badly when the feature is not well deployed in the network since BUs are bounced from one `LBM`-speaker to another. We note that the slope of all curves is not constant and slows down between 9 and 12 CNs, as for the *on-tree links*, particularly the two upper curves. On the *on-tree links* curves, the slope is higher before 9 CNs, and lower after 12 CNs; on the *mobility cost* curves, the slope is lower before 9 CNs and higher after 12 CNs. This seems to indicate first an optimal value when a BU is most likely to be duplicated, and second an optimal value when the list of CNs is most likely to be split into two equivalent number of CNs. As the BU progresses in the network, the combination of these two events is most unlikely to happen.

End-To-End Delay The maximum delay between the MN and the most distant CN is shown on fig.9.15 (e). The maximum delay is bounded to some limit when all routers or at least transit routers are LBM-enabled. It obviously tends to increase rapidly to a very large *end-to-end delay* when only CNs are able to process the LBM Header. In this case, the packet traverses all CNs before reaching its ultimate destination.

Conclusion This experiment shows that List-Based Multicast does not need to be deployed in all routers to perform well. However, it clearly shows that at least some well-dedicated List-Based Multicast routers should be deployed, at transit points like in the backbone network for instance. In this case, it scales to an increasing number of CNs, otherwise LBM performs very badly from a very low number of CNs

9.4.2 LBM vs all schemes

To conclude our simulations, we compare List-Based Multicast against the unicast delivery of BUs and the multicast delivery using a CBT centered on the HA, and a SPT. Curves show List-Based Multicast when only CNs are LBM-enabled, and when both CNs and transit routers are LBM-enabled.

On-Tree Links Fig.9.16 (a) and (b) shows *on-tree links*. If LBM-enabled nodes are well located in the network, we observe that List-Based Multicast compares to a SPT in terms of *on-tree links*, although all routers do not process the List-Based Multicast option. On the zoom, we can see that List-Based Multicast is more appropriate than unicast from a number of CNs around 5 on this curve.

Mobility Cost In terms of *mobility cost*, we can only compare List-Based Multicast with the unicast delivery, since there is no multicast signaling cost whatsoever ($cost\_multicast = 0$) in both cases. From fig.9.16 (c) and (d), we can see that List-Based Multicast compares very badly against unicast Mobile IPv6 when only CNs are able to process the extension. However, the benefit of List-Based Multicast over unicast is obvious when there exists a minimum number of LBM-enabled routers well located in the network.

Conclusion This experiment shows that List-Based Multicast performs very well against the unicast delivery of BUs without requiring any additional signaling. There is nevertheless provisions concerning the deployment of the LBM capability in well-located routers since List-Based Multicast is clearly inefficient in terms of delay and mobility cost when only CNs are able to process the option.

9.5 Conclusion

Our simulations have evidenced the fact that routing optimization is a necessary feature to reduce the network load and the transmission delays. However, routing optimization is made at the expense of periodic BUs that must be sent individually and periodically to every correspondent nodes. When these messages are sent to a large numbers of correspondent nodes, it causes a periodic signaling burst (*Binding Update explosion* that is propagated in the network. Links close to the node that emits these BUs are most likely to suffer from this periodic burst, particularly over the air where a significant amount of the available bandwidth is consumed as this is highlighted by our simulations on a MAC 802.11 wireless link.

It is obvious that the standard multicast technique consumes the same amount of bandwidth on a given link whatever the number of CNs, then the benefit of multicast is straightforward, particularly on the wireless link. The comparison between the two main multicast tree types shows that the SPT and the CBT perform nearly equivalently, but SPT is more subject to tree updates and therefore is not adequate to frequent moves. Thus, we would recommend to choose the CBT over the SPT.

The evaluation of List-Based Multicast shows that it performs rather well compared to the unicast delivery, all the more that the signaling cost of List-Based Multicast is null. It even exhibits good results for a large number of CNs provided a number of LBM-enabled routers are well deployed in the Internet. List-Based Multicast seems nevertheless better adapted to transmit Prefix Scope Binding Updates to a bounded number of CNs since the size of the packet is increasing with the number of CNs. Thus, for a very large number of CNs, traditional multicast may be better adapted to transmit Prefix Scope Binding Updates, provided a cheap *inter-domain multicast protocol* can be selected, in terms of *multicast overhead*. The simulation comparison between standard multicast and List-Based Multicast complies with the intuitive comparison in section 2.3.4.

We note that all our simulation results are applicable whether the mobile node is a *host* or a *router*. Given our mobility model, our results could also be applied to Hierarchical Mobile IPv6. The *mobility cost* for a MN that perform Wide-Area Mobility displacements is indeed equivalent whether the MN performs Mobile IPv6 or Hierarchical Mobile IPv6. It it performs Hierarchical Mobile IPv6, the MAP is a good candidate for the core. BUs could be encapsulated from the MN to the MAP, and then forwarded along a delivery tree rooted at the MAP.

Conclusions and Perspectives

Conclusions and Perspectives

Contributions

In this dissertation, we have studied *network mobility support* as opposed to *host mobility support*. We first presented all the pieces on which our work is based: the architecture of the Internet; the TCP/IP protocol suite, particularly IPv6. In chapter 2, we highlighted that the mobility problem is mainly an addressing translation problem and we advocated the need for a permanent and invariant node identifier, i.e. a separation between the identifier of the node (node identifier) and the identifier of the node's point of attachment to the network (location identifier). Our study of a number of *host mobility support* proposals and their classification into a taxonomy showed that the mobility problem can be overcome by several means. The most common means is *two-tier addressing* which binds two addresses to a mobile node, one used as a permanent and location invariant identifier, and one used as a routing directive. This approach is probably the best one since it does not require a redesign of the TCP/IP *reference model*. Indeed, although the advent of IPv6 offers the opportunity to make orthogonal design choices, making such orthogonal design changes could take mobility into better account, but it wouldn't allow smooth transition from the aging IPv4 network layer to its rightful heir, IPv6 and all the new protocols that assist it.

Since *network mobility support* has not yet been specifically considered in the literature, though sometimes mentioned, we have defined, in section 1.3, a new terminology specifically targeted for mobile networks. This terminology was particularly useful to define the architecture elements, characteristics, issues, goals, problems and requirements pertaining to *network mobility support* (chapter 5). This work has been submitted to the IETF as an Internet Draft [Ernst et al., 2001b].

We then investigated a few approaches, based on our study of *host mobility support* proposals. *Two-Tier Addressing* is the approach followed by Mobile IPv6, the most advanced *host mobility support* solution proposed for standardization at the IETF. Although Mobile IPv6 is not designed to support mobile networks, we have studied its suitability and we came to the conclusion that it could serve as a base protocol. There are, of course, a number of issues to address, the most important being providing continuous Internet connectivity, optimal routing, and scalability.

Thus, an important part of this study focused on enhancements that address Mobile IPv6 shortcomings. In particular, we address the questions of offering permanent Internet connectivity to the mobile network, optimal routing, while scaling to an important number of correspondent nodes. The mobile router that connects the mobile network to the Internet performs Mobile IPv6 and obtains subsequent careof addresses on each visited link. Our first proposal, Prefix Scope Binding Update is a necessary Mobile IPv6 extension for advertising the new careof address as the routing directive to route packets from correspondent nodes to the mobile network. It proposes binding the careof address obtained by the mobile router to the mobile network prefix, i.e. the prefix that is common to the addresses of all mobile network nodes. The advertisement of this binding is done by means of a Prefix Scope Binding Update message, i.e. an enhanced BU. As a result of this message, a *network-specific route* is recorded in the Binding Cache at the home agent and at the correspondent nodes as opposed to a *host-specific route* in the standard Mobile IPv6.

However, Prefix Scope Binding Update does not solve the scalability issues since BUs are sent individually to each correspondent node. For a very large number of correspondent nodes, as it is envisioned for some specific instances of mobile networks like a train or an aircraft, this would lead to a *Binding Update Explosion*. The second proposal we proposed is extending Mobile IPv6 with multicast capabilities whereby multicast protocols are used to deliver Prefix Scope Binding Updates to correspondent nodes. The distribution of BUs by multicast is an innovative way to manage mobility. Because the cost of multicast management might overcome the gain in terms of Mobile IPv6 signaling, we proposed using another multicast technique, List-Based Multicast, which records the addresses of the destinations in the packet itself. Traditional multicast is better envisioned for a very large number of correspondent nodes and also assumes that a scalable inter-domain multicast protocol is deployed in the near future. On the other hand, List-Based Multicast is better envisioned for a smaller number of correspondent nodes since the size of the packet is growing with the number of correspondent nodes recorded in the packet. To overcome this, the correspondent nodes could be split into several distinct lists. Our propositions could also be combined with various techniques in order to offer an overall panel of solutions as it is outlined in section 7.4.

