Studies to provide an end-to-end IP solution for a Mobile Extranet

C. Blondia and O. Casals

Dept. Mathematics and Computer Science
University of Antwerp, Universiteitsplein 1, B-2610 Antwerpen-Belgium
blondia@uia.ua.ac.be, http://win-www.uia.ac.be/u/pats/

Computer Architecture Department Department
Polytechnic University of Catalonia, C/Jordi Girona 1-3, E-08034 Barcelona-Spain
olga@ac.upc.es, http://research.ac.upc.es/XARXES/CompNet

Abstract

The European MOEBIUS project (Mobile Extranet Based Integrated User Services) is developing a platform for remote access to intranet services using mobile access technology and investigating its performance characteristics to demonstrate the advantages for health care, business and residential applications. This paper focuses on the studies that are being done in MOEBIUS on two particular items: (i) the applicability of the transport control protocol (TCP) used in Internet in a wireless environment and (ii) the provisioning of an end-to-end IP solution with the use of Cellular IP for micro-mobility.

Key words

Extranet, transport control protocol, micro-mobility

1. Introduction

The aim of the European IST project MOEBIUS (Mobile Extranet Based Integrated User Services) is to build an integrated mobile service platform supporting mainly health care applications. The reference architecture (Figure 1) for the Mobile Extranet contains the following main elements:

- The hospital Intranet with:
    - A generic server providing information for an Intranet user
    - A firewall function including features suitable for Mobile Extranet
    - A WAP server
- The fixed terminal location in the Intranet
- A mobile terminal location external to the Intranet and relayed to the Intranet through:
    - A GPRS network
    - The public Internet

The platform uses GPRS for Layer 2 mobility together with Mobile IPv6 technologies and provides the appropriate security mechanisms. GPRS has been chosen for the experimental phase of the project as it is the most attractive solution in a short-medium term. GPRS is an already deployed solution for radio access to IP-based services. GPRS offers a flexible resource reservation that is cost-efficient for supporting bursty data transfer.

Nevertheless, the Mobile Extranet concept defined in Moebius is not confined to the use of GPRS and evolutive solutions such as Cellular IP and UMTS have been considered by the system study part of the project. In particular, in this paper we present some results that have been obtained to better meet the global IP concept: the applicability of the transport control protocol (TCP) used in Internet in a wireless environment and the provisioning of an end-to-end IP solution with the use of Cellular IP for micro-mobility.

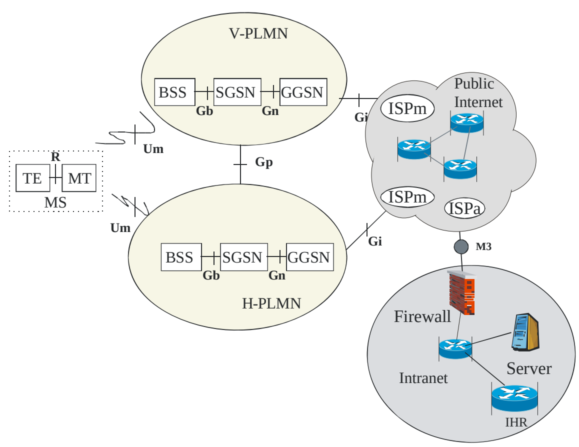

Figure 1: MOEBIUS reference architecture

1.1. TCP in a mobile environment

TCP is a variable window protocol where losses are considered as congestion signals. Basically, TCP increases the window as acknowledgements are received and reduces the window when losses are detected. Since TCP makes no distinction between losses due to congestion and losses due to errors in the transmission links, both produce window reduction. Therefore, losses due to transmission errors may unnecessarily produce a transmission rate reduction and make TCP perform poorly.

Different mechanisms have been proposed to improve the TCP performance over lossy links. These mechanisms can be classified into three categories: (i) end-to-end, (ii) link-layer and (iii) split-connection.

We analyze the TCP performance degradation introduced by a wireless media and the capability of different mechanisms to correct it. Since we assume that end-to-end IPsec is to be used, we have not considered solutions that break the TCP end-to-end semantics as split-connection. We evaluate the benefits of introducing: (i) improvements over the TCP Reno release, (ii) improvements over the classical "tail dropping" mechanism used by routers. We also study the unfairness that occurs when TCP sources having high transmission error rates compete with sources having low transmission error rates.

