Handoff initiation in Mobile IPv6

Torben W. Andersen, Anders Lildballe and Brian Nielsen†

Aalborg University
Department of Computer Science
Fredrik Bajersvej 7E
Aalborg Dk-9220, Denmark
{witter | andy | bnielsen}@cs.auc.dk

Abstract

The Mobile IPv6 protocol is expected to become the main carrier of traffic to mobile and wireless nodes. For a mobile node to experience seamless roaming it is essential that the *handoff latency* is reduced to a minimum and that a minimum of packets are lost during the handoff. It is therefore imperative that a handoff is initiated such that network connectivity is maintained for the longest possible period of time.

In this paper we examine the performance of two basic *handoff initiation* algorithms denoted Eager Cell Switching (ECS) [1] and Lazy Cell Switching (LCS) [1] respectively. We investigate how essential protocol parameters influence the handoff latency for ECS and LCS both theoretically and empirically.

Our results include *mathematical models* able to predict the handoff latency and a *thorough empirical study* in a local area testbed running Mobile IPv6. The models are found to conform to the latencies measured in the testbed. Using the mathematical models we propose a new set of optimized protocol parameters reducing the handoff latency. We also find that ECS produces the fastest handoff, but initiates far too many handoffs. Conversely, LCS produces fewer but much slower handoffs in the range of several seconds. We conclude that a good handoff initiation algorithm should be *proactive* and probably needs to take link layer information about signal quality into account.

Keywords

Mobile IPv6, handoff initiation, mathematical models, testbed, FreeBSD, KAME.

1. Introduction

The Internet Protocol (IP) is expected to become the main carrier of traffic to mobile and wireless nodes. This includes ordinary data traffic like http, ftp and email, as well as voice, video and other time sensitive data. To support mobile users, the basic Internet protocols have been extended with protocols (Mobile IP) for intercepting and forwarding packets to a mobile and possibly roaming node. Seamless roaming requires that users and applications do not experience loss of connectivity or any noticeable hick-ups in traffic. However, little is known about the performance of these protocols in an actual network. In particular, it is not understood how different *handoff initiation algorithms* influences essential performance metrics like the packet loss and the duration of a handoff.

The Mobile IPv6 specification [2] contains only a weak specification of handoff initiation algorithms. Two conceptually simple handoff initiation algorithms which have gained considerable interest are Eager Cell Switching (ECS) and Lazy Cell Switching (LCS) [1]. Both operate at the network layer without requiring information from the lower (link) layers. To discover new networks they depend on receiving *router advertisements* broadcasted from access routers. ECS *proactively* initiates a handoff every time a new network prefix is learned. Conversely, LCS acts *reactively* by not initiating a handoff before the primary network is confirmed to be unreachable. When the *lifetime* (which is specified in the broadcasted router advertisements) of the primary network expires, LCS probes the current default router to see if it is still reachable. If not, a handoff to another network is initiated.

The performance of ECS thus depends on the frequency with which access routers are broadcasting router advertisements. Similarly, LCS also depends on the frequency of broadcasted router advertisements, but additionally depends on the lifetime of network prefixes and probing time. The theory and data needed to decide what handoff initiation algorithms to use in what circumstances, how to tune protocol parameters, and where to put optimization efforts, are missing. The work presented here bridges this gap by proposing a *theoretical model* of handoff performance, and by presenting the results of a *thorough empirical study* of handoff performance in a *Mobile IPv6 testbed*. A further novelty is that our work deals with IPv6 for which implementations has only recently become available for empirical studies.

Currently, most work regarding handoff in Mobile IP is dedicated within the field of *handoff execution* using micro mobility schemes. This work aims at reducing the network load and reducing the handoff latency by minimizing the propagation delays of *binding updates* and *binding acknowledgements* such as in HAWAII [3] and in Cellular IP [4]. While reducing handoff execution time is important, we believe that this is insufficient to obtain seamless handoffs, especially when the network topology cannot be controlled as is the case in the Internet. We suggest taking a *proactive* approach to handoff initiation which has the potential of reducing the packet loss to zero. This strategy has also proven successful in existing wireless networks such as GSM [5].

The remainder of this paper is organized as follows. Section 2 introduces the mathematical models for handoff latency. Section 3 presents our experimental results and compares these to the theoretically predicted results. In Section 4 we apply the mathematical models to obtain a set of optimized protocol parameters. Finally, Section 5 concludes the paper.

