Mobile IPv6 Deployments: Graph-based Analysis and practical Guidelines

Guillaume Valadon [a,b,*] Clémence Magnien [a,b] Ryuji Wakikawa [c]

[a] *UPMC Univ Paris 06, UMR 7606 LIP6, F-75005, Paris, France*

[b] *CNRS, UMR 7606 LIP6, F-75005, Paris, France*

[c] *TOYOTA InfoTechnology Center, U.S.A., Inc.*

Abstract

The Mobile IPv6 protocol is a major solution to supply mobility services on the Internet. Many networking vendors have already implemented it in their operating systems and equipments. Moreover, it was recently selected to provide permanent IP addresses to end-users of WiMAX and 3GPP2. Mobile IPv6 relies on a specific router called the home agent that hides location changes of the mobile nodes from the rest of the Internet. To do so, the mobile nodes' traffic must flow through the home agent. This mandatory deviation produces longer paths and higher communication delays.

In order to solve these problems, we describe a new approach to address deployments of Mobile IPv6 based on graph theory and could be applied to any operator's network. In particular, we use notions of centrality in graphs to quantify increases of communication distances induced by dogleg routing and identify relevant home agents locations. We evaluate this approach using real-world network topologies and show that the obtained Mobile IPv6 performance could be close to direct paths ones. The proposed algorithm is generic and can be used to achieve efficient deployments of Mobile IPv6 as well as Home Agent Migration: a new Mobile IPv6 architecture using several distributed home agents.

Key words: Mobile IPv6, mobility management, graph, centrality

1. Introduction

The Mobile IPv6 protocol allows a mobile node to communicate using the same IPv6 address while it moves, thanks to a specific mobility router – the home agent. However, its performance is threatened by the dogleg routing: the mandatory deviation of the mobile node's traffic to the home agent. Choosing an accurate location of the home agent is consequently a critical aspect of Mobile IPv6 deployments since it impacts communication delays and path lengths.

Here, we model operator's networks as graphs and formally describe the impact of the dogleg routing according to the home agent's locations using well-known graph metrics. In particular, we use notions of degree and centrality in graphs to identify relevant home agents locations, i.e. such that the dogleg routing causes only a small increase to communications distances. Based on these observations, we present a generic methodology to address Mobile IPv6 deployment issues, which could be applied to any operator network.

Our proposal is singular compared to other solutions, such as the Return Routability Procedure [13], since the optimization of paths is only achieved with a correct placement of home agents, is non-intrusive protocol-wise, and is not affected by strict packet

* Corresponding author:
guillaume.valadon@lip6.fr

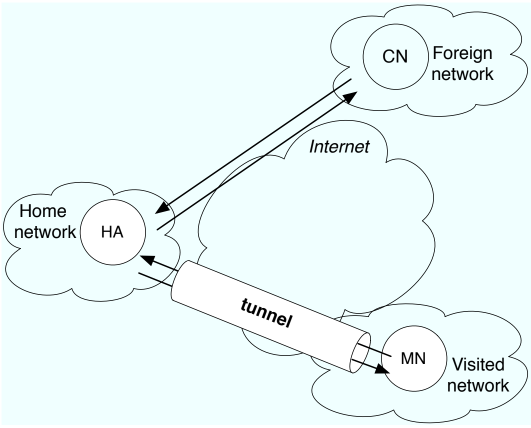

Fig. 1. Mobile IPv6: communication example

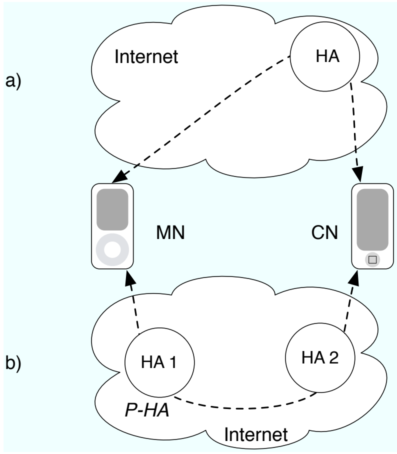

Fig. 2. Mobile IPv6 compared to Home Agent Migration

filtering policies. Consequently, it can immediately be used to enhance Mobile IPv6 deployments, and remain compatible with the current Internet architecture.

The rest of this article is organized as follows: we first describe the Mobile IPv6 protocol as well as Home Agent Migration in Section 2. We discuss notions of degree and centrality in Section 3. Then in Section 4, we formally define the impact of Mobile IPv6 and Home Agent Migration on path lengths, and discuss a methodology for communications comparing home agent locations. In Section 5, we numerically evaluate this methodology using real-world network topologies. Finally, prior to the conclusion, we examine related work in Section 6.

2. Considered mobility protocols

2.1. *Mobile IPv6*

The Mobile IPv6 protocol provides a unique permanent identifier (an IPv6 address) to a mobile node (MN) independently of its network of attachment. The key component of Mobile IPv6 is the home agent (HA), which is located in the mobile node's home network. It is a dedicated IPv6 router that manages the home IPv6 prefix, as well as the binding between the permanent IPv6 address (the identifier) and the IPv6 address acquired in the visited network (the locator). In the Mobile IPv6 terminology, these addresses are respectively referred as Home Address and Care-of Address. Packets sent to the Home Address, that belongs to the home prefix, by correspondent nodes [1] (CN) are routed to the home network and intercepted by the home agent which forwards them to the current location of the mobile node: its Care-of Address. Likewise, packets sent from the mobile node to its correspondent node must go through the home agent prior to being delivered. At any given time, the mobile node always communicates using its Home Address regardless of the network it is connected to. This deviation of packets to the home agent is called dogleg routing, and causes longer paths and higher communication delays between mobile nodes and their correspondents, see Figure 1.

The effects of the dogleg routing strongly depend on the home agent's locations. Due to the addressing and routing architectures, the location of a home agent is topologically and physically restricted by its home prefix. It must be in the correct physical location so that it can receive packets destined to the home prefix. Therefore, the home agent has to be placed where this prefix is advertised on the Internet.

2.2. *Home Agent Migration*

We proposed a new mobility architecture, called Home Agent Migration [22], that uses the traditional Mobile IPv6 protocol with an additional mobility management plane. In this new plane, home agents are distributed all over the Internet and exchange information about mobile nodes that they can reach. This deployment is performed with the help of any-

[1] from now on, we consider that correspondent nodes are always fixed nodes.