1.2. End-to-End IP solution

The new trends in mobile service provision push to have a common core network and to have IP protocol up to the base stations to minimize the number of technologies in the network, and therefore the cost.

Mobile IP has been widely accepted as protocol to provide macro-mobility (inter-domain layer 3 mobility) in the IETF world. Currently, IETF is proceeding forward proposing complete solutions for mobility management at layer 3 providing micro-mobility (intra-domain mobility) and enhancing inter-domain mobility and global roaming. The Mobile IP standard [1] has several drawbacks:

- Triangular routing: a node's home agent has to be notified of every change of location. These location updates incur the latency of traveling to the possibly distant home agent.
- Inefficient use of link resources through the use of tunneling that may span a large number of router hops.
- Support for QoS, e.g. through the use of RSVP, is not straightforward with triangular routing.

As currently defined, Mobile IP does not extend well to a large number of portable devices moving frequently between small cells. A lot of research is being done to solve the different problems and several approaches are being discussed within IETF. These approaches try to solve the problems by improving specific aspects of Mobile IP or by introducing new protocols for micro-mobility able to inter-work with Mobile IP.

We evaluate the efficiency of one of the protocols that have been proposed, namely HAWAII [2]. In this paper, we define models for the different HAWAII path setup schemes and use these models to evaluate the delay and delay jitter packets experience in a network that uses HAWAII for supporting intra-domain mobility.

2. Evaluation of mechanisms to improve TCP performance over wireless links

Among end-to-end improvements we have Selective Acknowledgments (SACK) [3], and New-Reno [4, 5]. These TCP enhancements try to avoid time-outs when multiple losses occur in the same window. SACK is more complex and needs collaboration between the TCP sender and transmitter. New-Reno, with a simpler implementation, can be used when SACK is not supported as only the TCP sender requires to be modified.

In this section we analyze the impact of a lossy wireless medium on the TCP performance and the benefits of using New-Reno and two implementations of SACK: a basic SACK implementation proposed in [5] that we will refer to simply as SACK, and a more complex implementation called Forward Acknowledgements (FACK) proposed in [6]. Furthermore, we study the unfairness that occurs when different transmission error rates are found among contending sources. We have also investigated the use of RED and ECN that show to be beneficial for this situation. Finally, to solve the drawbacks that arise when using TCP in a wireless environment we analyze a possible solution based on ECN routers.

2.1. New-Reno Overview

The fast retransmit and fast recovery implementation allows TCP to recover from a single segment loss in a window. When multiple losses occur, this implementation ends up waiting for a retransmission time-out. The New-Reno improvement is based on the observation that partial acks received during the fast recovery phase are an indication of segment losses. Partial acks are defined as those confirming part but not all of the segments that were outstanding at the start of the fast recovery period.

2.2. SACK Overview

With the cumulative acknowledgements used by TCP, the sender has little information to take retransmission decisions in case of multiple losses in a single window. A classical mechanism used by other ARQ protocols is the usage of selective acknowledgments by means of which the TCP receiver informs about out of order segments that have been correctly received. The RFC [3] describes the mechanisms to be used by the TCP sender and receiver to agree upon the usage of SACK option and the way by which the TCP receiver communicates the reception of correct segments received out of order. Using SACK the cumulative meaning of the ack number does not change. Additionally, the reception of out of order segments correctly received is communicated by specifying the edges of blocks of data correctly received that are contiguous and isolated. The space available in the TCP header to accommodate the options allows three of such blocks to be specified. Upon reception of SACK blocks the TCP sender may turn on a bit of the segments outstanding in the retransmission queue indicating that they have been correctly received. Therefore, when a retransmission is to be done, an unmarked segment is taken.

Communication of correct data reception by means of the SACK option does not prevent for the TCP sender the retransmission of this data. Actually, the TCP receiver is allowed to "renege" and discard blocks of data that have already been acknowledged by means of the SACK option. In order to retransmit these blocs the RFC [3] says that the TCP sender should reset the marked segments in the retransmission list upon a retransmit timeout.

Although a detailed explanation of the SACK indications is given in [3], no description of the TCP sender behavior is proposed. In our simulator we have used two implementation of SACK.