We validated the performance of our solutions by means of simulation, using NS-2, which required important enhancements to the publicly available code. We contributed our enhancements, refereed to as MobiWan [Ernst, 2001b], to the NS-2 community. Our simulations are mainly concerned with measuring signaling incurred in the core network. Interestingly, our simulations are not targeted to mobile networks specifically. This means that simulation results are applicable to both mobile nodes and mobile networks for a significant number of correspondent nodes. Moreover, our conclusions can be drawn alternatively for Mobile IPv6 and Hierarchical Mobile IPv6. Our simulation results showed that the use of multicast is beneficial from a small number of correspondent nodes while List-Based Multicast performs excellently for a bounded number of correspondent nodes.

The need to displace an entire network in the Internet is perceived for a number of years, but the issues pertaining to this topic were never, or nearly never considered. At first sight, protocols that support *host mobility* were thought to be adequate without prior investigation of their suitability. We then went to the IETF to discuss about this in August 2000. We presented some of the issues to the Mobile IP Working Group and proposed Prefix Scope Binding Update (as described in section 7.1) as an immediate solution for an immediate need (submitted as an IETF Internet Draft [Ernst et al., 2000b, 2001a]). At that time, this didn't raise concerns from the attendance; only some security concerns were raised. Attendees whether didn't understand the need for mobile networks, or didn't understand the problems and the issues, or were simply not interested in it. We had to wait until the following IETF meeting before voices started to raise on the mailing list. The discussion started around the question of authorization to send Prefix Scope Binding Updates, and gradually developed to other issues. More and more people have joined the discussion since then which proves the interest of the research community in this topic, to the point that some extensions similar to our Prefix Scope Binding Update proposal were brought to the Mobile IPv6 specification (version 14, July 2001).

At 51st IETF meeting, August 2001, London, we advocated that issues turning around *network mobility support* were not enough understood to be part of the Mobile IPv6 specification, and would only result in further delaying the progress of this specification to an IETF Proposed Standard RFC. We met consensus on this. As a result, the recent additions to the specification were taken out from the draft. We also presented our terminology, our list of requirements (as detailed in chapter 5, and submitted as an IETF Internet-Draft [Ernst et al., 2001b]) and argued whether or not the Mobile IP Working Group is the right place to discuss *network mobility support*. The outcome from the meeting is that we are now going to propose the creation of a new IETF Working Group (i.e. a BOF) to discuss this topic specifically. A starting premise of this new activity is that both Mobile IPv6 and Hierarchical Mobile IPv6 are candidate protocols for *network mobility support*, but a number of enhancements are obviously needed.

We have therefore contributed in this new subject in a number of ways. First, we raised this promising new area of research in a number of conferences, particularly at the IETF. We also had the opportunity to publish our work as conference papers [Ernst et al., 1999, 2000c], conference presentations [Ernst, 2000,

2001a], European Patents [Ernst and Lach, 2000; Ernst and Castelluccia, 2001], internet-drafts [Ernst et al., 2001b,a] and package distribution [Ernst, 2001b]. Most documents may be found on the web page of the author at time of writing of this dissertation [PLANETE].

Perspectives and Future Work

Our study has only focused on some aspects of the topic: providing continuous Internet connectivity to mobile network nodes, allowing optimal routing between the correspondent nodes and the mobile network nodes and minimizing signaling. Other aspects are left for future study. We would particularly like to consider security aspects, nested mobility, multi-homing, and efficient routing between two mobile networks. We have also not yet addressed a number of issues depicted in section 5.2. Particularly, the impact on the routing protocols running in the visited networks and their interaction with the routing protocol running within the mobile network deserves more consideration. Other potential approaches, as listed in section 5.5 must also be investigated. Among them, the routing-based approach must be considered in priority, in combination with our approach, as a solution to support Local-Area Mobility. Also, as demonstrated by our simulation results, our proposals could be applied to Hierarchical Mobile IPv6, which would require some extensions.

Our solutions themselves raise a number of issues and highly depend on the progress made at the IETF in various working Groups. Concerning security, Prefix Scope Binding Update faces authorization concerns while the multicast distribution of Prefix Scope Binding Updates faces authentication concerns. We are presently working on these security aspects and we monitor the current work and discussions from the Mobile IP, the IPv6[4], and SMUG working groups. Work at the IRTF Working Group about multicast security must also be investigated. Concerning multicast, the cost of traditional multicast must be evaluated when an appropriate *inter-domain multicast protocol* is standardized in the IDMR working group. The MAGMA working group which works on group membership and the SSM (Source Specific Multicast) working group must also be monitored.

More performance evaluations must also be conducted. We are particularly willing to perform more extensive simulation sets and to validate our solutions under different mobility and CN-inter-arrival-time models, with different scenarios. All the simulations outlined in this present document were ran with topologies generated with the same *transit-stub* GT-ITM model. We would like to perform our simulations with bigger topologies, and with other topology patterns as found in GT-ITM and in BRITE (Boston University Representative Internet Topology Generator) [Medina et al., 2001], a new topology generator. We also would like to evaluate the performance of List-Based Multicast when correspondent nodes are randomly distributed in a limited number of sites. On a more prospective basis, we are also working on the deployment of List-Based Multicast and a set of architectures that combine multicast with other mechanisms.

More questions were probably raised rather than answered during the course of this study. In a sense, the present document does the spade-work on the question of *network mobility support* and is meant to open the debate in a subject of growing interest. We have indeed opened the door of a large research topic which should drive an important research effort in the near future. This topic is already subject to the interest of car manufacturers, airline companies, public transportation companies, and wireless vendors of any kind. Many more should join.

[4]The IPv6 working group was previously called IPNG (IP Next Generation)

Appendix A

Abbreviations

- AR: access router.
- AP: access point.
- BR: border router.
- CN: correspondent node.
- FA: Forwarding Agent.
- HA: home agent.
- MA: Mobility Agent.
- MN: mobile node.
- MH: mobile host.
- MN_{ip}: home address of the mobile node MN.
- MN_{coa}: careof address of the mobile node MN.
- MNN: mobile network node.
- MNN_{ip}: IP address of the mobile network node MNN.
- MR: mobile router.
- MR_{ip}: home address of the mobile router MR.
- MR_{coa}: careof address of the mobile router MR.
- LA: Locating Agent.
- LFN: local fixed node.
- LMN: local mobile node.
- $LCoA$: local careof address.
- RA: Redirecting Agent.
- $RCoA$: regional careof address.
- UA: Updating Agent.

- VMN: visiting mobile node.
- VMN_{ip}: home address of the visiting mobile node VMN.
- VMN_{coa}: careof address of the visiting mobile node VMN.

Appendix B

Résumé Détaillé en Français

Dans cette dissertation, nous étudions de nouveaux aspects de la mobilité dans le réseau Internet. Nous consacrons cette étude à IPv6, la nouvelle version de Internet Protocol (IP), le protocole de la couche réseau qui régit les communications entre deux machines de l'Internet.

Introduction

Motivations et Objectifs Les travaux traditionnels dans le domaine de la mobilité se sont jusqu'à présent toujours intéressés au support des *stations mobiles*, c'est-à-dire des stations changeant de point d'ancrage dans la topologie Internet. Le but recherché est la fourniture d'une connectivité Internet permanente, sans que les communications en cours ne soient interrompues suite au changement de point d'ancrage. Or, il est également souhaitable qu'un réseau tout entier puisse se déplacer et changer son point d'ancrage. Nous ferons dorénavant référence à un *réseau mobile* comme étant un ou plusieurs sous-réseaux connectés à l'Internet par l'intermédiaire d'un ou plusieurs routeurs qui changent leur point d'ancrage au cours de leurs déplacements. Entre autres applications, un train, un avion, un bus ou une voiture, connecté à l'Internet et offrant à ses passagers des bornes d'accès à l'Internet, tout en assurant la transmission de données nécessaires à la navigation, sont les exemples les plus évidents de *réseaux mobiles.*

Malgré le besoin de fournir un accès Internet permanent à toutes les stations localisées dans un *réseau mobile*, suscité soit par les fabricants de véhicules, soit par les compagnies de transport, soit par les usagers eux-mêmes, aucun travail sérieux se préoccupant des problèmes spécifiques liés au déplacement d'un réseau n'a encore vu le jour dans ce domaine. D'autre part, les exemples de *réseaux mobiles* cités ci-dessus ont pour particularité d' accueillir des stations qui peuvent elles-mêmes être mobiles si on considère les téléphones portables, ordinateurs portables, et autres moyens techniques permettant de se connecter à l'Internet et transportés par les passagers. De tels exemples justifient le besoin de s'intéresser aux problèmes causés par plusieurs niveaux de mobilité, par des réseaux de très grande taille pouvant accueillir des centaines de stations, et par des déplacements dans des parties très éloignées de la topologie Internet, appartenant à des fournisseurs de services distincts.