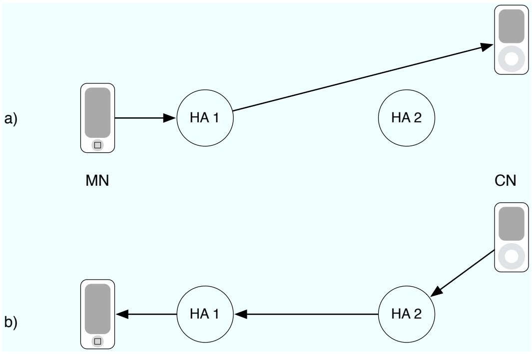

Fig. 3. Home Agent Migration: communication examples

cast routing [14] in which each home agent advertises the same home IPv6 prefix. Consequently, a mobile node will exchange traffic with its topologically closest home agent, reducing communication delays.

The aim of this new kind of home agent deployment is to provide an efficient route optimization scheme that (1) reduces communication latency, (2) is compatible with the current specification and implementation of Mobile IPv6's mobile nodes, and (3) is transparent for correspondent nodes.

Figure 2 compares the network architectures of Mobile IPv6 (a) and the proposed system (b). Due to the distribution of home agents in the Internet topology, a mobile node will always use the nearest home agent as if the home agent had migrated close to the mobile node's location. This closest home agent is referred to as the primary home agent (P-HA in Figure 2(b)). Similarly, packets sent by a second mobile node, here MN2, are routed to their closest home agent using generic IP routing mechanisms and are then forwarded to the primary home agent of MN1. With Mobile IPv6 (Figure 2(a)), the mobile nodes communicates through a single home agent.

Due to anycast routing, the communications between mobile and correspondent nodes are not symmetric. In Figure 3, we consider an architecture with two home agents HA1 and HA2 and two nodes, the mobile node MN, closer to HA1, and the correspondent node CN, closer to HA2. The notion of proximity is given by the regular routing of packets from the nodes to the home agents. From the MN to the CN (Figure 3(a)), there is no straightforward way to know which home agent is closer to the CN: HA1 directly sends the packet to the CN as a regular router does. From the CN to the MN (Figure 3(b)), packets sent from CN to MN1 are immediately intercepted by HA2 thanks to anycast routing then forwarded to MN1's primary home agent (HA1). Note that the communications between two mobile nodes are similar to communications from a correspondent node to a mobile.

3. Important graph vertices

In this article, we are looking for appropriate vertices for locating home agents. Note indeed that selecting the optimal locations for a given number of home agents is computationally expensive. Intuitively, such vertices should have an important position in the graph to avoid as much as possible the effects of the dogleg routing. In the literature, several metrics were defined to capture such characteristics. Here, we present the degree, the betweenness centrality and the shortest path centrality (that precisely identify vertices that gather most shortest paths within a graph). In following sections, we will use these three metrics as a way to find which home agent's locations are more likely to minimize communications distances with Mobile IPv6 and Home Agent Migration.

The degree of a vertex x is simply the number of edges that end in x. It is an efficient metric to identify important vertices. It is frequently used in graph analysis, as it is simple to calculate.

The betweenness centrality is used in various scientific fields such as sociology, see for instance [9]. For example, sociologists use it to pinpoint highly connected people in relationships graphs. Likewise, it is used to identify how a viral infection will spread in a given graph, and which highly connected vertices should be protected or cured to limit the epidemic [8]. The betweenness centrality intuitively estimates the importance of a vertex as a function of the number of shortest paths it lies on. This makes it very interesting in our context because if a home agent is located on a large number of shortest paths, it should have a small impact on communication distances.

More formally, the betweenness centrality of a vertex x, denoted by $C_B(x)$, is equal to:

$$C_B(x) = \sum_{y \neq z \neq x \in V} \frac{\rho_{yz}(x)}{\rho_{yz}},$$

where ρ_{yz} is the number of shortest paths from y to z, and $\rho_{yz}(x)$ the number of shortest paths from y to z that contain x.

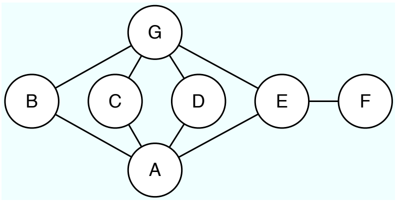

Fig. 4. Centrality: example graph

However, it turns out that the betweenness centrality is not completely adequate for our purpose. Indeed, when a vertex x is on a shortest path between two vertices u and v, if there are many shortest paths between u and v, this decreases the betweenness centrality of x.

Figure 4 illustrates this. Vertex A is on a shortest path between 18 pairs of vertices, and vertex E is on a shortest path between 12 pairs of vertices. Vertex A seems therefore to be a better home agent location than E. However, A has a lower betweenness centrality than E: their respective betweenness centrality are 9 and 10.5.

In our context what is essential is whether a vertex is on a shortest path between two vertices, regardless of the number of shortest paths between these vertices. We therefore introduce the shortest path centrality of a vertex x, noted $C_{SP}(x)$. It precisely describes the number of shortest paths that include x:

$$C_{SP}(x) = |\{(y, z), y, z \neq x, \rho_{yz}(x) \geq 1\}|$$

where $\rho_{yz}(x)$ is the number of shortest paths from node y to node z that contain x.

4. Methodology

In this section, we formally define the impact of Mobile IPv6 and Home Agent Migration on path lengths. Then, we discuss the placement strategies as well as the anycast routing emulation used afterward in the evaluation section.

4.1. *Impact on paths lengths*

We now discuss the influence of the Mobile IPv6 and Home Agent Migration protocols on the path length between two vertices in a graph. Home agents, mobile nodes, and correspondent nodes can be located in any vertex. We define two new notations: $d_{mip6}(A, B)$ and $d_{ham}(A, B)$ denote the length of a communication path from A to B when Mobile IPv6 or Home Agent Migration is used, respectively. We call it the *communication distance*. In general, it is longer than the *direct* distance between A and B, because the communication path has to go through one or two home agents depending on the case.

With Mobile IPv6, packets exchanged between two (correspondent or mobile) nodes A and B must go through the home agent HA. The resulting communication distance between A and B is therefore the sum of the distances between A and HA, and between HA and B:

$$d_{mip6}(A, B) = d(A, HA) + d(HA, B).$$

Note that with Mobile IPv6, the paths between a mobile node and a correspondent node are symmetric. Communications between mobile nodes and correspondent nodes are equivalent to communications between correspondent nodes and mobile nodes.