2.3. RED and ECN Overview

Traditionally, routers in the Internet have implemented a simple FIFO queue with a "tail dropping" management mechanism. Tail dropping consists of discarding incoming segments when no more room is available in the queue. This simple mechanism may lead to the following drawbacks:

- Unfairness: some unbalanced situations among contending TCP sources may lead to some sources monopolizing the queue space. For example, sources with shorter round trip delays are more favored than sources having higher delays.
- High end-to-end delays: tail dropping tends to maintain the queues full and thus with high end-to-end delays. Furthermore, synchronization effects may grow leading to periods of high and low transmission rate where the queue is respectively built up and drained. In this case end-to-end delays may have strong oscillations.

To solve these problems an "active queue management" has been recommended [7]. This consists of dropping packets at the routers before buffer overflows according to some algorithm which tries to maintain the queue around a low threshold. One of these algorithms is the Random Early Discard (RED) [8].

The Explicit Congestion Notification (ECN) [9] is a proposal to indicate a congestion state by setting a bit in the packet header. This bit-setting mechanism requires that the congestion control algorithm reacts to this bit indication in a similar way TCP reacts to packet drop and so [9] proposes to add this capability to TCP.

2.4. Numerical Results

Figure 2 shows the network topology we have simulated in this paper. A more extensive evaluation can be found in [10].

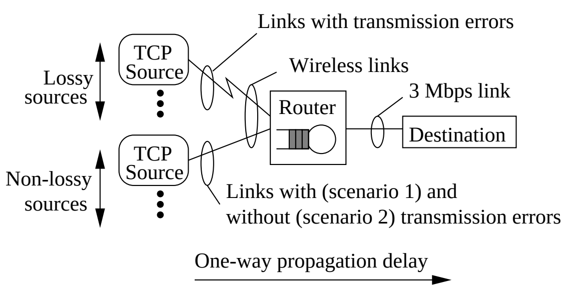

Figure 2: Simulation scenarios.

We have considered two scenarios:

- Lossy scenario with all wireless links having the same transmission error rate
- Half-lossy scenario where only half of the transmission links have transmission errors.

The following graphs show the *Gain*, for each group of sources having the same ratio of transmission errors, calculated as the average goodput over the fair goodput (link rate divided by number of sources) multiplied by 100.

The curves in the middle of the graph correspond to the lossy scenario while the ones at the top and the bottom correspond respectively to the lossy and non-lossy sources of the half-lossy scenario.

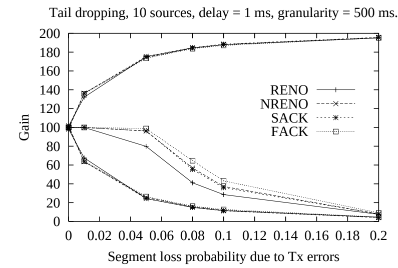

Figure 3: Gain obtained in the lossy and half-lossy scenarios using different TCP implementations.

In the following we derive some general guidelines from the simulation results.

Unfairness: the degradation of the gain obtained by the sources having losses due to transmission errors is higher when the congested link is shared with sources having no transmission losses. The sources without transmission errors have the tendency to lock-out the sources with transmission errors.

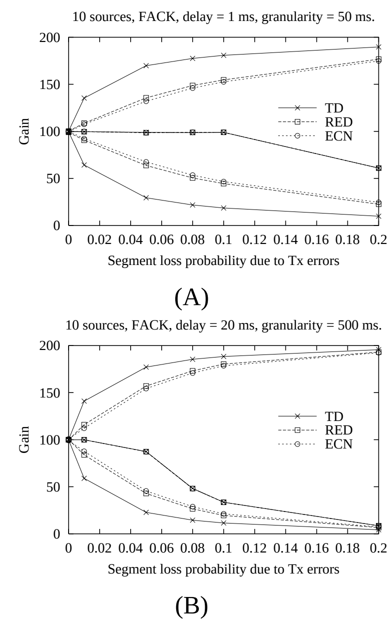

Figure 4: Gain obtained in the lossy and half-lossy scenarios varying the delay and granularity using a tail dropping (TD), RED and ECN-RED router.

RED and ECN advantage over tail dropping routers: the unfairness effect is reduced but the higher the granularity, the lower the benefit of using RED and ECN in terms of fairness.

Reno, New-Reno, SACK and FACK: these TCP implementations are not effective to solve the unfairness problem. Reno is clearly outperformed by the other TCP implementations, New-Reno achieves the same performance as the basic SACK implementation (the one we refer to as SACK), using the more complex TCP implementation (FACK) does not represent a significant improvement over the much simpler New-Reno implementation.