---

† Author for contact

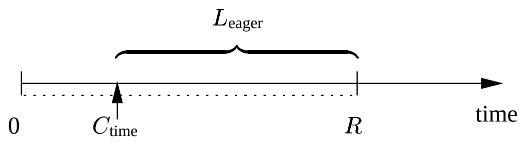

0: Router advertisement sent prior to R
R: First router advertisement received after C_{time}
$C_{\text{time}} \in [0, R]$
$L_{\text{eager}} = R - C_{\text{time}}$

Figure 1: *Model for computing L_{eager}.*

2. Mathematical models

Because our focus is the performance of handoff initiation algorithms, we do not include the propagation delay of *binding updates* and *binding acknowledgements* used in the complete handoff procedure[1]. We therefore define handoff (initiation) latency as the time from the current primary network gets out of range until an event occurs at the mobile node which triggers it to perform a handoff.

The goal is to predict the variation in handoff latency and its average as a function of the primary protocol parameters. The variation in latency is important to real-time sensitive traffic because it also indicates the best and worst case amount of time the mobile node risks being unreachable. We therefore derive a density function from which average and variation handoff latency can be computed.

2.1. Basic definitions

In the mathematical models we assume perfect cell boundaries, i.e., getting within range of a new network coincides with leaving the range of the primary network. According to [6] a router must pick a random delay between each broadcast of an unsolicited router advertisements in order to avoid routers synchronizing. We denote the minimum possible time between two consecutive router advertisements $R_{\min}$ and the maximum time between two consecutive router advertisements $R_{\max}$ meaning that the period between any two router advertisements must be found in the interval [$R_{\min}$,$R_{\max}$]. We make the simplifying assumption that all access routers broadcasts with the same frequency range. The time at which a mobile node enters the range of a new network is denoted by C_{time}.

The lifetime of broadcasted network prefixes is denoted by T_{l}. The time it takes to probe a default router is assumed to be uniformly distributed within the interval [$Q_{\min}$,$Q_{\max}$].

2.2. Eager Cell Switching

For ECS we define the theoretical handoff latency L_{eager} to be the period from getting out of range of the primary network until the reception of a router advertisement from a new network. This is illustrated in Figure 1.

We pursue a density function for L_{eager} from which we can derive the average value of L_{eager}. Because the primary variables C_{time} and R are random but dependent, we start by computing their joint density function $P_{C_{\text{time}},R}(c_{\text{time}}, r)$ which can be expressed as

$$P_{C_{\text{time}},R}(c_{\text{time}}, r) = P_{C_{\text{time}}|R}(c_{\text{time}}|r) \cdot P_R(r) \tag{1}$$

where $P_{C_{\text{time}}|R}(c_{\text{time}}|r)$ is the probability distribution for C_{time} given R. As C_{time} is evenly distributed in the interval [0,R] we can calculate $P_{C_{\text{time}}|R}(c_{\text{time}}|r)$ as

$$P_{C_{\text{time}}|R}(c_{\text{time}}|r) = \frac{1}{r} \cdot 1_{c_{\text{time}} \in [0,r]} \tag{2}$$

where $1_{c_{\text{time}} \in [0,r]}$ is an indicator function with a value of 1 when $c_{\text{time}} \in [0, r]$ and 0 otherwise.

The density function $P_R(r)$ expresses the probability that C_{time} should occur in an interval of size $R = r$. Intuitively, the probability of C_{time} occurring in an interval is proportional to the size of the interval. When the interval size is given as r, the function $f(r) = r \cdot 1_{r \in [R_{\min}, R_{\max}]}$ exhibits this intuitive property. We can obtain the density function $P_R(r)$ as $f(r)$ divided with the area of $f(r)$. This yields

$$\begin{aligned} P_R(r) &= \frac{r}{\int_{R_{\min}}^{R_{\max}} r\, dr} \cdot 1_{r \in [R_{\min}, R_{\max}]} \\ &= \frac{2r}{{R_{\max}}^2 - {R_{\min}}^2} \cdot 1_{r \in [R_{\min}, R_{\max}]} \end{aligned} \tag{3}$$