Contrary to Mobile IPv6, communications between correspondent and mobile nodes are in general not symmetric when Home Agent Migration is used. As a consequence, we need to consider separately three possible communication patterns: between two mobile nodes $MN1$ and $MN2$, from a mobile node MN to a correspondent node CN, and finally from a correspondent node CN to a mobile node MN. In the following equations, the notation HA_X represents the home agent that is closest to node X according to anycast routing.

Concerning communications between two mobile nodes, the communication distance between $MN1$ and $MN2$ is simply the sum of the distances between $MN1$ and its primary home agent HA_{MN1}, between HA_{MN1} and HA_{MN2}, and between $MN2$ and its primary home agent HA_{MN2}:

$$\begin{aligned} d_{ham}(MN1, MN2) = d(MN1, HA_{MN1}) \\ + d(HA_{MN1}, HA_{MN2}) \\ + d(HA_{MN2}, MN2). \end{aligned} \quad (1)$$

For communications between a correspondent node and a mobile node, the behavior is the same:

$$\begin{aligned} d_{ham}(CN, MN) = d(CN, HA_{CN}) \\ + d(HA_{CN}, HA_{MN}) \\ + d(HA_{MN}, CN). \end{aligned} \quad (2)$$

Finally concerning mobile nodes to correspondent nodes communications, the path goes only through the primary home agent of the mobile node:

$$d_{ham}(MN, CN) = d(MN, HA_{MN}) + d(HA_{MN}, CN). \quad (3)$$

Our goal is to find locations for home agents such that the communication distances are as short as possible. Regarding Mobile IPv6, the path length between two nodes A and B is not altered by the dogleg routing, i.e. $d(A, B) = d_{mip6}(A, B)$, when the home agent is located on the shortest path between A and B. Consequently, good home agent locations correspond to vertices that belong to a large number of shortest paths in the graph.

A similar discussion applies to Home Agent Migration: when two mobile nodes $MN1$ and $MN2$ communicate [2], the two primary home agents HA_{MN1} and HA_{MN2} must be on the shortest path between $MN1$ and $MN2$ to obtain performance similar to the *direct* communication, i.e. $d(MN1, MN2) = d_{ham}(MN1, MN2)$. With Home Agent Migration, increasing the number of home agents enables the mobility system to control more shortest paths and, as a result, allows to achieve better performance. Finding out the best arrangement of a limited number of home agents is a difficult and computationally expensive task. Indeed, it implies computing all of the possible home agents arrangements, then recomputing paths lengths according to home agents locations, and finally comparing the resulting modified distances with the *direct* distances.

4.2. *Placement strategies*

Three different strategies are used to identify relevant locations for home agents: one uses the degree, the other one the betweenness centrality, and the last one uses shortest path centrality. We study home agents locations by placing them in vertices sorted by decreasing degree, betweenness centrality and shortest path centrality, and comparing the distance between two vertices to the communication distances with d_{mip6} and d_{ham}.

In Internet-like graphs, the betweenness and the degree centrality can be correlated. It has been shown that, to some extent, the betweenness of a vertex is proportional to its degree [7]. This is an important result concerning the upcoming evaluation as the degree can be computed faster than betweenness.

4.3. *Anycast routing emulation*

With Home Agent Migration, a mobile node is always associated with its topologically closest home agent. Likewise, packets sent from a correspondent node to a mobile node are intercepted by the home agent that is the closest to the correspondent node. In practice, these automatic selections of home agents are accomplished thanks to anycast routing, and the routing protocols used to advert the mobile prefix to the network.

We emulate anycast routing by considering that primary home agents are the closest to mobile or correspondent nodes. When two home agents are equally distant from the node, the preference is given to the first home agent according to some fixed order.

5. Evaluation

In this Section, we provide an evaluation of home agent locations based on the methodology previously described, and compare the three location strategies (according to degree, betweenness centrality [3], and shortest path centrality in decreasing order). We first describe the three graphs that we will use for this evaluation. Then, we separately discuss Mobile IPv6 and Home Agent Migration, and show that it is possible to minimize their impact on communication distances by judiciously choosing the home agent locations.

5.1. *Studied networks*

Here, we consider three communications networks represented as graphs:

(i) `GEANT` network, a European research network spread over thirty countries. The graph was built using information about the layer 2 topology released in December 2004 on the GÉANT web page [2]. It contains 63 vertices. Edges are weighted using routing costs [4] that express specific routing optimizations made

[2] also true with communications from a correspondent node to a mobile node.

[3] computed using the Brandes algorithm [20].

[4] the ISIS routing protocol is used in GÉANT.

by networks administrators. In practice, these costs help to differentiate network links and reflect that some are more important than others: high routing costs are for example used to refrain the usage of transit links in favor of peering ones.

(ii) `WIDE` project's network: a Japanese research network connecting university campuses and companies. The graph was obtained using layer 2 topology information privately available to members of the project, then refined based on personal discussions with the network administrators. It contains 28 nodes that correspond to Network Operating Centers, campus networks, as well as peering points. Edges are not weighted.

(iii) `V6-MAP`: the IPv6 topology of the Internet produced on June 2003 using the network cartographer (nec) tool [3]. It is publicly available on the tool's web page. This tool sends traceroute queries to a set of distributed servers and is able to identify IPv6 addresses that belong to the same router in order to combine them into a single vertex. Moreover, it merges the different topologies discovered by the traceroute servers. The graph contains 4256 vertices. Edges are not weighted.

5.2. *Mobile IPv6*

In Mobile IPv6, the communications between two mobile nodes, and between a mobile node and a correspondent node are equivalent: all packets must go through the home agent. As a consequence, in this evaluation, we will only consider the influence of home agent locations on communications between two mobile nodes.

From a deployment point of view, Mobile IPv6 is still a young protocol. As of today, it has never been commercially deployed; its practical usage is mainly limited to testbeds or research networks. There is therefore no such thing as a typical Mobile IPv6 deployment. Nevertheless, we noticed that network administrators often locate home agents in subnetworks for management convenience. Such a setup is also quite common in the literature. We will now study the impact of the home agent location on communication distances when it is located in a subnetwork, i.e. a vertex of degree one.