3. Evaluation of HAWAII

3.1. HAWAII Path setup schemes

HAWAII is a domain-based approach for supporting mobility [2] and [11]). The mobile host keeps it network address unchanged while moving within a domain. The Corresponding Hosts (CH) and the Home Agent (HA) do not need to be aware of the host's position within the domain. To reach the Mobile Host (MH), HAWAII uses specialized path setup schemes that update forwarding entries in specific routers. From the four schemes considered in [11], we consider only the forwarding schemes: Multiple Stream Forwarding (MSF) and Single Stream Forwarding (SSF). These schemes forward packets from the Old Base Station (BSO) to the New Base Station (BSN) before being diverted at the cross-over router (i.e. a router where the path from CH to BSO and the path from CH to BSN cross. They are modeled using a queuing network consisting of M/M/1 queues representing the delay experienced in the different routers. By considering the different paths packets follow, it is possible to compute the delay components and the expected number of dropped packets (due to late arrival at the MH when the play-out delay has experienced) or lost packets (when the MH is not reachable any longer from BSO).

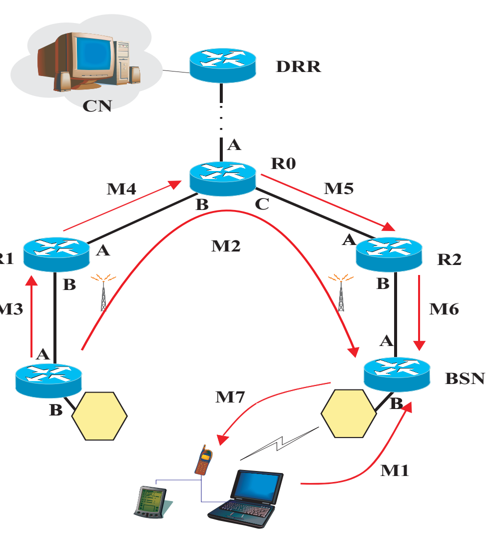

Figure 5: The MSF

3.1.1. *Multiple Stream Forwarding (MSF)*

In order to describe the operation of the MSF path setup scheme, we define the following messages (see Figure 5). At the instant of handoff, BSO loses contact with the MH and at the same time the MH sends a MIP registration message (M1) to BSN. The latter sends a path setup update message M2 to the BSO. When M2 arrives at BSO, the BSO starts to forward all packets with destination MH via router R1, including those packets that arrive after the handoff instant and that were queued in the buffers of the BSO. For that purpose, BSO adds a forwarding entry to its routing table indicating that packets for MH should leave the BSO via interface A.BSO sends the path setup message (M3) to R1, who adds a forwarding entry to its routing table indicating that packets for the MH should leave the R1 via interface A. R1 sends the path setup message (M4) to R0, who adds a forwarding entry indicating that packets for the MH should leave the R0 via interface C. From this instant on, all packets arriving at router R0 are sent directly to BSN. The path setup message continues (M5 and M6) triggering similar actions until it reaches BSN. Remark that MSF can create transient routing loops (for example after BSO has changed its entry to forward packets but before R1 processes M3).

3.1.2. *Single Stream Forwarding (SSF)*

This path setup scheme is similar to the MSF. The difference is that now packets can only be forwarded from BSO and not from intermediate routers, resulting in a single stream of forwarded packets. No loops can occur here, due to the special technique of interface-based forwarding. For more details, we refer to [11].

3.2. A queuing model for path setup schemes

The use of the HAWAII MSF path setup scheme implies that packets with destination the MH will follow a route, depending on their time of arrival at certain routers. In what follows, the packets are divided in classes according to the path they follow. The timing of these classes will be given from the point of view of a packet originating from the corresponding host that arrives at R0. We need the following notations. Let Rx(X) denote the random variable that stands for the time needed for a packet to be processed by router Rx, leaving via interface X. In other words, it denotes the time between the arrival at the router and the departure through output interface X. (Same notations applies to "routers" BSO and BSN). Furthermore, let (Rx,Ry) denote the propagation time on the link between router Rx and router Ry. Let t_{ho} be the time instant of handoff (i.e. the time instant that BSO loses contact with the MH and that M1 is generated in the MH). We assume that the time needed to update the forwarding entries is about equal to the service time of any packet in any router. *Class 1 packets*: These packets arrive at the BSO after the handoff took place and before the new forwarding entry is added at the BSO. Such packets arrive at the R0 after

$$t_0 = t_{ho} - R0(B) + (R0,R1) + R1(B) + (R1,BSO)$$

but before

$$t_1 = t_{ho} + \{(MH,BSN) + BSN(A) + (BSN,R2) + R2(A) + (R2,R0) + R0(B)\}$$

These packets are queued in the BSO until they are forwarded to the MH using the route R0-R1-BSO-[queuing]-R1-R0-R2-BSN.