The joint density function $P_{C_{\text{time}},R}(c_{\text{time}}, r)$ obtained by combining these results becomes

$$\begin{aligned} P_{C_{\text{time}},R}(c_{\text{time}}, r) &= P_{C_{\text{time}}|R}(c_{\text{time}}|r) \cdot P_R(r) \\ &= \frac{2}{{R_{\max}}^2 - {R_{\min}}^2} \cdot 1_{c_{\text{time}} \in [0,r]} \cdot 1_{r \in [R_{\min}, R_{\max}]} \end{aligned} \tag{4}$$

The density function $P_{L_{\text{eager}}}(l)$ for L_{eager} can be obtained by integrating over the joint density function $P_{C_{\text{time}},R}(c_{\text{time}}, r)$ for all possible values of c_{time} and r. As r can be expressed as $r = c_{\text{time}} + l$ (which follows from $l = r - c_{\text{time}}$), we have

$$\begin{aligned} P_{L_{\text{eager}}}(l) &= \int_{-\infty}^{\infty} P_{C_{\text{time}},R}(c_{\text{time}}, c_{\text{time}} + l) dc_{\text{time}} \\ &= \frac{2}{{R_{\max}}^2 - {R_{\min}}^2} \cdot 1_{l \in [0, R_{\max}]} \left[c_{\text{time}}\right]_{\max(0, R_{\min} - l)}^{R_{\max} - l} \\ &= \begin{cases} \frac{2}{R_{\max} + R_{\min}} & \text{if } 0 \le l \le R_{\min} \\ \frac{2(R_{\max} - l)}{{R_{\max}}^2 - {R_{\min}}^2} & \text{if } R_{\min} < l \le R_{\max} \\ 0 & \text{otherwise} \end{cases} \end{aligned} \tag{5}$$

In Figure 2 the density function $P_{L_{\text{eager}}}(l)$ from formula 5 is plotted. We observe that using ECS the handoff latency is bounded by the value of $R_{\max}$ and that there is the highest probability of obtaining handoff latencies in the range [0,$R_{\min}$].

Given the density function $P_{L_{\text{eager}}}(l)$ it is easy to obtain the average value of L_{eager}, denoted by $\overline{L}_{\text{eager}}$, by integrating over the product of $L_{\text{eager}} = l$ and $P_{L_{\text{eager}}}(l)$ for all possible values of L_{eager}:

$$\begin{aligned} \overline{L}_{\text{eager}} &= \int_0^{R_{\max}} l \cdot P_{L_{\text{eager}}}(l) dl \\ &= \frac{{R_{\max}}^3 - {R_{\min}}^3}{3({R_{\max}}^2 - {R_{\min}}^2)} \end{aligned} \tag{6}$$

---

[1] To obtain the total handoff latency, the rountrip time to the home agent must be added to the figures stated in this paper.

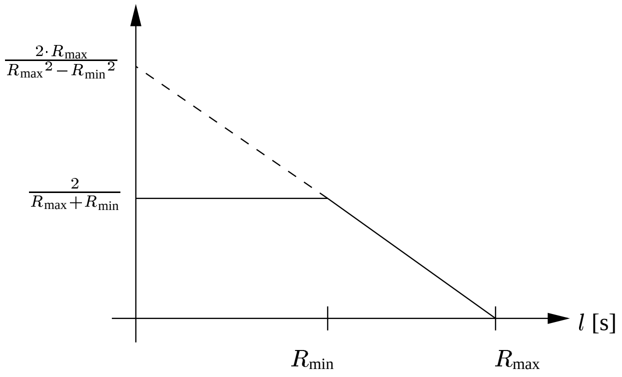

Figure 2: *Density function $P_{L_{eager}}(l)$ for ECS.*

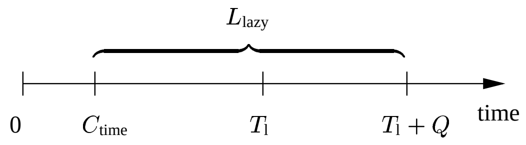

0: Last router advertisement received from primary network

C_{time}: Leaving the range of the primary network

T_{l}: Lifetime of network prefix

Q: Probing time

$L_{\text{lazy}} = T_{\text{l}} - C_{\text{time}} + Q$

Figure 3: *Model for computing L_{lazy}.*

2.3. Lazy Cell Switching

For LCS we define the theoretical handoff latency L_{lazy} as the time from leaving the range of the primary network until concluding that the primary network is unreachable. This occurs when the life time of the primary network has expired and probing it has failed. We use the model depicted in Figure 3 where the last router advertisement from the primary network was received at time 0. We pursue a density function for L_{lazy} from which we can calculate the average value of L_{lazy}.