Figures 5 and 6 show the modification of path lengths when the home agent is located in a subnetwork, for the `WIDE`, `GEANT` and `V6-MAP` graphs. The x-axis represents the difference between the communication distances and *direct* distances. The y-axis represents the number of pairs of vertices in the graph for which the increase of the *direct* distance is less than or equal to the corresponding x value. For example, the point at coordinates $(2, 73)$ means that 73% of the pairs of vertices have their direct distances increased by at most 2. In these four figures, *min* corresponds to the subnetwork which modifies distances the most, and *max* to the subnetwork which modifies distances the least, i.e. the subnetwork that gives the best performance.

For the `WIDE` network (Figure 5(a)), when the home agent is located in the *max* subnetwork, 55% of all pairs of vertices distances are increased by at most 3. The corresponding fraction is 13% with *min*. We observe a similar behavior in the `GEANT` network (Figure 6(a)). Concerning the `V6-MAP` network (Figure 5(b)), the difference is even more significant. When the home agent is in *min*, only 0.18% of pairs of vertices are increased by at most of 4. In this case, all shortest paths are seriously modified: their distances are increased up to 72. With *max*, 75% of all shortest paths are only increased by at most 4. Finally when weights (see Section 5.1) are used to compute the distance in the `GEANT` network (Figure 6(b)) the two considered subnetworks deliver even more dissimilar results: with *min* 94% of all pairs of vertices are increased by more than 2, and some are increased up to 13210. In contrast, with *max*, this fraction is higher: 52%.

It is therefore clear that the subnetworks are not identical from a performance point of view: they do not equally modify distances. There is, however, no straightforward solution to find out which vertex will perform the best comparing to the other vertices of degree one.

Using the same method, we now compare the subnetwork that modify paths lengths the least (called *max* to keep the same naming convention) to vertices with the highest degree, betweenness and shortest path centrality. Results are presented in Figures 7(a), and 7(b) for the unweighted graphs `WIDE` and `V6-MAP` and in Figures 8(a) and 8(b) for the graph `GEANT`.

When the home agent is located in one of the two vertices with the highest degree, betweenness or shortest path centrality, the number of paths that

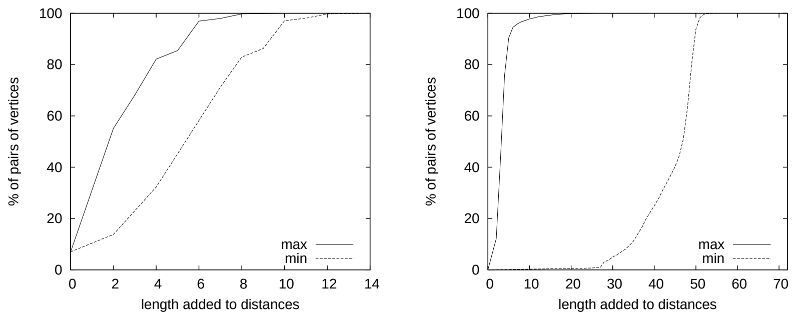

(a) `WIDE` (b) `V6-MAP`

Fig. 5. Mobile IPv6: one home agent located in a subnetwork − CDF of communications distances increase compared to direct distances. *min*: subnetwork that modifies path lengths the most (worst). *max*: subnetwork that modifies path lengths the least (better).

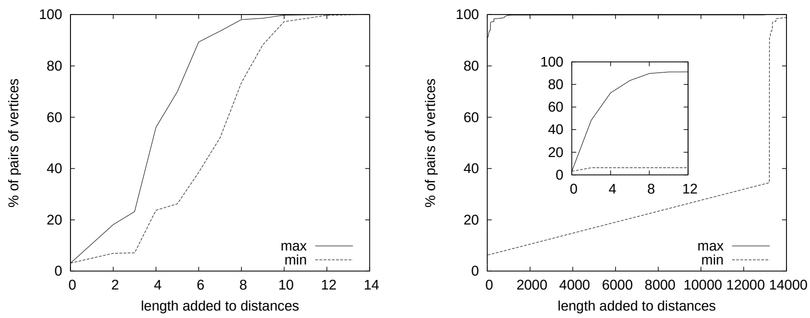

(a) `GEANT` **unweighted** (b) `GEANT` **weighted**

Fig. 6. Mobile IPv6: one home agent located in a subnetwork (`GEANT`) − CDF of communications distances increase compared to direct distances. *min*: subnetwork that modifies path lengths the most (worst). *max*: subnetwork that modifies path lengths the least (better).

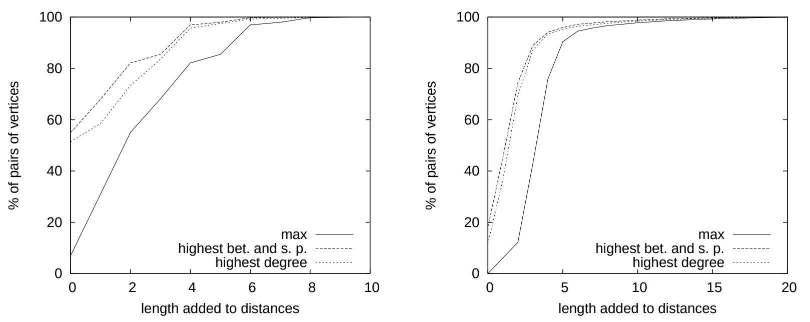

(a) `WIDE` (b) `V6-MAP`

Fig. 7. Mobile IPv6: degree, betweenness and shortest path centrality − CDF of communication distances increase compared to direct distances. *max*: subnetwork that modifies path lengths the least.

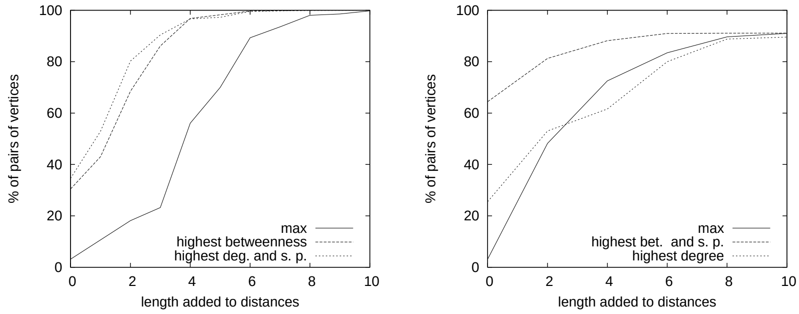

(a) GEANT **unweighted** (b) GEANT **weighted**

Fig. 8. Mobile IPv6: degree, betweenness and shortest path centrality (GEANT) − CDF of communication distances increase compared to direct distances. *max*: subnetwork that modifies path lengths the least.

are not modified is drastically increased in the unweighted graphs. In WIDE, 51% and 54% of all shortest paths are not modified when the degree and the betweenness, and shortest path, centrality are respectively used to select the home agent location. This is 7% with *max*. Similarly in V6-MAP, 17% and 26% are not modified with the degree and the betweenness, and the shortest path centrality; this is 0.04% when the home agent is located in *max*.