Cette dissertation a donc pour objet l'étude du support de la mobilité pour les *réseaux mobiles* et leur problématique. Nous avons pour but le défrichage du terrain en vue de faire avancer la réflexion dans ce sujet et de proposer des solutions spécifiques. Comme nous allons le voir, cette étude s'intéresse en particulier aux questions de localisation, de routage et d'adressage.

Structure du Document Dans un premier temps, nous présentons les fondements de l'Internet, c'est-à-dire la suite de protocoles qui permet la communication entre deux stations distinctes, ainsi que le prob-

lème d'ordre général posé par la mobilité. Est ensuite présenté l'État de l'Art dans le domaine de la mobilité, c'est-à-dire les propositions permettant la gestion des *stations mobiles*. Faisant suite à l'étude des diverses propositions, nous dressons une taxinomie qui permet de mettre en avant leurs points communs et leur différences. La deuxième partie entre dans le vif du sujet et s'intéresse donc spécifiquement au support des *réseaux mobiles*. Tout d'abord, nous définissons une terminologie permettant de dresser les caractéristiques des *réseaux mobiles* et leur problématique. Plusieurs approches, identifiées dans la seconde partie, sont envisageables pour résoudre les problèmes de localisation, de routage, et d'adressage. Une solution basée sur `Mobile IPv6`, la solution standard de l'`IETF` pour le support des *stations mobiles*, semblant abordable, nous étudions les forces et les faiblesses de ce protocole pour supporter également les *réseaux mobiles*. Cette étude nous permet de proposer un certain nombre d'extensions. En particulier, nous nous intéressons à réduire le coût des messages de contrôle généré par ce protocole et nécessaire à la localisation du *réseau mobile*. Partant d'un certain nombre d'observations, nous proposons de réduire le coût des messages de contrôle induit par ce protocole au moyen de deux techniques multipoint. La première, dite traditionnelle, établit un arbre de distribution entre le *réseau mobile* et ses correspondants. La deuxième enregistre directement la liste des correspondants dans le message de contrôle. Nous concluons cette partie par une vue d'ensemble d'une nouvelle architecture de gestion de la mobilité rassemblant diverses techniques rencontrées lors de l'étude de l'État de l'Art ainsi que nos extensions multipoint. La dernière partie est consacrée à l'évaluation de la performance de nos extensions multipoint, par simulation. Nous concluons cette dissertation par un certain nombre d'observations, de perspectives et d'axes de recherche futurs. Enfin, notre conclusion récapitule notre contributions et discute des perspectives de notre travail.

Première Partie: Internet et la Mobilité

Chapitre 1: Terminologie Dans ce chapitre, nous présentons l'essentiel des définitions nécessaires à notre étude, c'est-à-dire les termes permettant de décrire le problème de la mobilité. Tout d'abord, nous présentons la terminologie en usage et l'architecture de l'`Internet`. Nous définissons ensuite ce que l'on entend par *mobilité dans la topologie `Internet`*, c'est-à-dire *mobilité dans la couche `IP`*. Il convient de différencier la *mobilité locale* (mobilité entre sous-réseaux topologiquement proches dans la topologie `Internet`), et la *mobilité globale* (mobilité entre sous-réseaux topologiquement éloignés dans la topologie `Internet`). Ensuite, nous définissons une nouvelle terminologie adaptée à la problématique des *réseaux mobiles*.

Chapitre 2: L'Internet et la Suite de Protocoles TCP/IP Dans ce chapitre, nous présentons dans un premier temps l'ensemble des protocoles qui permettent la communication entre deux machines distantes. Tout d'abord, nous présentons succinctement le modèle `OSI` sur lequel est basé la suite de protocoles `TCP/IP` qui régit les communications de l'`Internet` et nous décrivons évidemment en détails le modèle d'adressage de `TCP/IP` et la couche *réseau*. Nous présentons ensuite les protocoles qu'il nous est nécessaire de connaître dans la suite de cette étude, c'est-à-dire `IPv6`, les protocoles de routage point à point, et surtout les protocoles de routage multipoint puisque nous allons largement y faire référence dans les prochains chapitres de cette étude. Nous abordons ensuite le problème de la mobilité. Le problème vient en fait essentiellement du modèle d'adressage de `TCP/IP` qui confond le rôle d'identifiant d'interface de l'adresse `IP`, et son rôle d'identifiant de la localisation dans la topologie `Internet` qui est hiérarchisée. Nous expliquons donc quel est l'impact de la mobilité sur le modèle d'adressage. En conclusion de ce chapitre, nous dressons une liste de *services* qui doivent être fournis par les mécanismes de support de la mobilité.

Chapitre 3: État de l'Art du Support de la Mobilité Dans ce chapitre, nous présentons un certain nombre de propositions pour le support de la mobilité des *stations mobiles*. Les propositions sont présentées séparément bien qu'elles aient des points communs tel que nous allons le voir dans le chapitre suivant. En premier lieu, nous présentons les standards actuellement développés à l'`IETF`, c'est-à-dire `Mobile IPv6` et `Hierarchical Mobile IPv6`. Suivent d'autres propositions qui servirent de base au développement des standards actuels ou qui proposent une approche radicalement différente.

Chapitre 4: Taxinomie Dans ce chapitre, nous dressons une analyse comparative et une taxinomie des propositions étudiées dans le chapitre précédent. Dans un premier temps, nous définissons un modèle d'abstraction qui permet d'extraire les différents composants de chaque proposition, et la fonction de chacun de leur composants. Nous définissons ensuite un modèle d'abstraction qui nous permet de classer les propositions en un nombre limité d'architectures. Après avoir présenté le modèle de Bhagwat [Bhagwat et al., 1996] qui est quelque peu trop restrictif pour extraire l'ensemble des fonctions et leur localisation, nous définissons notre propre modèle. Les propositions se divisent en deux catégories. La première est une catégorie qui fait essentiellement usage de services offerts par le réseau et qui ne nécessite pas de changement d'adresse de la part de la *station mobile*. Deux architectures tombent dans cette catégorie: l'architecture basée sur le routage, et l'architecture basée sur l'inondation. La deuxième catégorie fait usage de deux adresses, l'une utilisée en tant qu'identifiant d'interface, et la deuxième utilisée en tant qu'identifiant de la position dans la topologie. Plusieurs architectures tombent dans cette catégorie: l'architecture *répertoire de position*, l'architecture *Agent Mère*, l'architecture *Hiérarchique*, l'architecture *Réseau Virtuel*, et l'architecture *Multipoint*. En fait, la plupart des propositions peuvent être classées dans plus d'une architecture simultanément et l'architecture *Hiérarchique* semble être nécessaire pour différencier la *mobilité locale* de la *mobilité globale*. L'architecture *multipoint* semble pouvoir être d'un grand intérêt mais nécessite d'importants progrès dans ce domaine particulier.

Seconde Partie: Le Support des Réseaux Mobiles

Chapitre 5: Définition de la Terminologie et de la Problématique Dans ce chapitre, nous définissons les caractéristiques d'un *réseau mobile* et la problématique spécifique résultant de leur mobilité. Nous faisons usage de la terminologie décrite dans le premier chapitre et définie spécifiquement dans le but de leur étude. En effet, aucune terminologie n'existait auparavant pour différencier un *réseau mobile* d'une *station mobile*, ni pour faire la différence entre le routeur qui connecte le *réseau mobile* à l'`Internet` (*routeur mobile*), d'une station résidant de manière permanente dans le *réseau mobile* (*station fixe locale*), d'une station appartenant au *réseau mobile* mais se déplaçant dans ou au-dehors du *réseau mobile* (*station mobile locale*), et d'une station n'appartenant pas au *réseau mobile* et qui s'y attache (*station mobile étrangère* Suite à la définition de cette terminologie, des caractéristique d'un *réseau mobile*, et de la problématique, nous nous attardons à définir un contexte de travail pour la suite de cette étude et un nombre de points que les solutions futures doivent considérer. Nous concluons ce chapitre par un certain nombre d'approches potentielles qui découlent plus ou moins des conclusions tirées dans le chapitre précédent. Parmi l'ensemble des approches envisagées, nous concluons que la capacité de `Mobile IPv6,` le standard de l'`IETF` pour le support des *stations mobiles*, doit être évalué en priorité. L'essentiel de ce chapitre a fait l'objet d'un rapport technique destiné au groupe de travail `Mobile IP` de l'`IETF` (*internet-draft*) et sert de document de base à la suite des travaux dans ce domaine.