Intuitively, the handoff latency consists of the remaining lifetime of the primary network plus the probing time Q used to determine that the primary network is unreachable. The remaining lifetime is the lifetime T_{l} of the last router advertisement minus the time the network was reachable, C_{time}. The handoff latency L_{lazy} can thus be expressed as

$$\begin{aligned} L_{\text{lazy}} &= \text{Lifetime remaining of primary network} + Q \quad (7) \\ &= T_{\text{l}} - C_{\text{time}} + Q \end{aligned}$$

Because Q is assumed to be uniformly distributed within the interval $[Q_{\min}, Q_{\max}]$, the density function for Q is

$$P_Q(q) \quad = \quad \frac{1}{Q_{\max} - Q_{\min}} \cdot 1_{q \in [Q_{\min}, Q_{\max}]} \qquad (8)$$

Using a similar line of reasoning and method of calculation as outlined for ECS, the density function for LCS denoted $P_{L_{\text{lazy}}}(l)$ can be computed. That is, we integrate over the joint density function of the involved parameters for all possible values. However, the intermediate calculations are more involved, and we shall here only state the result in Equation 9. Further details on its derivation can be found in [7].

$$\begin{aligned} P_{L_{\text{lazy}}}(l) \quad &= \quad \int_{-\infty}^{\infty} P_{\text{Q}}(l - l_{\text{np}}) \cdot P_{C_{\text{time}}}(T_{\text{l}} - l_{\text{np}}) dl_{\text{np}} \quad (9) \\ & \text{where } P_{C_{\text{time}}}(l) = P_{L_{\text{eager}}}(l) \end{aligned}$$

The average handoff latency $\overline{L}_{\text{lazy}}$ can be computed by integrating over the density function, or more simply, directly from Equation 7 which states that $L_{\text{lazy}} = T_{\text{l}} - C_{\text{time}} + Q$. The average latency $\overline{L}_{\text{lazy}}$ can be computed as the average remaining lifetime plus the average probing time $\overline{Q}$:

$$\begin{aligned} \overline{L}_{\text{lazy}} &= T_{\text{l}} - \overline{C}_{\text{time}} + \overline{Q} \qquad (10) \\ &= T_{\text{l}} - \frac{{R_{\max}}^3 - {R_{\min}}^3}{3({R_{\max}}^2 - {R_{\min}}^2)} + \frac{1}{2}(Q_{\max} - Q_{\min}) \end{aligned}$$

The average remaining lifetime of the primary network is the lifetime of the last received router advertisement minus the average amount of time that the primary network was reachable. On average the primary network is reachable for an amount of time corresponding to the time used to discover a new network. This corresponds exactly to the average handoff latency for ECS. Thus, $\overline{C}_{\text{time}}$ equals $\overline{L}_{\text{eager}}$ calculated in Equation 6.

3. Experimental Results

In Section 2 we presented mathematical models needed to compute the handoff latency as a function of essential protocol parameters. In this section we present the design of a Mobile IPv6 testbed and compare the theoretically predicted handoff latency with the handoff latency experienced by a mobile node in the Mobile IPv6 testbed.

3.1. The Mobile IPv6 testbed

The testbed is depicted in Figure 4 and consist of four nodes; three routers and one host. The three routers, *iridium*, *platin* and *nikkel*, are assigned an IPv6 prefix for each network device. The mobile node, *lantan*, is manually assigned an IPv6 address at the fec0:0:0:1::/64 network, its home network. When *lantan* is not at its home network, it uses *stateless auto-configuration* to obtain an IPv6 address as its care-of address. The home agent is located at *iridium* which also hosts an application corresponding with an application at the mobile node. The mobile node can roam between the two access routers *platin* and *nikkel*. The link media used in the experiments reported here are standard 802.3 10 Mbit/s Ethernet devices. However, the testbed also runs 802.11b 11 Mbit/s Wireless LAN connections. The connection between the access routers *platin* and *nikkel* allows them to coordinate on whom should offer access to the mobile node. It also allows experiments with route optimizations because it offers an alternative path to the mobile node.