From now on, we respectively call HB, HD, and HSP the vertices with the highest betweenness centrality, the highest degree, and the highest shortest path centrality. In the general case, placing the home agent in HB provides slightly better results than placing it in HD. In GEANT however, the performance of HD is slightly better than HB. For instance, the fraction of pairs for which the communication distance is equal to the real distance is 34% for HD and 30% for HB. This behavior indicates that the vertex HD is indeed on a shortest path between more pairs of vertices than HB. As a matter of fact, in all of the considered graphs, when the home agent is located in HSP, the results are always the best, i.e. either equal to HB or HD.

When weights are used with GEANT, HB and HSP as a home agent location gives better results than either HD or *max*. The comparison between HD and *max* is however not as straightforward, and reveals some interesting features. The number of pairs for which the communication distance is strictly equal to the *direct* distance is higher for HD than *max*. Notice that in this case, the vertex with the highest degree is the one with the second highest betweenness centrality. This shows that even though the betweenness centrality does not always capture the fact that a vertex is a good location for a home agent, it is still a very good indicator. However, this tendency is reversed when we consider cases in which the communication distance is larger than the *direct* distance. For instance, with *max* 83% of all pairs of vertices are increased by 6 or less; this is 79% with HD. This indicates that when the home agent is located in the subnetwork *max*, the communication distances are smaller than when it is located in HD. Indeed, even though the vertex *max* has as small betweenness, it is the neighbor of the vertex with the second highest betweenness centrality. Consequently, it delivers slightly better results than HD.

5.3. *Home Agent Migration*

Here, we only consider the betweenness centrality as the placement strategy; the observations are similar with the degree and the shortest path centrality. The three placement strategies are compared and discussed in Section 5.4.

Unlike in the Mobile IPv6 protocol, the communications between mobile and correspondent nodes are not symmetric with Home Agent Migration (see Equations 2 and 3 on page 5). Prior to other studies, we will therefore evaluate this difference by comparing the communication distances from a mobile node to a correspondent (*MN-CN*) and from a mobile node to another mobile node (*MN-MN*) [5].

By definition, the communication distance *MN-MN* is larger than or equal to the *MN-CN* one . Figures 9 and 10 represent the difference of communication distances between these two communication patterns for the WIDE, GEANT, and V6-MAP graphs.

[5] which is the same as the communication distance from a correspondent node to a mobile node.

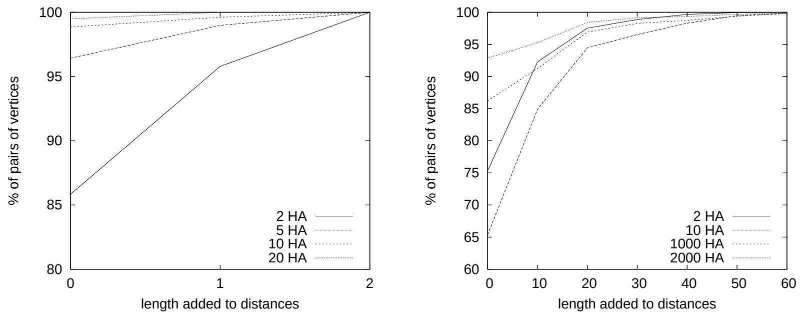

(a) `WIDE` (b) `V6-MAP`

Fig. 9. Home Agent Migration: correspondent and mobile nodes communications − CDF of communication distances increase between the two communication patterns *MN-CN* and *MN-MN*. Home agents are selected using the betweenness centrality.

The x-axis represents the difference of communication distances. The y-axis represents the number of pairs of vertices for which the difference between the communication distance is lesser than or equal to the corresponding x value. The case of one home agent is not represented because it is equivalent to Mobile IPv6 and the communication distances are equal in this case.

We can see that the difference of the two communications patterns is very small. In all graphs except `V6-MAP`, this difference is at most 2. Note, that `V6-MAP` provides worse performance than other graphs but this is also the case with Mobile IPv6. In Figures 9(a) and 9(b), when the number of home agents increases, the difference between the two communication patterns becomes less important. In the `WIDE` graph, when five home agents are used, 97% of communications distances are equal. We observe comparable results with the `GEANT` graph. Concerning the `V6-MAP` graph, we chose to represent a broad range of home agent sets. However, if we consider a realistic deployment of Home Agent Migration, only small values are important. Deploying and managing a high number of home agents is indeed not feasible in practice.

Finally, even though the difference between communication distances from a mobile node to a correspondent node and from a mobile node to a mobile node (or, equivalently from a correspondent node to a mobile node) can be important in some cases, in practice it is quite small. In the following discussions, we will therefore only consider communications between two mobile nodes.

We now make some observations concerning the the impact of the number of home agents on communication distances. We select home agents by decreasing betweenness centrality: when a single home agent is considered, it is placed on the vertex with the highest betweenness centrality, when two home agents are considered, they are placed on the two vertices with the highest betweenness centrality, and so on [6] .

Figures 11 and 12 show the modification of communication distances when the number of home agents is increased for the `WIDE`, `GEANT` and `V6-MAP` graphs. The x-axis represents the difference between communication distances and direct distances. The y-axis represents the number of pairs of vertices in the graph for which the increase of the distance is lesser than or equal to the corresponding x value.

From these figures, it is clear that increasing the number of home agents increases the number of shortest paths controlled by the Home Agent Migration infrastructure. As a consequence, less communication distances are modified when more home agents are added to the system. However, in the case of the `V6-MAP` graph (Figure 11(b)) with less than five home agents, Home Agent Migration is less efficient than Mobile IPv6 (one home agent). This is linked to the fact that, as already explained, communications between two mobile nodes in Home Agent Migration must go through two home agents, instead of one in Mobile IPv6. By adding one home agent, it is therefore possible that the performance is degraded. For instance, if there was a home agent on a shortest path between two mobile nodes, the new home agent can be closer to one of the mobile nodes but not on the shortest path. As a result, with Home Agent Migration, when home agents are

[6] when only one home agent is considered, Home Agent Migration is equivalent to Mobile IPv6.