Chapitre 6: Les Faiblesses de Mobile IPv6 Étant donnée la conclusion du chapitre précédent, le présent chapitre étudie les avantages et les limitations de `Mobile IPv6` pour le support des *réseaux mobiles*. Le prédécesseur de `Mobile IPv6` fait état d'une solution pour supporter des *sous-réseaux mobiles*, mais de manière très succincte et sans considérer les problèmes mis en avant dans le chapitre précédent. L'idée de base est que le *routeur mobile* fasse usage de `Mobile IPv6` comme une *station mobile* et obtienne une adresse temporaire sur son nouveau point d'ancrage, en plus de son adresse permanente. Les paquets destinés aux stations dans le *réseau mobile* peuvent alors être envoyés via cette adresse temporaire. `Mobile IPv6` ne fait plus mention de cette possibilité, aussi nous attardons nous à vérifier si les *réseaux mobiles* peuvent néanmoins être considérés comme dans `Mobile IPv4.` Cette étude démontre que la spécification actuelle de `Mobile IPv6` ne le permet pas et que cela pose un certain nombre de problèmes, notamment de sécurité, et de passage l'échelle (nombre de *réseaux mobiles*, taille des *réseaux mobiles*, et de nombre de stations extérieures communiquant avec des stations du *réseau mobile*).

Chapitre 7: Propositions d'Extensions de Mobile IPv6 Dans ce chapitre, nous proposons un certain nombre d'extensions qui traitent des limitations de `Mobile IPv6` mises à jour dans le chapitre précédent.

Dans un premier temps, nous proposons une extension qui permet d'établir une relation entre le préfixe d'adresse `IP` commun à toutes les stations résidant dans le *réseau mobile* et l'adresse temporaire obtenue par le *routeur mobile* à chacun de ses points d'ancrage. Cette extension permet tout d'abord de rediriger des paquets destinés aux stations résidants dans le *réseau mobile*. Elle permet également un routage optimal en informant toutes les stations communiquant avec le *réseau mobile* que les paquets destinés à n'importe quelle station résidant dans le *réseau mobile* doivent être redirigés via l'adresse temporaire du *routeur mobile*.

Pour le long terme, nous considérons les problèmes de passage à l'échelle. Nous proposons de réduire le coût des messages de contrôle induit par `Mobile IPv6` pour faire parvenir l'adresse temporaire du *routeur mobile* à tous les correspondants du *réseau mobile* au moyen de deux techniques multipoint. La première, dite traditionnelle, établit un arbre de distribution entre le *routeur mobile* et le groupe multipoint constitué par l'ensemble des correspondants du *réseau mobile*. La deuxième repose sur une nouvelle technique de routage multipoint qui enregistre directement la liste des correspondants dans le message de contrôle.

Nous concluons ce chapitre par une vue d'ensemble d'une nouvelle architecture de gestion de la mobilité rassemblant diverses techniques rencontrées lors de l'étude de l'État de l'Art ainsi que nos extensions multipoint.

Troisième Partie: Évaluation des Performances

Chapitre 8: Procédé de Simulation Dans ce chapitre, nous décrivons notre méthode d'évaluation de la performance de nos propositions. Dans un premier temps, nous décrivons nos besoins puis notre plateforme de simulation basée sur `NS-2.1b6` et nos extensions. Ensuite, nous présentons notre scénario de simulation. Les topologies utilisées contiennent plusieurs centaines de noeuds répartis en une centaine de sites. Le mobile se déplace dans un certain nombre de sites et communique avec un nombre croissant de correspondants. La dernière partie de ce chapitre présente les métriques utilisées lors de l'analyse des résultats.

Chapitre 9: Analyse Nous avons regroupé tous nos résultats de simulations dans ce chapitre pour faciliter leur présentation et leur interprétation. Nos simulations ne traitent que de la *mobilité globale*, ce qui signifie que l'ensemble de nos résultats est applicable à la fois à des extensions portées sous `Mobile IPv6` et à des extensions portées sous `Hierarchical Mobile IPv6`.

Notre première expérience vise à mettre en avant notre scénario de mobilité et à démontrer que le déplacement d'un mobile dans des parties topologiquement éloignées d'un réseau n'a pas d'impact sur le nombre moyen de liens entre le mobile et chacun de ces correspondants. Notre deuxième expérience démontre que l'optimisation de routage dans `Mobile IPv6` est nécessaire pour réduire la charge du réseau, en particulier lorsque le nombre de correspondants croît, et démontre également que les messages de contrôle ont un coût important, notamment sur le lien sans fil. L'expérience qui suit compare l'envoi des messages de contrôles de `Mobile IPv6` de manière point-à-point et de manière multipoint, avec les protocoles multipoint traditionnels. Nous démontrons que l'usage d'une technique multipoint est avantageuse, à condition que le coût induit par l'usage de la technique multipoint elle-même soit réduit, et qu'un arbre de type *arbre centré sur le noyau* est plus avantageuse qu'un *arbre des plus courts chemins*. L'évaluation de la deuxième technique multipoint montre clairement que celle-ci est bénéfique à partir d'un nombre très faible de correspondants, à condition que les routeurs capables de traiter la liste des correspondants contenus dans les paquets soient suffisamment nombreux et bien placés.

Conclusion et Perspectives

Dans cette dissertation, nous avons étudié le support des *réseaux mobiles*, au contraire des travaux habituels qui traitent du support des *stations mobiles*.

Contribution Dans le chapitre 2, nous avons mis en avant que le problème causé par la mobilité est essentiellement un problème d'adressage et nous avons défendu le besoin d'un identifiant permanent ne dépendant pas de la position du mobile dans la topologie `Internet`. Notre étude des différentes propositions pour la gestion des *stations mobiles* (chapitre 3) et leur classification dans une taxinomie (chapitre 4) a démontré que ce problème d'adressage peut être résolu de plusieurs manières. Le moyen le plus répandu fait usage de deux adresses, l'une utilisée en tant qu'identifiant d'interface, et la deuxième utilisée en tant qu'identifiant de la position dans la topologie. C'est aussi l'approche qui nous semble la plus évidente et la plus simple car elle ne remet pas en cause la pérennité des différents protocoles de la couche réseau et ne nécessite pas de réelles modifications dans le réseau.

Bien que parfois mentionné dans la littérature, le support des *réseaux mobiles* est un sujet nouveau. Aucune terminologie spécifique n'a pour l'instant été établie, ni leurs caractéristiques et leurs problèmes particuliers. Nous comblons ce manque dans le chapitre 5, dans lequel nous définissons également un cahier des charges pour les solutions futures. Faisant suite à l'étude des propositions pour la gestion des *stations mobiles*, et en prenant en compte les caractéristiques propres aux *réseaux mobiles*, nous avons envisagé un certain nombre d'approches pour supporter ces derniers.

Nous concluons qu'aucune approche ne peut offrir une solution simple et immédiate à ce problème, mis à part peut-être une approche basée sur `Mobile IPv6,` le standard de l'`IETF` pour le support des *stations mobiles*. Conséquemment, nous avons étudié dans quelles limites `Mobile IPv6` peut-être adapté au support des *réseaux mobiles* (chapitre 6). Outre un certain nombre de problèmes qui nécessiterait simplement la clarification de la spécification `Mobile IPv6,` nous avons également relevé des problèmes de sécurité, et un problème concernant le passage à l'échelle pour les *réseaux mobiles* ayant un nombre important de correspondants. Pour résoudre ces problèmes, nous proposons trois groupes d'extensions (chapitre 7). Le premier groupe d'extensions a pour but d'établir une connexion permanente entre les stations situées dans le *réseau mobile* et l'`Internet` et de permettre un routage optimal depuis leurs correspondants. Ces propositions ont également fait l'objet d'une publication à l'`IETF` dans le groupe de travail `Mobile IP`. Les deux groupes d'extensions suivants ont pour but de résoudre le problème de passage à l'échelle. A cette fin, nous faisons usage de deux techniques multipoint. La première, dite traditionnelle, établit un arbre multipoint entre le *routeur mobile* et tous les correspondants du *réseau mobile* en vue d'y transmettre les paquets de contrôle permettant le routage optimal. La deuxième, plus novatrice et plutôt destinée à un nombre assez restreint de correspondants consiste à enregistrer la liste des correspondants directement dans le paquet de contrôle, ce qui évite l'établissement de l'arbre, généralement coûteuse, mais limite aussi le nombre de correspondants dû à l'accroissement de la taille du paquet.