All nodes run FreeBSD version 4.1 [9]. On top of FreeBSD the KAME package [10] is installed. The KAME package includes Mobile IPv6 support and IPsec support. The KAME package installed is the weekly snap-release of 25/9-2000. A snap-release is the newest version of the package and may include functionality that is still under development and is not fully tested. The Mobile IPv6 code supplied with KAME is an example of such functionality. The Mobile IPv6 implementation included in KAME can be configured to use either the ECS or the LCS handoff initiation algorithm.

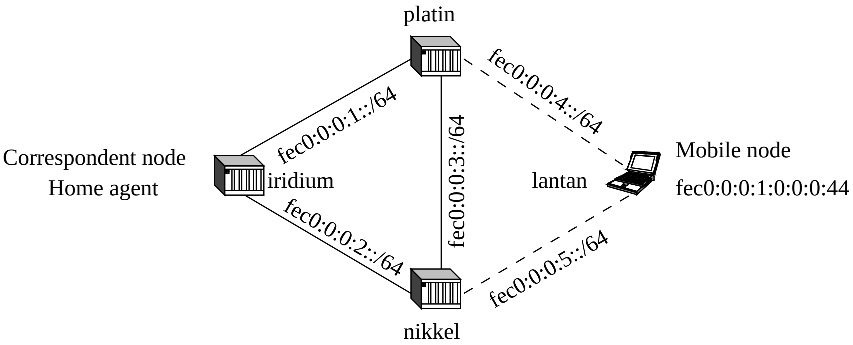

Figure 4: *Testbed with IPv6 addresses and network prefixes.*

3.2. Experimental approach

Two different scenarios have been emulated in the testbed. In the *no network overlap scenario* perfect cell boundaries are assumed. A mobile node moving out of the range of one network therefore coincides with the mobile node moving within range of another network. This scenario corresponds to the one for which mathematical models were derived in Section 2. In the *network overlap scenario* cell boundaries are overlapping between the two networks and the mobile node is always able to reach at least one network. This scenario was applied to investigate whether ECS was able to avoid packet loss when able to receive and send packets via two networks at the same time. LCS behaves identically in both scenarios.

The mobile node moving in and out of the range of a network has been emulated by operating a firewall at the access routers. When the firewall is enabled, the mobile node is not able to receive or send any traffic through that particular access router. Accordingly, when the firewall is disabled all traffic is allowed to pass to and from the mobile node.

To determine the handoff latency in an experiment we applied the following method:

- UDP packets are send from the correspondent node to the mobile node. Each packet contains a send timestamp and a sequence number. The UDP packets are send with a random interval between 95 ms and 105 ms. The interval is randomized to make sure the network does not adjust itself to any particular sending frequency.
- The UDP packets are received and time stamped at the mobile node. The sequence number, the send and receive timestamp are stored upon reception of a packet as an entry in a log file.
- A handoff is registered from packets missing in the log file. We compute the measured handoff latency by multiplying the number of lost packets with the average period between sending packets (0.1 seconds). The precision of the measured latency is thus ± 0.1 seconds. If a handoff is performed without losing packets it will therefore not be registered. Both the average and the frequency distribution of handoff latencies can be computed by inspecting the log.

By reducing the interval between UDP packets (increasing frequency) a higher accuracy will be obtained, and the measured latencies will approach the theoretical latencies defined in in Section 2. The interval of 95 ms to 105 ms was chosen to avoid too many UDP packets being sent. Due to a memory leakage in the KAME Mobile IPv6 software only a limited number of packets can be sent from a correspondent node before it crashes. In the experiments presented in this paper we were able to perform 300 to 400 handoffs in sequence before the correspondent node crashed[2].

3.3. Overview of performed experiments

Using the Mobile IPv6 testbed we have performed the following experiments

Default configuration: In this experiment the router advertisement interval and network prefix lifetime is set as recommended in [2]. This means a router advervisement interval randomly chosen between 0.5 and 1.5 seconds and a lifetime of 4 seconds. The purpose of this experiment is to reveal handoff latency using the default configuration.

Latency as a function of router advertisement interval: Handoff performance is measured for different router advertisements intervals, but with an identical network prefix lifetime. The purpose of this experiment is to investigate how the interval between sending router advertisements affects the handoff latency.

Latency as a function of network prefix lifetime: In this experiment the handoff latency is measured for different network prefix lifetimes, but with a fixed range for the intervals between sending router advertisements. The purpose is to investigate how the lifetime of router advertisements affects the handoff latency.