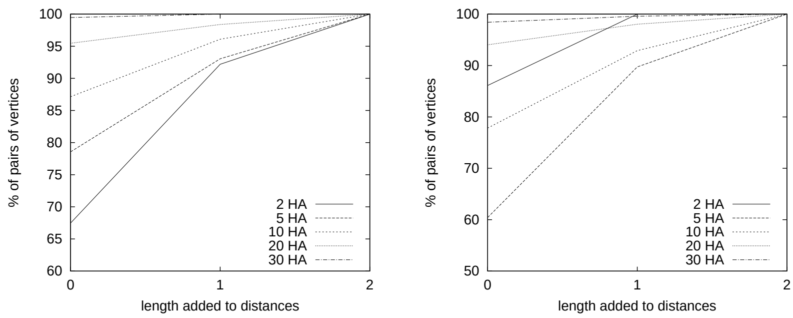

(a) GEANT **unweighted** (b) GEANT **weighted**

Fig. 10. Home Agent Migration: correspondent and mobile nodes communications (GEANT) − CDF of communication distances increase between the two communication patterns *MN-CN* and *MN-MN*. Home agents are selected using the betweenness centrality.

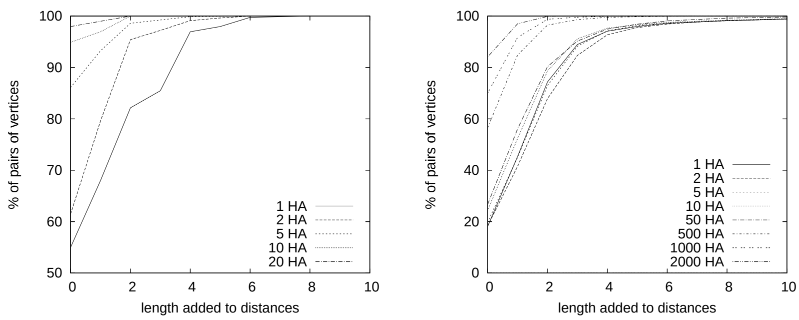

(a) WIDE (b) V6-MAP

Fig. 11. Home Agent Migration: increasing the number of home agents − CDF of communications distances increase compared to direct distances.

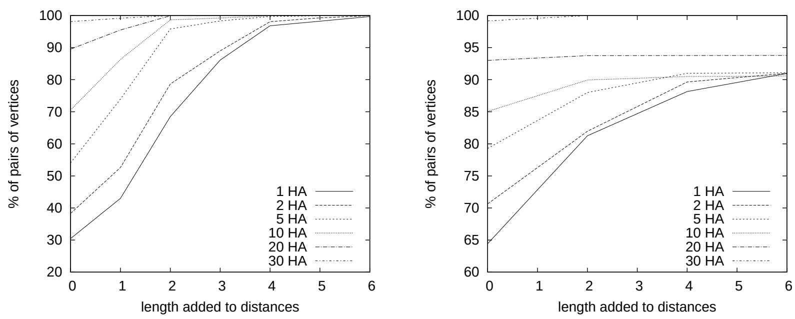

(a) GEANT **unweighted** (b) GEANT **weighted**

Fig. 12. Home Agent Migration: increasing the number of home agents (GEANT) − CDF of communications distances increase compared to direct distances.

far from each other, their impact on communication distances is also more important.

As expected, adding more home agent leads, in general, to better performance but this is not always this simple and we will study this in more details in the next section.

5.4. *Degree, betweenness and shortest path centrality*

Finding out the right number of home agents is a difficult task that also depends on the placement strategies for the home agents. Therefore, we now study together the number of home agents and the placement strategies consisting in placing them in vertices with highest degrees, betweenness centralities and shortest path centralities. Figures 13 and 14 show the percentage of the communication distances increased by 2 or less with a given number of home agent for the `WIDE`, `GEANT` and `V6-MAP` graphs.

When the number of home agents is low, we observe a remarkable behavior in two cases. In Figure 14(a), considering the degree, if three home agents are used instead of two, the performance are worse. We observe a similar behavior with the `V6-MAP` graph and a large number of home agents, see Figure 13(b). This is an important result that shows that care should be taken when deploying Home Agent Migration.

For the `GEANT` graphs, when home agents are selected using the shortest path centrality, the communication distances are less modified than when the degree or the betweenness are used.

With Mobile IPv6, we showed that the shortest path centrality always selects the best home agent location. The same conclusion does not apply for Home Agent Migration. Indeed, in the case of multiple home agents, whereas the shortest path centrality delivers good results, it might not be the best placement strategies especially when the number of home agents is high.

The shortest path centrality gives good performance. However, it does not always give better ones than the betweenness centrality or the degree. Moreover, these metrics do not have the same computation cost. Indeed the degree is directly retrieved using the graph description. The betweenness and the shortest path centralities are more expensive to compute since they require computing all shortest paths between all pairs of vertices. Their complexity is $\mathcal{O}(n * m)$. We will now study how these measures are related.

Figures 15(a) and 15(b) respectively represent the betweenness and the shortest path centrality values, in log scale, plotted against the degree in the `V6-MAP` graph. Note that while the results of the `WIDE` and `GEANT` graphs are similar to the ones given in this section, their visual representation is not significant enough to be provided.

Both figures indicate that vertices that have the highest degrees also have very high centrality values. This is a remarkable property concerning our placement strategies as we proposed to locate home agents in vertices that have the highest degree and centrality values. Such a property is however not symmetric. From the two figures, it is clear that vertices with a small degree do not have necessarily a small centrality value.

This is a strong result that is particularly pertinent in practice. Indeed, when the full network topology is not known, or when it is not possible to compute centrality metrics, home agents can be located in, or close to, highly connected routers (i.e. high degree nodes) in order to obtain good performance.