La performance de nos propositions multipoint est évaluée par simulation en faisant usage du simulateur `NS-2` (chapitre 8) et montre que nos propositions diminuent très fortement la proportion de bande passante consommée par les messages de contrôle dans l'`Internet` (en particulier le lien sans fil par lequel le *routeur mobile* s'attache généralement à l'`Internet`). Nos simulations démontrent également que la deuxième technique multipoint (`List-Based Multicast`), ne nécessite pas d'être déployée dans l'ensemble de l'`Internet`, mais en revanche nécessite de bien choisir les routeurs capables de traiter les paquets contenant une liste de destinataires (chapitre 9).

Perspectives et travaux futurs Afin de permettre la standardisation de notre premier groupe de propositions, il sera nécessaire de traiter les problèmes de sécurité plus en profondeur, notamment les problèmes d'autorisation et d'authentification. Nos solutions basées sur les techniques multipoint, quant à elles, font aussi face à des problèmes ouverts liés à la sécurité et nécessitent de suivre activement les travaux faits dans ce sens à l'IETF. L'usage des protocoles de routage multipoint traditionnels, en particulier, est plus problématique et nécessite de sélectionner un protocole suffisamment peu coûteux pour ne pas perdre le bénéfice de la diminution du coût des messages de contrôle des extensions portées à `Mobile IPv6.` En l'occurrence, le protocole candidat doit permettre la transmission multipoint entre domaines administratifs, ce qui est un problème activement débattu à l'`IETF` en ce moment. Lorsque les travaux de l'`IETF` auront suffisamment avancé dans ce sens, nous pourrons conduire de plus amples simulations. Nous devons également conduire d'autres simulations de la technique `List-Based Multicast` sous d'autres scénarios et d'autres configurations.

Nos solutions n'ont pas la prétention de résoudre tous les problèmes définis dans le chapitre 5, bien au contraire. Un grand nombre de problèmes subsistent, et nous souhaitons particulièrement travailler sur les aspects de la communication entre deux *stations mobiles* toutes les deux situées dans un *réseau mobile*, sur les interactions entre le *routeur mobile* et le protocole de routage du réseau visité, sur d'autres aspects de sécurité, etc. Nous souhaitons également approfondir les autres approches que nous avons dans un premier temps envisagées, pour faire la place à l'étude de solutions basées sur `Mobile IPv6`. Nous souhaitons aussi détailler une architecture plus générale de la gestion de la mobilité basée à la fois sur nos techniques multipoint et d'autres techniques identifiées lors de l'étude des propositions pour le support des *stations mobiles*. En particulier, nous pensons qu'une technique hiérarchique combinée avec la distribution multipoint des messages de contrôle dans un réseau multipoint virtuel pourrait apporter une solution d'ensemble à la mobilité.

Avec un regard un peu plus général, nous pouvons dire que notre étude a ouvert une brèche dans la gestion habituelle de la mobilité et mis à jour de nouveaux problèmes qui nécessiteront de plus amples travaux de recherche. En effet, le besoin de déplacer des réseaux est perçu depuis un certain nombre d'années, mais aucun travail particulier ne leur avait pour le moment été consacré. Découlant de cette constatation, les problèmes qui leur sont propres n'avaient donc pour le moment pas été abordés. Notre contribution à l'`IETF` a eu pour résultat une certaine prise de conscience, assez faible dans un premier temps, mais grandissante dans les rangs de la communauté `Internet`. Le support des *réseaux mobiles* est en effet un sujet qui intéresse de nombreux industriels. Nul doute que nombre d'entre eux travaillent d'ailleurs déjà sur le sujet; quelques informations dispersées au compte-gouttes le laisse penser. Nous sommes d'ailleurs sur le point, avec le soutien de quelques industriels, de créer un nouveau groupe de travail à l'`IETF` pour traiter spécifiquement le cas des *réseaux mobiles*.

Bibliography

K. C. Almeroth. The Evolution of Multicast: From the MBone to Interdomain Multicast to Internet2 Deployment. *IEEE Network - Magazine of Global Information Exchange*, 14(1), January 2000.

B. Awerbuch and D. Peleg. Concurrent Online Tracking of Mobile Users. In *Proc. ACM SIGCOMM*, pages 221–33, Zurich, Switzerland, 1991. M.I.T, MA, USA and Weizemann Institute, Israel.

T. Ballardie, J. Crowcroft, C. Diot, C. Lee, R. Perlman, and Z. Wang. On Extending the Standard IP Multicast Architecture. Research Report RN/99/21, University College London, October 1999.

T. Ballardie, P. Francis, and J. Crowcroft. Core Based Trees (CBT) An Architecture for Scalable Inter-Domain Multicast Routing. *IEEE Network - Magazine of Global Information Exchange*, 23(4):85–95, 1993.

T. Bates, R. Chandra, D. Katz, and Y. Rekhter. Multiprotocol Extensions to BGP-4. Request For Comments 2283, IETF, February 1998.

P. Bhagwat, C. Perkins, and S. Tripathi. Network Layer Mobility: an Architecture and Survey. *IEEE Personal Communications*, 3(3):54–64, June 1996.

P. Bhagwat, S. Tripathi, and C. Perkins. Network Layer Mobility: an Architecture and Survey. Technical Report CS-TR-3570, UMIACS-TR-95-117, Computer Science Department and IBM, T.J. Watson Research Center, Sept 1995.

L. Blazevic and J.-Y. Le Boudec. Distributed Core Multicast (DCM): A Multicast Routing Protocol for Many Groups with Few Receivers. *SIGCOMM - Computer Communication Review*, 29(5), 1999.

L. Blazevic and J.-Y. Le Boudec. Distributed Core Multicast (DCM). Internet Draft draft-blazevic-dcm-mobility-01.txt, IETF, June 2000.

R. Boivie. A New Multicast Scheme for Small Groups. Technical report, IBM Research Report, 1999.

R. Boivie, N. Feldman, Y. Imai, W. Livens, D. Ooms, and O. Paridaens. Explicit Multicast Xcast Basic Specification. Internet Draft draft-ooms-xcast-basic-spec-02.txt, IETF, October 2001. Work in progress.

R. Boivie, N. Feldman, and C. Metz. Small Group Multicast. Internet Draft draft-boivie-sgm-01.txt, IETF, July 2000a. Work in progress.

R. Boivie, N. Feldman, and C. Metz. Small Group Multicast: A New Solution for Multicasting on the Internet. *IEEE Internet Computing*, 4(3):75–79, 2000b.

J. Bound and C. Perkins. Dynamic Host Configuration Protocol for IPv6 (DHCPv6). Internet Draft draft-ietf-dhc-dhcpv6-14.txt, IETF, February 1999. Work in Progress.

R. Braden. Requirements for Internet Hosts - Communication Layers. Request For Comments 1122, IETF, October 1989.

T. Braun. Multicast for Small Conferences. Technical Report IAM-00-008, University of Berne, Switzerland, July 2000. http://www.iam.unibe.ch/ rvs/publications.

R. Caceres and V. N. Padmanabhan. Fast and Scalable Handoffs for Wireles Internetworks. In *Proc. of the Second ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom)*, Rye, New York, USA, November 1996. Bell Laboratories and University of California at Berkeley.

K. Calvert, M. Doar, and E. W. Zegura. Modeling Internet Topology. *IEEE Communications Magazine*, 35 (6):160–163, June 1997.

K. Calvert and E. Zeggura. GT-ITM: Georgia Tech Internetwork Topology Models. Technical report, Georgia Institute of Technology, College of Computing, 1996. http://www.cc.gatech.edu/projects/gtitm/.

A. Campbell and J. Gomez. IP Micro-mobility Protocols. *ACM SIGMOBILE Mobile Computing and Communication Review (MC2R)*, April 2001.