All experiments have been performed using both the network overlap and the no network overlap scenario.

In this section we present a selection of the empirically obtained results. The full set of results is given in [7]. In all plots we have added the theoretically predicted handoff latency such that the theoretical and empirical results can easily be compared.

3.4. Default Settings

First we compare the theoretically predicted probability distributions with the measured frequency distributions using the default configuration of access routers in the no overlap scenario. Next we compare their performance numerically.

The histogram in Figure 5 depicts the frequency distribution of the experimentally measured handoff latencies for ECS. The

---

[2] Further confidence in the mathematical models have been obtained by implementing a simulator in JAVA. Using this simulator we have confirmed the theoretically predicted density functions for a range of configurations for both ECS and LCS.

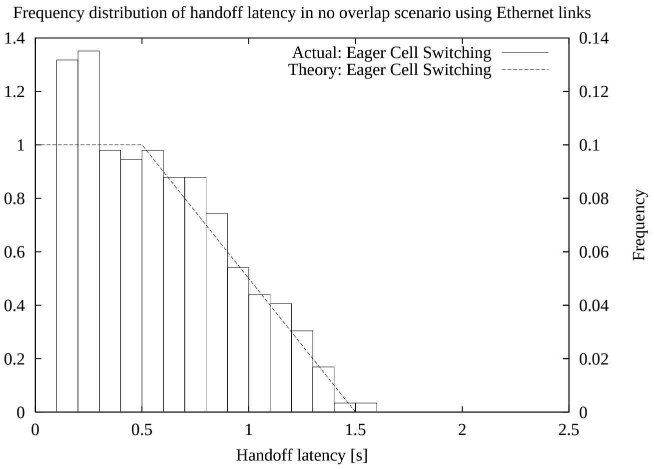

Figure 5: *Frequency distribution for handoff latency using ECS and default configuration in the no network overlap scenario.*

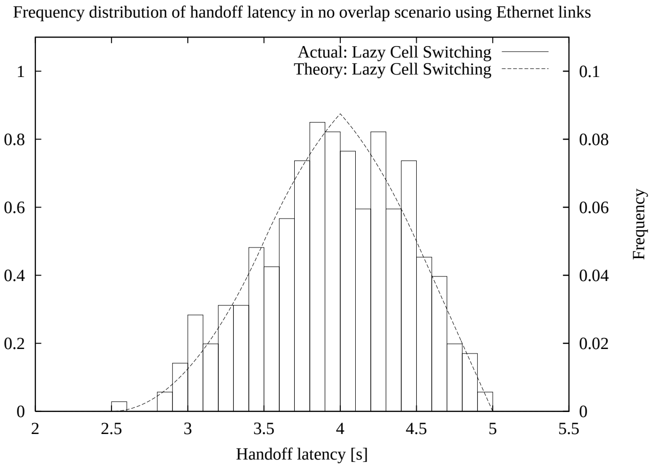

Figure 6: *Frequency distribution for handoff latency using LCS and default configuration in the no network overlap scenario.*

continuous line shows the theoretically predicted density function. The observed handoff latencies lie in the range 0.1 to 1.5 seconds. We note also that no handoff latencies in the interval [0,0.1] are present. This is caused by the experimental setup in which the precision is limited by the frequency of packets from the corresponding node, i.e., latencies below 0.1 seconds cannot be observed. Instead these handoff latencies are recorded in the [0.1,0.2] and the [0.2,0.3] intervals. We conclude that the experimental results for ECS conforms well to those predicted by the mathematical models.

Figure 6 depicts the histogram obtained for LCS. Here the observed latencies range from 2.5 seconds to 5 seconds with most values centered around 4 seconds. It can be seen that LCS is generally unable to avoid packet loss as it does not initiate a handoff before after the primary network has become unavailable. We also conclude that the empirically obtained frequency distribution conforms well to the density function predicted by the mathematical model.

The comparison of ECS and LCS is summarized in Figure 7. Here it is seen that, when no overlap between network ranges exists, ECS yields an average handoff latency of 0.54 seconds which corresponds to the time it takes to discover the new network. LCS yields a much worse latency with an average of 3.97 seconds. Also best and worst case values are higher. We conclude that with respect to handoff latency ECS outperforms LCS. Furthermore, we have found from the experiments with overlapping networks that ECS is able to avoid packet loss altogether during a handoff, provided that a sufficient overlap between network ranges exists.