6. Related work

While placement strategies of servers' instances is a well studied topic both in research and industrial fields, little has been done, strictly speaking, regarding the Mobile IPv6 protocol and the placement of home agents in network topologies. However, several proposals were defined at the IETF to solve some of Mobile IPv6 limitations. Their main goals are to reduce the number of signaling packets sent from the mobile node to the home agent. For example, Fast Handover for Mobile IPv6 (FMIPv6) [15] pushes some of Mobile IPv6 functionalities into access routers to permit mobile nodes to communicate immediately after a handover prior to rebinding to the home agent. Similarly, the multiple Care-of Address (mCoA) extension [21] describes another approach that allows mobile nodes with multiple network interfaces to anticipate handovers by binding different Care-of Addresses to the same Home Address. On the other hand, Hierarchical Mobile IPv6 [4,19] proposes to locate home agents closer to mobile nodes using a dedicated hierarchy of home agents that minimize the impact of mobile nodes locations, and movements, on the communication per-

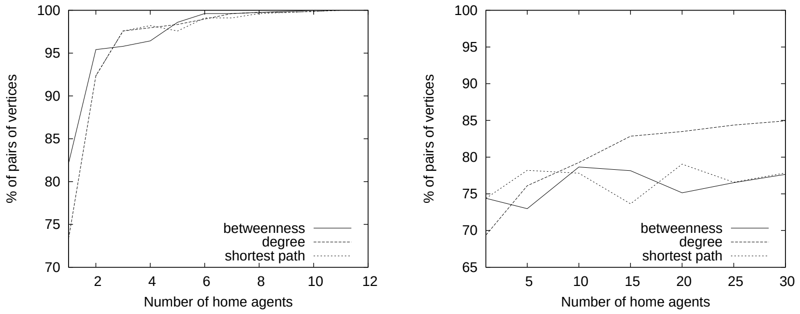

(a) `WIDE` (b) `V6-MAP`

Fig. 13. Home Agent Migration: comparison of degree, betweenness and shortest path centrality − CDF of communication distances modified by 2 or less, when the corresponding number of home agents is used (higher is better).

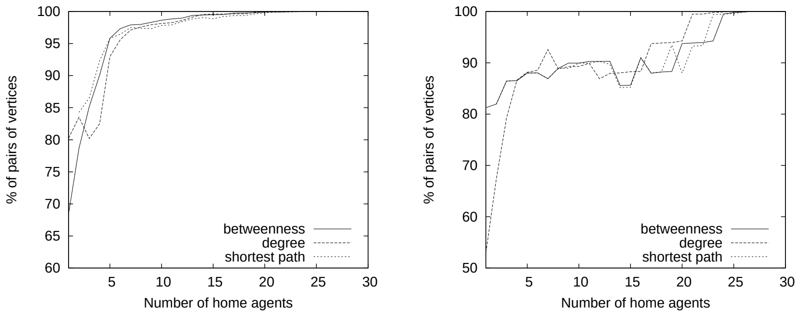

(a) `GEANT` **unweighted** (b) `GEANT` **weighted**

Fig. 14. Home Agent Migration: comparison of degree, betweenness and shortest path centrality (`GEANT`) − CDF of communication distances modified by 2 or less, when the corresponding number of home agents is used (higher is better).

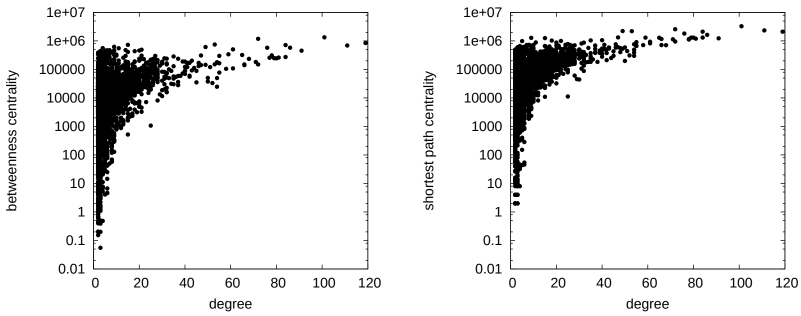

(a) Betweenness centrality (b) Shortest path centrality

Fig. 15. Relationship between degree, betweenness and shortest path centrality in the `V6-MAP` graph.

formance. However, this proposal does not specifically describe how to locate home agents but provides good performance and keeps the network load low as the network's size and the number of home agents increases.

The influence of home agents locations on the performance has been specifically studied in Universal Mobile Telecommunications System (UMTS) [18]. Using a network simulator, the author evaluated different home agents locations in a given UMTS architecture. The results clearly outline that concerning data transmission, there is a direct relation between the home agent location and the performance: the closer the home agents are to mobile nodes, the higher the effective bandwidth is. While this paper uses an approach different from our proposal, its outcome is complementary and indicates that placing home agents based on degree, or betweenness or shortest path centrality is not only improving the paths lengths in theory but will also deliver enhanced performance in practice.

Finding the judicious locations of servers' instances is a topic that has already been addressed in different networking fields, ranging from Content Delivery Networks to Internet mapping architectures. The common concept is to distribute multiple instances of the service in the network topology in order to maximize a quality function. Depending on the context, such a function aims at minimizing the distances between the multiple instances and end-nodes, maximizing end-nodes bandwidth and users experience, reducing the number of retransmitted packets, or using high quality paths.

Content Delivery Networks deliver media content to end-users by transparently distributing duplicated mirrors in the Internet architecture. Here, the quality function aims at optimizing the placement based on network metrics. Indeed, a common strategy is to identify relevant mirror locations based on the estimated path lengths to end-nodes [12]. Other works [17,6] use more metrics in the quality function such as client bandwidth, request rates, or path quality, and show that, when possible, taking network metrics into account enhances the accuracy of the placement strategy.

Concerning Internet mapping using traceroute like approaches, performing measurements from multiple locations allows to discover more routers and links, and as a result improve the overall quality of the mapping [16]. In this case, the quality function aims at finding monitor locations, as well as destinations to probe, that maximize the number of discovered routers and links. To some extent, this problem is similar to the efficient placement of home agents in which we want to control a maximum number of shortest paths. Several strategies based on random selection, or vertices degrees sorted by increasing and decreasing degree can be used to place the traceroute monitors in the Internet topology [10]. However, this work indicates that the most efficient solution is to select servers and locations by decreasing degrees, similarly to what was done in this article.