A. Campbell, J. Gomez, S. Kim, Z. Turanyi, C.-Y. Wan, and A. Valko. Comparison of IP Micro-Mobility Protocols. *IEEE Wireless Communications Magazine*, 9(1), February 2002.

A. Campbell, J. Gomez, and A. G. Valko. An Overview of Cellular IP. In *IEEE Wireless Communications and Networking Conference (WCNC)*, September 1999a.

A. Campbell, J. Gomez, C.-Y. Wan, Z. Turanyi, and A. Valko. Cellular IP. Internet Draft draft-ietf-mobileip-cellularip-00,txt, IETF, December 1999b. Expired. Available http://www.comet.columbia.edu/cellularip/.

K. G. Carlberg. A Routing Architecture that supports Mobile End Systems. In *IEEE Proceedings of the Military Communications Conference (MILCOM'92)*, pages 159–164, 1992.

C. Castelluccia. A Hierarchical Mobility Management Scheme for the Internet. Technical Report CSL-TR-97-736, Stanford University, California, September 1997.

C. Castelluccia. A Hierarchical Mobile IPv6 Proposal. Technical Report 226, Institut National de la Recherche en Informatique et en Automatique (INRIA), November 1998a.

C. Castelluccia. A Hierarchical Mobility Management Scheme for IPv6. In *Third Symposium on Computers and Communications (ISCC'98)*, June 1998b.

C. Castelluccia. Toward a Hierarchical Mobile IPv6. In *Eighth IFIP International Conference on High Performance Networking (HPN'98)*, Vienna, Austria, September 1998c.

C. Castelluccia. A Hierarchical Mobile IPv6 Proposal. In *AMOS ACTS Mobile Summit*, Sorento, Italy, June 1999. INRIA.

C. Castelluccia. HMIPv6: A Hierarchical Mobile IPv6 Proposal. *ACM Mobile Computing and Communication Review (MC2R)*, April 2000.

C. Castelluccia and L. Bellier. Toward a Unified Hierarchical Mobility Management Framework. Internet Draft draft-castelluccia-uhmm-framework-00.txt, IETF, June 1999. Work in progress.

C. Castelluccia and D. Jacquemin. CBTM: A Core Based Tree Mobility Management Scheme for IPv6. Internal report, INRIA, 1998.

S. Cheshire and M. G. Baker. Internet mobility 4x4. In *Proceedings of the ACM SIGCOMM'96 Conference*, pages 318–329, Stanford University, CA, August 1996.

G. Cho. A Location Management Scheme Supporting Route Optimization for Mobile Hosts. *Journal of Network and Systems Management*, 6(1):31–50, 1998.

G. Cho and L. F. Marshall. An Efficient Location and Routing Scheme for Mobile Computing Environments. *IEEE Journal on Selected Areas in Communications*, 13(5), June 1995.

G. Cizault. *"IPv6"*. Editions O'Reilly, 2nd edition, 1999.

D. Cohen, J. B. Postel, and R. Rom. IP Addressing and Routing in a Local Wireless Network. Technical report, University of Southern California, July 1991.

COMET. Web page at Columbia University. http://www.comet.columbia.edu.

A. Conta and S. Deering. Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification. Request For Comments 2463, IETF, December 1998.

M. Crawford, A. Mankin, T. Narten, J. W. Stewart, and L. Zhang. Separating Identifiers and Locators in Addresses: An Analysis of the GSE Proposal for IPv6. Internet Draft draft-ietf-ipngwg-esd-analysis-03.txt, IETF, November 1998.

S. Deering. *Multicast Routing in Datagram Internetwork*. PhD thesis, Stanford University, US, December 1991.

S. Deering and D. Cheriton. Multicast Routing in Datagram Internetworks and Extended LANs. *ACM Transactions on Computer Systems*, pages 85–111, May 1990.

S. Deering, D. Estrin, D. Farinacci, V. Jacobson, C.-G. Liu, and W. Liming. An Architecture for Wide-Area Multicast Routing. *ACM SIGCOMM Computer Communication Review*, 24(4):126–135, August 1994.

S. Deering, W. Fenner, and B. Haberman. Multicast Listener Discovery (MLD) for IPv6. Internet Draft draft-ietf-ipngwg-mld-02.txt, IETF, June 1999. Work in Progress.

S. Deering and R. Hinden. Internet Protocol Version 6 (IPv6) Specification. Request For Comments 2460, IETF, December 1998.

M. Diagne and T. Noël. A Protocol for IPv6 Nomadic Communications. In *5th IEEE Malaysia International Conference on Communications (MICC)*, October 2001a.

M. Diagne and T. Noël. IPv6 Unicast Protocol for Small Group Communications. In *9th IEEE International Conference on Software, Telecommunications annd Computer Networks (Softcom)*, October 2001b.

T. Ernst. The Mobile Next Generation Internet. In *5th Cabernet Radicals Workshop*, Valadares, Portugal, July 1998.

T. Ernst. Extending Mobile IPv6 with Multicast to support Mobile Networks in IPv6. In *IP based Cellular Network conference (IPCN)*, Paris La Defense, France, May 2000. Upper Side.

T. Ernst. Mobile Routers in IPv6. In *IP based Cellular Network conference (IPCN)*, Paris, France, May 2001a. Upper Side.

T. Ernst. MobiWan: NS-2 extensions to study mobility in Wide-Area IPv6 Networks. Technical report, MOTOROLA Labs and INRIA Rhône-Alpes, 2001b. http://www.inrialpes.fr/planete/pub/mobiwan/.

T. Ernst, L. Bellier, C. Castelluccia, and H.-Y. Lach. Mobile Networks Support in Mobile IPv6. Internet Draft draft-ernst-mobileip-v6-network-01.txt, IETF, November 2000a. Work in progress.

T. Ernst, L. Bellier, C. Castelluccia, and H.-Y. Lach. Mobile Networks Support in Mobile IPv6. Internet Draft draft-ernst-mobileip-v6-network-00.txt, IETF, July 2000b. Work in progress.

T. Ernst, L. Bellier, A. Olivereau, C. Castelluccia, and H.-Y. Lach. Mobile Networks Support in Mobile IPv6. Internet Draft draft-ernst-mobileip-v6-network-02.txt, IETF, June 2001a. Work in progress.

T. Ernst and C. Castelluccia. Mobility Management in Communications Networks. European Patent 01401241.3-, Motorola Labs, May 2001.

T. Ernst, C. Castelluccia, and H.-Y. Lach. Les Réseaux Mobiles dans IPv6. In *13ème congrès DNAC (De Nouvelles Architectures pour les Communications)*, Ministère chargé des Télécoms, Paris, France, December 1999.

T. Ernst, C. Castelluccia, and H.-Y. Lach. Extending Mobile IPv6 with Multicast to Support Mobile Networks in IPv6. In *1st European Conference on Universal Multiservice Networks (ECUMN)*, Colmar, France, October 2000c.

T. Ernst and H.-Y. Lach. Communication System and Method Therefor. European Patent 00401441.1-2216, Motorola Labs, May 2000.

T. Ernst, H.-Y. Lach, and C. Castelluccia. Network Mobility Support in IPv6: Problem Statement and Requirements. Internet Draft draft-ernst-mobileip-monetv6-00.txt, IETF, July 2001b. Work in progress.

D. Estrin, D. Farinacci, D. Thaler, S. Deering, M. Handley, V. Jacobson, C. Liu, P. Sharma, and L. Wei. Protocol Independant Multicast-Sparse Mode (PIM-SM): Protocol Specification. Request For Comments 2362, IETF, June 1998.

K. Fall and K. Varadhan. NS notes and documentation. Technical report, The Vint Project, UC Berkeley, LBL, USC/ISI, Xerox PARC, 2000. "http://www.isi.edu/nsnam/ns/index.html".

C. Hedrick. Routing Information Protocol - RIP. Request for Comments 1058, IETF, June 1988.

A. Helmy. A Multicast-based Protocol for IP Mobility Support. In *Proc. of the 2nd International Workshop on Networked Group Communication (NGC)*, Stanford University, Palo Alto, California, USA, November 2000.

R. Hinden and S. Deering. IP Version 6 Addressing Architecture. Request For Comments 2373, IETF, July 1998.

R. Hinden, M. O'Dell, and S. Deering. An IPv6 Aggregatable Global Unicast Address Format. Request For Comments 2374, IETF, July 1998. Work in progress.

H. Holbrook and D. Cheriton. IP multicast channels: EXPRESS support for large-scale singlle-source applications. In *ACM Sigcomm*, Cambridge, Massachusetts, USA, August 1999.