Handoff strategy	**L [s] (Theory)**			**Latency [s]**		
	Avg	**Min**	**Max**	**Avg**	**Min**	**Max**
Eager	0.54	0	1.5	0.54	0.10	1.52
Lazy	3.95	2.5	5.0	3.97	2.54	4.97

Figure 7: *Summary of expected and actual results for the no network overlap setup using default router configuration.*

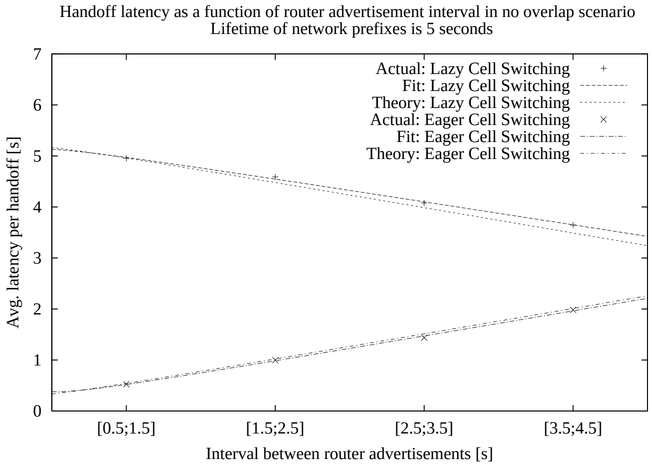

Figure 8: *Handoff latency as a function of router advertisement interval for ECS and LCS in the no network overlap scenario.*

However, a disadvantage of ECS is that it always performs a handoff when discovering a new network, whether or not this offers stable connectivity. Consequently, we expect ECS to perform unnecessarily many handoffs resulting in an increased network load and loss of connectivity. We conclude that both ECS and LCS have serious performance lacks, but that the performance of ECS indicates that proactive handoff initiation has the potential to avoid packet loss.

3.5. Varying Advertisement frequency

One of the primary protocol parameters is the frequency of router advertisements. Its effect on handoff latency is shown in Figure 8 which plots handoff latency as a function of the interval between broadcasting router advertisements.

ECS behaves like what we intuitively would expect: A higher frequency implies that networks are discovered sooner, which again implies faster handoffs. Surprisingly however, we observe that the LCS latency is actually decreasing when the interval between broadcasting router advertisements is increased. The explanation for this is that the lifetime is fixed at a constant value of 5 seconds in this experiment. This result thus indicates that the handoff latency for LCS can be minimized by configuring access routers with a prefix lifetime very close to the maximum interval between broadcasting router advertisements.

Handoff strategy	Router configuration			L [s] (Theory)		
	R_{min}	R_{max}	T_{l}	Avg	Min	Max
Eager	0.5	1.5	4	0.54	0	1.5
Lazy	0.5	1.5	4	3.95	2.5	5.0
Eager	0.9	1.1	1.1	0.50	0	1.1
Lazy	0.9	1.1	1.1	1.17	1.0	2.1
Lazy	0	2	2	1.83	1.0	3.0

Figure 9: *Mathematically predicted handoff latency in the no network overlap scenario for different configurations.*

4. Optimizing protocol configuration

In this section we show that the default configuration of access routers proposed in [2] does not result in optimal handoff performance for neither ECS or LCS. In [2] it is suggested that a router should broadcast unsolicited router advertisements distanced by a random period chosen from the interval [0.5,1.5]. This yields an average network load of one router advertisement every second. We propose a new set of parameters which reduce handoff latency *without* increasing the network load.

Using the same network load as the suggested rate of an average of one advertisement per second, the average ECS latency $\overline{L}_{\text{eager}}$ can be minimized by adjusting R_{min} and R_{max} in Equation 6 subject to the constraint that the sum of R_{min} and R_{max} must equal 2. Close inspection reveals that ECS performs best when R_{min} and R_{max} are configured with values as close together as possible. Optimal performance for ECS can therefore be obtained when both R_{min} and R_{max} are set to a value of 1.

The same method applied to LCS reveals that LCS performs better when R_{min} and R_{max} are configured with values far from each other. For LCS optimal performance can therefore be obtained by configuring R_{min} to a value of 0 and R_{max} to a value of 2. This is in direct contradiction to the optimal configuration for ECS. However, for LCS we have observed that the dominating factor for the handoff latency is the lifetime of broadcasted network prefixes. As this lifetime cannot be configured to be lower than the value of R_{max}, a reduction of R_{max} (which is the case when R_{min} and R_{max} is configured to have values close to each other) can also benefit the performance of LCS, if the lifetime is configured close to the value of R_{max}.

