TKN Telecommunication Networks Group

Technical University Berlin

Telecommunication Networks Group

Performance Evaluation of a QoS-Aware Handover Mechanism

S. Sroka and H. Karl

sroka,karl@ft.ee.tu-berlin.de

Berlin, March 2003

TKN Technical Report TKN-03-006

TKN Technical Reports Series

Editor: Prof. Dr.-Ing. Adam Wolisz

Abstract

Mobile communication is increasingly oriented towards the usage of all IP networks as fixed-network components. An open problem is how to provide QoS guarantees that are competitive with that of existing cellular networks. In particular, an appropriate handover support is missing: a handover should not be performed to a base station that is not able to support the desired QoS. This paper presents a performance evaluation of a mechanism that integrates QoS support and handover mechanisms in IP networks such that a handover is conditionalized upon availability of sufficient QoS resources. We show that this scheme is able to efficiently provide such conditional handover support, that it is competitive with standard Hierarchical Mobile IPv6 regarding the amount of carried traffic, and that it outperforms Hierarchical Mobile IPv6 in arranging traffic within the network, resulting in a superior efficiency especially for mixed QoS/non-QoS traffic.

Contents

Abstract

Mobile communication is increasingly oriented towards the usage of all IP networks as fixed-network components. An open problem is how to provide QoS guarantees that are competitive with that of existing cellular networks. In particular, an appropriate handover support is missing: a handover should not be performed to a base station that is not able to support the desired QoS. This paper presents a performance evaluation of a mechanism that integrates QoS support and handover mechanisms in IP networks such that a handover is conditionalized upon availability of sufficient QoS resources. We show that this scheme is able to efficiently provide such conditional handover support, that it is competitive with standard Hierarchical Mobile IPv6 regarding the amount of carried traffic, and that it outperforms Hierarchical Mobile IPv6 in arranging traffic within the network, resulting in a superior efficiency especially for mixed QoS/non-QoS traffic.

Chapter 1

Introduction

Public mobile communication systems like GSM or its successors are currently based on expensive, special-purpose hard- and software in their fixed communication network. With growing cost pressure, a tendency to replace custom-made systems with standard, off-the-shelf equipment is attractive. This tendency resulted in the concept of "all IP" mobile communication systems: IP protocols and hardware are easily and cheaply available from the mass market for data communication.

However, it is not clear whether this IP-based protocol family will be able to fulfill all the requirements that are currently placed on mobile communication systems. In particular, current GSM-type solutions provide a seamless mobility support and even guarantee certain levels of Quality of Service while a user moves around (as far as "guarantees" are possible over a wireless channel). Providing such guarantees using IP-based networks is still a challenging task: Much research has been done on QoS in fixed IP networks, but how to support QoS guarantees in a *mobile* IP network is not clear; even the precise notion of such a requirement is not entirely well defined.

Starting from user expectations, such a requirement would be that a once established communication session, e.g. a phone call, should be continued with a negotiated quality as long as this is possible. A user does not want to be bothered with variations in the quality just because he moves around; least of all would he be willing to pay for a call that is not of an acceptable quality.

The technical questions that arise from this requirement are manifold, e.g., related to resource allocation on the radio link. In this paper, we concentrate on the repercussions of this requirement for the IP network: Consider a cellular system, consisting of a set of base stations, interconnected by an IP network. A user of such a system could negotiate a certain "contract" with the system about the quality of a communication session; such a contract could for example including prescriptions about a preferred and a minimally acceptable quality of a session. The crucial point is then the handover from one base station to another: In order to fulfill this contract, the base station has to provide sufficient resources to the user, or must not accept the call at all if it would only be possible to do so with non-satisfactory quality (if not even the minimal quality could be provided for).

If, during a handover, there is only a choice between the old and a single new base station, this approach would only guarantee that an established session does not degrade in quality, but if the new base station is overloaded, there is no way but to drop the session. In fact, as the areas far from a base station are likely to be the most crucial ones regarding quality, it is to be expected that coverage areas of multiple base stations will overlap, providing a choice among several base stations from which the best one should be selected. Moreover, this concept also extends to base stations that belong to different types of systems (e.g., classical cellular or WLAN-based systems), jointly forming a heterogeneous access system. Such a heterogeneous system with so-called vertical handovers is likely to considerably raise the probability that more than a single new candidate base station is available.

In an all-IP system, ensuring the data flow to a newly selected base station is already solved by existing mobility solutions, e.g., Mobile IP. The remaining question is *how* to select a best suited base station; how to signal the communication needs using IP protocol means; how to ensure that the mobility solution is well integrated with this signaling process to guarantee that a handover to a new base station is only performed when sufficient resources are available? In brief: how to make an IP-based mobility solution QoS-aware?