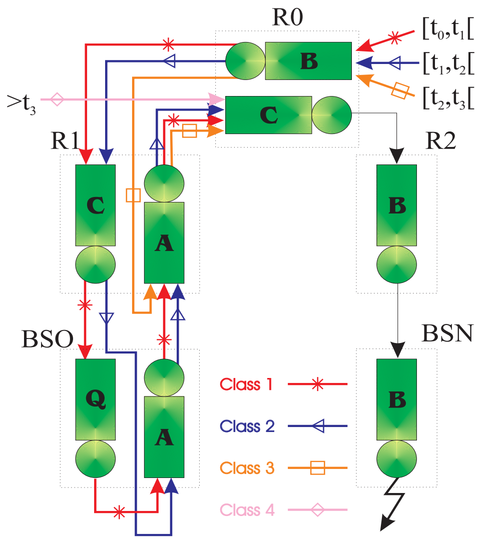

Figure 6 : A queuing model for the MSF scheme

Class 2 packets: These packets arrive at the R1 before M3 causes the adding of the new forwarding entry at R1, and they arrive at the BSO after the new forwarding entry has been added. They reach the MH via the route R0-R1-BSO-R1-R0-R2-BSN. At time instant

$$t_1' = t_1 + (R0,R1) + R1(B) + (R1,BSO) + BSO(A) + BSO,R1) + R1(A)$$

message 3 is processed and router R1 changes its forwarding entries for packets with destination the MH. Hence, the packets that arrive at R0 in the interval [t_1, t_2], with

$$t_2 = t_1' - \{R0(B) + (R0,R1)\}$$

belong to class 2.

Class 3 packets: These packets arrive at the R0 before M4 causes the adding of the new forwarding entry at R0, and they arrive at the R1 after the new forwarding entry has been added. They reach the MH via the route R0-R1-R0-R2-BSN. At time instant

$$t_3 = t_2 + R0(B) + (R0,R1) + (R1,R0) + R0(C)$$

router R0 changes its forwarding entries for packets with destination the MH (based in message M4). Therefore, packets arriving at R0 in the interval [t_2, t_3], belong to class 3.

Class 4 packets: These packets are forwarded to the BSN directly, i.e. via the route R0-R2-BSN. They arrive at R0 after time instant t_3.

Remark: The MSF scheme may result in the creation of routing loops (for example, after BSO has changed its entry to forward packets but before Router 1 processes M3). However these loops exist only for extremely short periods of time. In our model we only consider possible loops for packets that belong to class 1, as the probability for the occurrence of other types of loops seems to be negligible.

Now consider a stream of packets originating from the corresponding node that arrive at R0 with constant inter-arrival time *T*. All routers are modeled as M/M/1 queues. Therefore the probability distribution of the additional time introduced by the MSF (and also SSF) scheme can be computed easily. This leads to the expected number of packets dropped at the MH when a certain playout delay *t* is expired.

3.3. Numerical results

Consider a packet stream with inter-arrival time T = 20ms. The propagation delay between every two routers equals 5 ms and the load of each router is ρ= 0.8.