In Figure 9 theoretical values of handoff latency for three different configurations of access routers are shown. The probing time for LCS is assumed to be in the interval [0,1]. We observe that both ECS and LCS performs better with the proposed configuration of access routers with $R_{\text{min}} = 0.9$ and $R_{\text{max}} = 1.1$. For ECS the average handoff latency is reduced from 0.54 seconds to 0.50 seconds and the worst case handoff latency is reduced from 1.5 to 1.1 second. Similarly, for LCS the average handoff latency is reduced from 3.95 to 1.17 seconds and the worst case handoff latency is reduced from 5.0 to 2.1 seconds. Our proposed settings thus simultaneously improves on both average, best, and worst-case.

Observe from the last row in Figure 9, that for LCS the advantage of configuring R_{min} and R_{max} far from each other is out-weighted by the fact that the lifetime T_{l} has to be configured at a higher value.

The performance of the new settings have been tried out in the testbed. The experiment confirmed the theoretically predicted values [7].

An alternative to reducing the lifetime of router advertisement messages is to exploit the *advertisement interval option* in router advertisements proposed in [2]. This option contains the maximum time (R_{max}) between router advertisements that mobile nodes should expect. This would allow a mobile node to probe its default router if no router advertisement has been received for a period corresponding to the value of R_{max}. This in effect forces LCS to become more proactive.

5. Conclusions and future work

In this paper we have presented mathematical models for the Eager Cell Switching (ECS) and the Lazy Cell Switching (LCS) handoff initiation algorithms able to predict handoff performance. Using a testbed installed with FreeBSD 4.1 and the KAME Mobile IPv6 software, these models were shown to accurately reflect the handoff latency experienced by an actual roaming node. By applying the mathematical models we showed how the Mobile IPv6 protocol configuration can be optimized to reduce the handoff latency without increasing network load due to router advertisements.

We have argued that both ECS and LCS have serious performance lacks, but that ECS has the potential to avoid packet loss. We are currently in the process of deploying access routers using Wireless LAN as the link media in a building wide experiment. Initial results [7] indicate that ECS does not perform well in an actual wireless network because sporadic router advertisements from new but unstable networks barely within reach causes ECS to handoff to the new network.

We are therefore working on the implementation of a more advanced handoff initiation algorithm. Our work indicates that this algorithm should be proactive, but should be better informed about the quality of the new network before deciding on performing a handoff. Our algorithm takes link layer information about signal quality into account as well as its throughput and price. We expect to report on this work in a future paper.

6. References

[1] Charles E. Perkins, *'Mobile IP - Design Principles and Practices'*, Addison-Wesley, 1998, ISBN 0-13-850505-5.

[2] David B. Johnson and Charles Perkins, *'Mobility Support in IPv6'*, Internet-Draft, draft-ietf-mobileip-ipv6-13.txt, November 2000. Work in progress.

[3] R. Ramjee, T. La Porta, S. Thuel, K. Varadhan and S.Y. Wang, *'HAWAII: A Domain-Based Approach for Supporting Mobility in Wide-area Wireless networks'*, 1999.

[4] Zach D. Shelby, Dionisios Gatzounas, Andrew Campbell and Chieh-Yih Wan *'Mobility Support in IPv6'*, Internet-Draft, draft-shelby-seamoby-cellularipv6-00.txt, November 2000. Work in progress.

[5] Michel Mouly, Marie-Bernadette Pautet *'The GSM System for Mobile Communications'*, 1992, ISBN 2-9507190-0-7.

[6] S. Deering, *'ICMP Router Discovery Messages'*, RFC 1256, September 1991.

[7] Torben Wittrup Andersen and Anders Lildballe, *'Seamless Handoff in Mobile IPv6'*, Masters Thesis, Department of Computer Science, Aalborg University, Denmark, 2001.

[8] ORiNOCO homepage, *'ORiNOCO PC Card (Silver)'*. Available at http://www.orinocowireless.com/.

[9] FreeBSD homepage. Available at http://www.freebsd.org.

[10] KAME homepage. Available at http://www.kame.net.