In previous work, we have described one solution for this problem: Add sufficient QoS information to Mobile IP's binding update procedure and conditionalize the execution of a handover upon this information. This mechanism has been described as a "QoS-conditionalized Binding Update" extension to Hierarchical Mobile IPv6 [2]. This work, however, has only described the mechanism as such, without providing quantitative results. The contribution of the present paper is to provide these results, using a simulation-based performance evaluation. We will compare our mechanism with a QoS-unaware mobility solution and introduce some additional refinements to it. We will show, using these results, that it is indeed feasible to implement a notion of QoS-aware handovers in a fashion that matches intuitive user expectations without over-taxing the network with signaling traffic or reducing the network's traffic capacity in an undue fashion. The QoS-conditionalized Binding Update mechanism turns out to be efficient, intuitive, and simple to implement.

The remainder of this paper is structured as follows. Section 2 will briefly summarize Mobile IPv6 and Hierarchical Mobile IPv6 to the degree that is necessary here and describe other approaches to QoS support in mobile networks. Section 3 describes our QoS-conditionalized Binding Update approach in more detail. In Section 4, the evaluation scenarios and the structure of the simulation are discussed; Section 5 contains the performance evaluation results. Finally, Section 6 concludes the paper and discusses options for future work.

Chapter 2

Related work

2.1 Mobile IPv6 and Hierarchical Mobile IPv6

Mobile IPv6 [3] is probably the best known mobility mechanism for IPv6 networks. It solves the basic IP mobility problem of a node's address representing both identity and location: it keeps the node's address constant to the "outside world" and assigns a temporary address to a mobile node (MN), matching the node's current location. Additionally, a so-called "home agent" (HA) is introduced, located within the mobile node's home network, which can intercept packets destined to the mobile node's home address and tunnel them towards the mobile node. As the mobile node has to inform the home agent of its current temporary address, the commonly claimed drawbacks of Mobile IPv6 are its high signaling load and long handover latencies.

To overcome these drawbacks for local handovers, Hierarchical MIPv6 (HMIPv6) [7] has been developed. HMIPv6 introduces a new entity, the Mobility Anchor Point (MAP) which works as a proxy for the HA in a foreign network as shown in Fig 2.1. When an MN moves into a network controlled by a new MAP, it is assigned two new CoAs: a Regional CoA on the MAP's subnet (RCoA) and an on-link address (LCoA), which is the same as used for MIPv6.

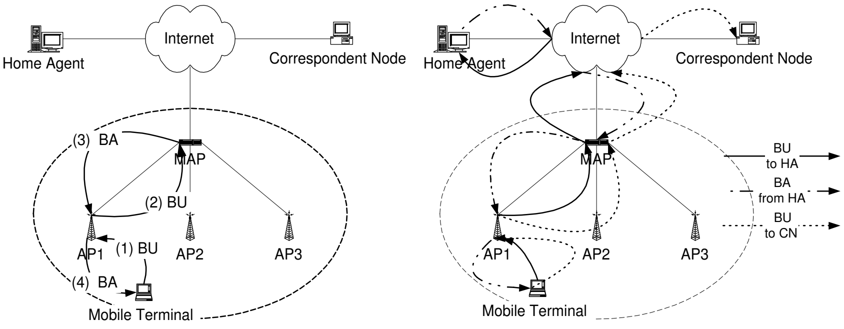

Figure 2.1: Process operations for HMIPv6. Local handover on the left, additional messages for global handover on the right hand side.

The MN then sends a binding update (BU) to the MAP specifying its RCoA in the Home Address field, using its LCoA as its source address. After receiving a binding acknowledgment (BA) from the MAP it also requests a BU (RCoA, Home Address) from its home agent and correspondent nodes

(CNs). If the MN moves locally, only the LCoA is changed and a registration packet is sent to the MAP. Since the RCoA is unchanged the HA and CN do not need to be informed about the new location. While this enhancement is more efficient for mobility support, it does not provide QoS support and QoS-aware authorization for mobile users.

In practical deployment, the MAP would usually be located in a gateway of an administrative domain; such an arrangement will be assumed in the remainder of this paper.

2.2 QoS Support for Mobile IP

In order to provide end-to-end QoS, IntServ and DiffServ have been designed; RSVP [9] has been developed as a signaling protocol. However, RSVP is difficult to use in mobility scenarios, e.g., due to its inability to build proper reservation state over the tunnel between the HA and the visited network.

Shen et al. [6] extended RSVP by proposing a "flow transparency" concept, i.e., identifying the flow address (source and destination address) with the MN's home address, regardless of the change of the MN's CoA. Paskalis et al. [5] proposed putting a "mobility proxy" at the edge of the access domain (e.g., a MAP domain) on behalf of RSVP messages: inside and outside the access domain, LCoA and RCoA of the MN will be used to identify the same session. The mobility proxy will either change the session information in Path/Resv messages accordingly and forward it (inter-domain handover), or generate a Path toward the mobile node/respond with a Resv message upon receipt of a Path message from the mobile node (intra-domain handover).