Y.-C. Hu and D. B. Johnson. Caching Strategies in On-Demand Routing Protocols for Wireless Ad Hoc Networks. In *Proc. of the Sixth ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom)*, Boston, MA, USA, August 2000. Carnegie Mellon University.

C. Huitema. *Routing in the Internet*. Prentice Hall, 1995.

C. Huitema. *IPv6 The New Internet Protocol*. Prentice Hall, 2nd edition, 1998.

IEEE. Guidelines for 64-bit Global Identifier (EUI-64) Registration Authority. Technical report, IEEE, March 1997. http://standards.ieee.org/regauth/oui/tutorials/EUI64.html.

Y. Imai. Multiple Destination option on IPv6(MDO6). Internet Draft draft-imai-mdo6-02.txt, IETF, September 2000. Work in progress.

J. Ioannidis. *Protocols for Mobile Internetworking*. PhD thesis, Columbia University, 1993.

J. Ioannidis, D. Duchamp, and G. Q. Maguire Jr. IP-based Protocols for Mobile Internetworking. In *Proc. ACM SIGCOMM*, pages 233–45. Department of Computer Science, Columbia University, September 1991.

J. Ioannidis and G. Q. Maguire Jr. The Design and Implementation of a Mobile Internetworking Architecture. In *Proc. Winter USENIX*, pages 491–502, San Diego, January 1993.

M. Ishiyama, M. Kunishi, K. Uehara, H. Esaki, and F. Teraoka. LINA: A New Approach to Mobility Support in Wide Area Networks. *IEICE Transactions on Communications*, E84-B(8), August 2001.

ISO. Data Processing - Open Systems Interconnection - Basic Reference Model. Technical Report ISO IS 7498-1984, International Organisation for Standardisation, 1984.

D. B. Johnson. Mobile Host Internetworking Using IP Loose Source Routing. Technical Report CMU-CS-93-128, Carnegie Mellon University, Pittsburgh, PA, February 1993.

D. B. Johnson. Scalable Support for Transparent Mobile Host Internetworking. *Wireless Networks, special issue on Recent Advances in Wireless Networking Technology, ACM and Baltzer Science Publishers*, 1 (3):311–321, October 1995.

D. B. Johnson and C. Perkins. Mobility Support in IPv6. Internet Draft draft-ietf-mobileip-ipv6-13.txt, IETF, November 2000. Work in progress.

U. Jonsson, F. Alriksson, T. Larsson, P. Johansson, and G. Maguire. MIPMANET - Mobile IP for Mobile Ad-Hoc Networks. In *MobiHOC*, Boston, Massachusetts, USA, August 2000.

S. Kent and R. Atkinson. IP Authentication Header. Request For Comments 2402, IETF, November 1998a.

S. Kent and R. Atkinson. IP Encapsulating Security Payload (ESP). Request For Comments 2406, IETF, November 1998b.

G. Kirby. Locating the User. *Communication International*, October 1995.

S. Kumar, P. Radoslavov, D. Thaler, C. Alaettinoglu, D. Estrin, and M. Handley. The MASM/BGMP Architecture for inter-domain multicast routing. In *ACM Sigcomm*, Vancouver, Canada, August 1998.

D. C. Lee, D. L. Lough, S. F. Midkiff, N. J. Davis, and P. E. Benchoff. The Next Generation of the Internet: Aspects of the Internet Protocol Version 6. *IEEE Network*, January 1998. Virginia Polytechnic and State University.

J. Leyden. Get online from low-earth-orbiting research craft. *The register*, October 2001. http://www.theregister.co.uk/content/5/22556.html.

G. Malkin. RIP version 2 - Carrying Additional Information. Request For Comments 1723, IETF, November 1994.

A. Medina, A. Lakhina, I. Matta, and J. Byers. BRITE: An Approach to Universal Topology Generation. In *Proc. International Workshop on Modeling, Analysis and Simulation of Computer and Telecommunications Systems (MASCOTS)*, 2001.

P. Mockapetris. Domain Names - Concepts and Facilities. Request for Comments 1034, IETF, November 1987a.

P. Mockapetris. Domain Names - Implementation and Specification. Request for Comments 1035, IETF, November 1987b.

J. Moy. Multicast Routing Extensions for OSPF. *Communications of the ACM*, 37(8):61–66, August 1994.

J. Moy. Open Shortest Path First - OSPF version 2. Request For Comments 2178, IETF, July 1997.

A. Myles, D. B. Johnson, and C. Perkins. A Mobile Host Protocol Supporting Route Optimization and Authentication. *IEEE Journal on Selected Areas in Communications, special issue on Mobile and Wireless Computing Networks*, 13(5):839–849, June 1995.

A. Myles and D. Skellern. Comparing Four IP Based Mobile Host Protocols. In *Joint-European Networking Conference*, pages 191–196, Trondheim, Norway, May 1993a. Macquarie University, Sydney, Australia.

A. Myles and D. Skellern. Comparison of Mobile Host Protocols for IP. *Journal of Internetworking Research and Experience*, pages 175–194, December 1993b.

J. Mysore and V. Bharghavan. A New-Multicasting-based Architecture for Internet Host Mobility. In *Proc. of the Third ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom)*, pages 161–172, Budapest, Hungary, September 1997. University of Illinois at Urbana-Champaign.

T. Narten, E. Nordmark, and W. Simpson. Neighbor Discovery for IP version 6 (IPv6). Request For Comments 2461, IETF, December 1998.

C. Navas, Julio and T. Imielinski. Geographic Addressing and Routing. In *Proc. of the ThirdACM/IEEE International Conference on Mobile Computing and Networking (MobiCom)*, Budapest, Hungary, September 1997. University of Rutgers, New Jersey, USA.

T. Noël. *Une Architecture d'Adressage et de Routage Logiques pour les Mobiles*. PhD thesis, Université Louis Pasteur, Strasbourg, France, Juin 1998.

T. Noël, D. Grad, and J. Pansiot. Logical Framework for Group Communications with Mobile Hosts. In *AFRICOM-CCDC*, Tunis, Tunisie, Octobre 1998.

T. Noël, D. Grad, and J. Pansiot. Communications Multipoint pour les Mobiles. *TSI*, 6, Juin 2001.

M. O'Dell. GSE - An Alternate Addressing Architecture for IPv6. Internet Draft draft-ietf-ipngwg-gseaddr-00.txt, IETF, March 1998.

D. Ooms. Taxonomy of xcast/sgm proposals. Internet Draft draft-ooms-xcast-taxonomy-00.txt, IETF, July 2000. Work in progress.

D. Ooms, W. Livens, and O. Paridaens. Connectionless Multicast. Internet-Draft draft-ooms-cl-multicast-02.txt, IETF, April 2000. Work in Progress.

V. Padmanabhan N. and R. Katz. Using DNS to Support Host Mobility. Slides of the presentation made at ILP, The Daedalus Group, University of California at Berkeley, March 1998.

C. Partridge. Technical Criteria for Choosing IP the Next Generation (IPng). Request For Comments 1726, IETF, December 1994.

C. Perkins. IP Mobility Support. Request For Comments 2002, IETF, October 1996a.

C. Perkins. Mobile-IP Local Registration with Hierarchical Foreign Agents. Internet Draft draft-perkins-mobileip-hierfa-00.txt, I.B.M, February 1996b.

C. Perkins and D. B. Johnson. Mobility Support in IPv6. In *Proc. of the SecondACM/IEEE International Conference on Mobile Computing and Networking (MobiCom)*, pages 27–37, Rye, New-York, USA, November 1996. IBM Corporation and Carnegie Mellon University.

C. Perkins and D. B. Johnson. Route Optimization in Mobile IP. IETF Internet Draft draft-ietf-mobileip-optim-09.txt, Sun Microsystems and Carnegie Mellon University, February 2000. Work in progress.

C. E. Perkins. *Mobile IP, Design Principles and Practices*. Wireless Communications Series. Addison-Wesley, 1998. ISBN 0-201-63469-4.

R. Perlman, C.-Y. Lee, J. Crowcroft, Z. Wang, C. Diot, J. Thoo, and M. Green. Simple multicast: A design for simple, low-overhead multicast. Internet-Draft draft-perlman-simple-multicast.txt, IETF, February 1999. Work in progress.

PLANETE. Web page at INRIA. http://www.inrialpes.fr/planete.

Postel. Internet Protocol DARPA Internet Program Protocol Specification. Request For Comments 791, IETF, September 1981a.