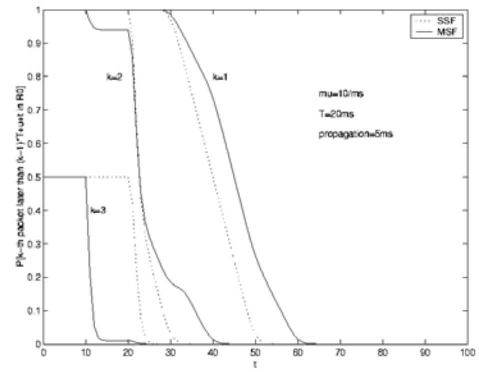

Figure 7:Drop probability first packets

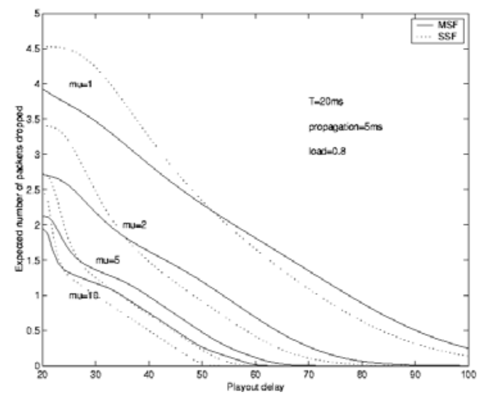

Figure 8 : Expected number of dropped packets

Figure 7 shows the drop probability in the MSF scheme (as computed in Section 3.1.2) for the first three packets as a function of the allowed play-out delay t. The drop probability in the SSF scheme is also depicted (dotted lines). The service rate of each router equals 10/ms. The graphs illustrate that both schemes perform similar, with some minor differences.

For example, the first packet has a slightly smaller drop probability in the case of SSF (due to the occurrence of loops in MSF), while it is the other way around for the 3rd packet. Figure 8 shows the expected number of packets that are dropped in the MSF and SSF schemes for a certain play-out delay, as a function of the service rate μ. It appears that the MSF scheme performs better when a small play-out delay is considered, while the SSF scheme performs slightly better for larger delays.

4. Conclusions

This paper presents results obtained in the European MOEBIUS project on particular issues related to the problem how to meet the global IP concept: the applicability of TCP in a wireless environment and the provisioning of an end-to-end IP solution with the use of HAWAII for micro-mobility.

The impact of a lossy wireless medium on the TCP performance and the benefits of using New-Reno and two implementations of SACK are analysed. Also the use of active queue management schemes like RED and ECN are investigated.

Secondly, a queueing model for two HAWAII setup schemes is defined and used to compute the expected number of packets dropped for a given play-out buffer at the destination.

5. References

[1] C. Perkins, ed., “IP Mobility Support”, IETF RFC 2002, October 1996.

[2] R. Ramjee, T. La-Porta, S. Thuel, K. Varadhan, L. Salgarelli, “IP micro-mobility support using HAWAII”, Internet draft, July 2000.

[3] M. Mathis, J. Mahdavi, S. Floyd, and A. Romanow. "TCP Selective Acknowledgment Options". RFC 2018, October 1996.

[4] S. Floyd and T. Henderson. "The NewReno Modification to TCP's Fast Recovery Algorithm". RFC 2582, April 1999.

[5] K. Fall and S. Floyd. "Simulation-based Comparisons of Tahoe, Reno and SACK TCP". ACM Computer Communication Review, July 1996. ftp://ftp.ee.lbl.gov/papers/sacks.ps.Z

[6] M. Mathis and J. Mahdavi, S. Floyd. "Forward Acknowledgement: Refining TCP Congestion Control". ACM Computer Communication Review, 26(4), October 1996.

[7] B. Braden, D. Clark, J. Crowcroft, B Davie, S. Deering, D. Estrin, S. Floyd, V. Jacobson, G. Minshall, C. Partridge, L. Peterson, K. Ramakrishnan, S. Shenker, J. Wroclawski, and L. Zhang. "Recommendations on Queue Management and Congestion Avoidance in the Internet". RFC 2309, April 1998.

[8] S. Floyd and V. Jacobson. "Random Early Detection Gateways for Congestion Avoidance". IEEE/ACM Transactions on Networking, August 1993.

[9] K. Ramakrishnan and S. Floyd. "A proposal to add Explicit Congestion Notificacion (ECN) to IP". RFC 2481, January 1999.

[10] L. Cerdá, O. Casals, “Study of the TCP Unfairness in a Wireless Environment”, IEEE ICT 2001, June 2001, Bucharest (Rumania).

[11] Ramjee, R., La Porta, T., Thuel, S., Varadhan, K., and Wang, S., HAWAII: a domain based approach for supporting mobility in wide–area wireless networks, Proceedings of International Conference on Network Protocols, ICNP’99.

[12] Blondia, C., De Cleyn. P., Willems, G., and Casals, O., An analytical model for path setup schemes in HAWAII, submitted (http://win-www.uia.ac.be/u/pats)