Distributed server instances is nowadays a key element of the Internet infrastructure that is used commercially to place the media content closer to the users [1], or sharing network load among the instances. One of the closest approaches to ours comes from the root DNS server deployments using anycast routing. Its efficiency in performing load balancing among the multiple instances was the key feature that drastically mitigated the Distributed Denial of Service attack against the DNS root servers in February 2007. Similarly, studies [11,5] showed that anycast routing is an effective solution to decrease the latency of replies. This is an important practical result that strengthens the conjoint use of the Home Agent Migration architecture and our placement proposal.

7. Conclusion

In this article, we formally describe the Mobile IPv6 and Home Agent Migration protocols using graph metrics in order to quantify their impact on communication distances. We proposed and evaluated three placement strategies that use the vertices' degree, betweenness and shortest path centrality to locate home agents. Our proposal has the advantage of providing an accurate description of genuine operational issues raised by these two mobility architectures.

Concerning the Mobile IPv6 protocol, we described the dogleg routing in terms of distances, and confirmed that all home agent locations do not provide equivalent performance. Locating home agents in subnetworks, as is often done in the literature, could indeed lead to bad performance and seriously increase communication distances. We also provided a detailed analysis of Home Agent Migration that highlights that using even a small number of home agents will in general, but not always, drastically decrease communication distances. Selecting the best locations for a given number of home agents being computationally expensive, we proposed three solutions to finding good locations that use the degree and the betweenness and the shortest path centrality, and showed that they allow to efficiently find home agent locations that provide good overall performance.

For Mobile IPv6, the shortest path centrality is the most efficient metric. For Home Agent Migration however, it is difficult to identify the metric that will always be the most effective. Nevertheless, in practice, the relation between high degree and high betweenness, or shortest path, centrality is especially interesting. Indeed, it could be used by system administrators who want to deploy Mobile IPv6 or Home Agent Migration, but cannot apply our graph based strategies because they do not have a graph model of their network. They could instead locate the home agent close to highly connected Network Operating Centers or routers; i.e. nodes with a high degree.

During our experiments, we identified several issues in the placement strategies that are linked to the underlying graph topology. In some specific cases, vertices with a high betweenness centrality deliver inferior performance than vertices with a lower betweenness centrality. In order to solve this issue we proposed a new centrality metric called shortest path centrality. It succeeds at precisely identifying nodes included in the maximum number of shortest paths; which is the fundamental property of efficient home agent locations. In future work, we would like to study this metric in depth in order to understand if it can be applied to other topics than home agents placement.

We believe that taking into account network metrics such as the Round Trip Time between routers, by using them as edge weights, is an important follow-up work. It will not only provide more accurate results, but will also ease their interpretation for real deployments. Likewise, the knowledge of mobile node locations would surely improve the outcome of our placement strategies. This will allow focusing only on communication paths that are specifically used by mobile nodes. Traffic matrices, such as available for the `GEANT` graph, could also help locating which routers gather most of the traffic, and which communication paths must be improved in priority.

Acknowledgment

The authors would like to acknowledge Matthieu Latapy and Laurent Bernaille for their comments and valuable discussions.

References

[1] Akamai. http://www.akamai.com/.

[2] The GÉANT Website. http://www.geant.net/.

[3] network cartographer (nec). http://dpt-info.u-strasbg.fr/∼magoni/nec/.

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

[5] L. Colitti, E. Romijn, H. Uijterwaal, and A. Robachevsky. Evaluating the effects of anycast on DNS root name servers. In *RIPE document RIPE-393*, 2006.

[6] Mark Crovella and Robert Carter. Dynamic Server Selection in the Internet. In *Proceedings of the Third IEEE Workshop on the Architecture and Implementation of High Performance Communication Subsystems (HPCS '95)*, 1995.

[7] L. Dall'Asta, I. Alvarez-Hamelin, A. Barrat, A. Vazquez, and A. Vespignani. Exploring networks with traceroute-like probes: Theory and simulations. In *Theoretical Computer Science 355*, 2006.

[8] David C. Bell, John S. Atkinson and Jerry W. Carlson. Centrality measures for disease transmission networks. *Social Networks*, 21:1–21, 1999.

[9] Linton C. Freeman. A Set of Measures of Centrality Based on Betweenness. *Sociometry*, 40(1):35–41, March 1977.

[10] Jean-Loup Guillaume, Matthieu Latapy, and Damien Magoni. Relevance of massively distributed explorations of the Internet topology: Qualitative results. *Computer Networks*, 50(16):3197–3224, 2006.

[11] ICANN. Root server attack on 6 February 2007. ICANN factsheet, ICANN, 2007.

[12] Sugih Jamin, Cheng Jin, Yixin Jin, Danny Raz, Yuval Shavitt, and Lixia Zhang. On the Placement of Internet Instrumentation. In *INFOCOM 2000 (1)*, pages 295–304, 2000.

[13] D. Johnson, C. Perkins, and J. Arkko. Mobility Support in IPv6. RFC 3775 (Proposed Standard), June 2004.

[14] D. Katabi and J. Wroclawski. A framework for scalable global IP-anycast (GIA). In *ACM SIGCOMM'2000*, August 2000.

[15] R. Koodli. Fast Handovers for Mobile IPv6. RFC 4068 (Experimental), July 2005.

[16] Damien Magoni and Mickal Hoerdt. Internet core topology mapping and analysis. *Computer Communications*, 28(5):494–506, 2005.

[17] L. Qiu, V. N. Padmanabhan, and G. M. Voelker. One The Placement of Web Server Replicas. In *INFOCOM 2001*, 2001.

[18] Behcet Sarikaya. Home Agent Placement and IP Address Management for Integrating WLANs with Cellular Networks. *Wireless Communications, IEEE*, 13:77–86, December 2006.

[19] H. Soliman, C. Castelluccia, K. El Malki, and L. Bellier. Hierarchical Mobile IPv6 Mobility Management (HMIPv6). RFC 4140 (Experimental), August 2005.

[20] Ulrik Brandes. A Faster Algorithm for Betweenness Centrality. *Journal of Mathematical Sociology*, 25:163–167, 2001.

[21] R. Wakikawa, T. Ernst, K. Nagami, and V. Devarapalli. Multiple Care-of Addresses Registration (work in progress, draft-ietf-monami6-multiplecoa-06). Internet Draft, Internet Engineering Task Force, February 2008.

[22] Ryuji Wakikawa, Guillaume Valadon, and Jun Murai. Migrating Home Agents Towards Internet-Scale Mobility Deployments. In *CoNext06, Lisbonne, Portugal*, December 2006.