Postel. Transmission Control Protocol DARPA Internet Program Protocol Specification. Request For Comments 793, IETF, September 1981b.

T. Quinot. An IPv6 architecture for Aeronautical Telecommunication Network. Master's thesis, Ecole Nationale Supérieure des Télécommunications Paris, EUROCONTROL - European Organization for the Safety of Air Navigation - ISA project (IPv6, Satellite communication and ATMode for ATN), 1998. http://www.eurocontrol.fr/.

M. Ramalho. Intra- and Inter-Domain Multicast Routing Protocols: A Survey and Taxonomy. *IEEE Communications Surveys and Tutorials*, 3(1), 2000.

R. Ramjee, T. La Porta, S. Thuel, K. Varadhan, and L. Salgarelli. IP micro-mobility support using HAWAII. Technical report, IETF, March 1999a. Work in Progress.

R. Ramjee, T. La Porta, S. Thuel, K. Varadhan, and S. Wang. HAWAII: A Domain-Based Approach for Supporting Mobility in Wide-Area Wireless Networks. In *IEEE International Conference on Network Protocols*, 1999b.

Y. Rekhter and P. Gross. Application of the Border Gateway Protocol in the Internet. Request for Comments 1772, IETF, March 1995.

Y. Rekhter and T. Li. A Border Gateway Protocol 4 - (BGP-4). Request For Comments 1771, IETF, March 1995.

P. Roberts and J. Loughney. Local Subnet Mobility Problem Statement. Internet Draft draft-proberts-local-subnet-mobility-problem-02.txt, IETF, November 2001. Work in progress.

S. Seshan and R. H. Balakrishnan, H. Kats. Handoffs in Cellular Wireless Networks: The Daedalus Implementation and Experience. *ACM/Balzer Journal on Wireless Networks*, 1995.

Z. Shelby, D. Gatzounas, A. Campbell, C.-Y. Wan, Z. Turanyi, and A. Valko. Cellular IPv6. Internet Draft draft-shelby-seamoby-cellularip-00,txt, IETF, November 2000. Expired. Available http://www.comet.columbia.edu/cellularip/.

A. C. Snoeren and H. Balakrishnan. An End-to-End Approach to Host Mobility. In *Proceedings of the Sixth Annual International Conference on Mobile Computing and Networking, MOBICOM 2000*, pages 155–166, 2000.

H. Soliman, C. Castelluccia, K. El-Makki, and L. Bellier. Hierarchical MIPv6 mobility management. Internet Draft draft-ietf-mobileip-hmipv6-03.txt, IETF, February 2001. Work in progress.

J. D. Solomon. *Mobile IP, The Internet Unplugged*. Prentice Hall Series in Computer Networking and Distributed Systems. Prentice Hall PTR, 1998. ISBN 0-13-856246-6.

J. W. Stewart. *"BGP4"*. Addison-Wesley, 1998.

C. Sunshine and J. Postel. Addressing mobile hosts in the ARPA Internet environment. Internet Engineering Note IEN 135, USC-ISI, March 1980.

A. S. Tanenbaum. *Computer Networks - Third Edition*. Prentice-Hall International, Inc, 1996.

F. Teraoka. *A Study on Host Mobility in Wide Area Networks*. Phd dissertation, Keio University, Japan, January 1993.

F. Teraoka, K. Claffy, and M. Tokoro. Design, Implementation and Evaluation of virtual Internet Protocol. In *Proceedings of the 12th International Conference on Distributed Computing Systems*, pages 170–177, Yokohama, Japan, June 1992. Sony CSL.

F. Teraoka and M. Tokoro. Host Migration Transparency in IP Networks: The VIP approach. *ACM Computer Communication Review*, 23(1):45–65, January 1993.

F. Teraoka, K. Uehara, H. Sunahara, and J. Murai. VIP: A Protocol Providing Host Mobility. *Communications of the ACM*, 37(8):67–75, August 1994.

F. Teraoka, Y. Yokote, and M. Tokoro. A Network Architecture Providing Host Migration Transparency. In *Proceedings of SIGCOMM'91 SYMPOSIUM, Communications Architectures and Protocols*, pages 209–220. Sony CSL, September 1991.

S. Thomson and T. Narten. IPv6 Stateless Address Autoconfiguration. Request For Comments 2462, IETF, December 1998.

A. G. Valkó. Cellular IP: A New Approach to Internet Host Mobility. *ACM Computer Communication Review*, January 1999.

P. Vixie, S. Thomson, Y. Rekhter, and J. Bound. Dynamic Updates in the DNS. Request for Comments 2136, IETF, February 1998.

D. Waitzman, C. Partridge, and S. Deering. Distance Vector Multicast Routing Protocol. Request For Comments 1075, IETF, November 1988.

L. Wei and D. Estrin. The Tradeoffs of Multicast Trees and Algorithms. In *Proceedings of the International Conference on Computer Communications and Networks*, San Francisco, September 1994.

L. Wei and D. Estrin. Multicast Routing in Dense and Sparse Modes: Simulation Study of Tradeoffs and Dynamics. Technical Report 95-613, USC, 1995.

E. W. Zegura, K. Calvert, and S. Bhattacharjee. How to Model an Internetwork. In *Proceedings of IEEE Infocom*, San Francisco, CA, 1996.

List of Figures

List of Tables

Index

Network Mobility Support in IPv6

Abstract: This thesis is devoted to the study of *network mobility support* in IPv6, the new generation of Internet Protocol. *Network mobility support*, unlike *host mobility support*, is concerned with situations where an entire network changes its point of attachment in the Internet topology. The purpose is to provide continuous and optimal Internet access to all nodes located in the mobile network. *Network mobility support* must be considered separately from *host mobility support* because it raises a number of new issues concerning the question of addressing, locating and routing. Our first contribution is the definition of a taxonomy that is used to describe all existing *host mobility support* schemes. Our second contribution is the definition of a new terminology, set of issues and set of requirements specifically targeted to *network mobility support*. Among possible approaches, we focus on Mobile IPv6, the IETF *host mobility support* standard and we study its ability to support mobile networks. We consider both short and long-term solutions. We propose extending Mobile IPv6 with new features as an immediate solution to address its shortcomings. As for the long term, we address the question of scalability to a large number of mobile networks communicating with a potentially large number of nodes. We propose the use of two distinct multicast techniques as solutions to reduce signaling incurred by Mobile IPv6. The performance of our multicast extensions is evaluated by simulation. We conclude this dissertation with a prospective architecture framework which combines our multicast extensions with a number of other techniques identified during this study.

Key-words: IPv6 - Network Mobility - Host Mobility - Mobility Support - Mobile Networks - Network in Motion - Mobile Routers - Routing - Multicast - Internet - Simulation - Mobile IPv6

Le Support des Réseaux Mobiles dans IPv6

Résumé: Cette thèse est dédiée à l'étude du support des réseaux mobiles dans IPv6, la nouvelle génération du protocole qui régit les communications dans l'Internet. Les travaux traditionnels dans ce domaine se préoccupent de fournir une connectivité permanente pour les stations mobiles. En revanche, l'objet de la présente étude est de traiter séparément le cas d'un réseau tout entier qui migre dans la topologie Internet, ce qui pose un certain nombre de nouveaux problèmes. Nous étudions tout d'abord l'État de l'Art dans le domaine traditionnel du support de la mobilité des stations mobiles. Cette étude nous permet de définir une taxinomie des propositions. En second lieu, nous définissons une nouvelle terminologie dédiée au support des réseaux mobiles, ainsi que leurs caractéristiques et les problèmes spécifiques causés par leur mobilité. Parmi un ensemble d'approches envisagées, nous nous consacrons tout particulièrement à l'usage de Mobile IPv6, le standard de l'IETF pour le support des stations mobiles. Dans un premier temps, nous proposons un certain nombre d'extensions nécessaires à ce protocole. Pour le long terme, nous proposons de réduire le coût des messages de contrôle induit par ce protocole au moyen de deux techniques multipoint. La première, dite traditionnelle, établit un arbre de distribution entre le réseau mobile et ses correspondants. La deuxième enregistre directement la liste des correspondants dans le message de contrôle. La performance de ces extensions multipoint est évaluée par simulation, et nous concluons cette dissertation par une vue d'ensemble d'une nouvelle architecture de gestion de la mobilité rassemblant diverses techniques, dont nos extensions multipoint.

Mots-Clés: IPv6 - Réseaux - Support Mobilité - Réseaux Mobiles - Routeurs Mobiles - Routage - Communications Multipoint - Internet - Simulation - Mobile IPv6