These RSVP-based approaches still have problems regarding latency for QoS signaling ($\geq$ two round-trip times) and signaling overhead. Chaskar and Koodli [1] proposed a QoS option that allows to include QoS-related data within existing mobility management messages. This approach allows to trigger a one-pass check for the required resources along the path toward the destination node in MIPv6. The QoS option is a hop-by-hop IPv6 header option containing one or more so-called QoS objects, which code the QoS desired by a flow. The QoS option is attached to the Mobile IPv6 binding update message. Hence, resource availability can be checked while the handover process is in progress. This approach allows a one-pass check for the required resources along the path toward the destination node, simultaneously with the binding update. Its drawback is that the MN does not receive any feedback on whether the desired QoS is available at all along the new path and that a handover always takes place, irrespective of available QoS resources.

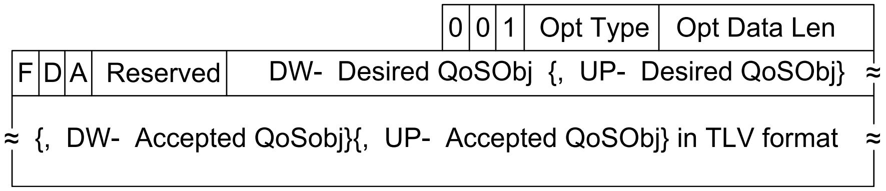

Figure 2.2: Structure of a QoS option

This problem is solved in [2], where a handover is conditionalized upon the ability of providing the required QoS along the new path. If a router along this path cannot provide adequate QoS, a handover request will be returned as unsuccessful and another QoS-conditionalized handover process could be initiated afterwards (if there is yet another potential path available). Furthermore, the QoS object does not need to be transmitted by every host, but could be obtained by the new access router from the old AR, e.g., by way of Context Transfer [4].

Chapter 3

QoS-conditionalized Handover

3.1 Overview

The QoS-conditionalized handover approach is based on HMIPv6[1] in order to share its advantages regarding short handover latency and low signaling load. Additionally, this approach allows to limit the QoS negotiations to the new path.

The goal of this approach is to maintain the established QoS flows during handover, ensuring that the QoS-requirements are fulfilled if the handover actually takes place. The approach assumes that the network consists of routers which may also be responsible for the maintenance of QoS resources, in which case we call them QoS entities. Within an access network are at least the APs and the MAP QoS entities. These entities receive QoS requests for a particular link and check if the QoS requirements can be satisfied. Typical resources are bandwidth of the link, buffer space of the router or CPU resources of the router.

The QoS entity has three opportunities to react to a QoS request: accept the request and reserve the required resources, reject the request, or offer a lower QoS level.

3.2 Protocol Operations

The QoS-conditionalized handover approach covers only local handover if the QoS flows between the CN and the MN are already established. For global handover an end to end renegotiation for the QoS flows is necessary.

For the signaling of the QoS requirements the QoS option is used (see Section 2.2). This option is attached to the IP header of the signaling packets. Thus the QoS information can be processed at the same time as the binding update. If the MN has QoS-sensitive connections to different CNs it has to attach one QoS option for each CN.

The message which carries the binding update and the QoS option is referred to as BU+QoS and the message which carries the binding acknowledgment is called BA+QoS.

If a MN wants to change its access point (AP) to the network it sends a BU+QoS to the MAP via the desired AP. How a MN detects the necessity to change its AP is out of scope. Each entity between the MN and the MAP (including the MAP) will pass the QoS requirements represented by the QoS option to internal QoS mechanisms and check its resource availability. If resources are

---

[1] More precisely, on HMIPv6's basic mode.

available locally, they are reserved and the message will be forwarded along its route. If resources are not available, negative feedback will be provided to the MN by a BA+QoS (which is marked as “failed”). If a BU+QoS message has reached the MAP and passed the local QoS test as well, the handover will take place (the binding cache in the MAP is updated to reflect the new LCoA) and a positive BA+QoS message is returned to the MN. Otherwise, no handover is performed and a negative BA+QoS message is returned to the MN. When observing a negative BA+QoS message, intermediate QoS entities can release reservations that could not be granted further upstream.

In order to allow both upstream and downstream QoS requirements to be considered, this approach assumes that packets for both directions follow the same route.

Figure 3.1 shows on example of a QoS-conditionalized handover procedure. At the beginning the MN has established QoS connections to the CN via AP 2 and intermediate router 1. Now it recognizes that it would be useful to change the access point to AP 3. Therefore, it sends a BU+QoS message to the potential new AP 3 (step 1), using an LCoA that reflects its (pending) attachment to AP 3. On the path to the MAP, there could be additional “intermediate” routers which have to accept the QoS requirements, too. If an intermediate router is not capable of providing requested QoS it sets the F-Flag in the QoS option to indicate that the QoS request can not be granted (2). Whether the QoS test fails or succeeds, the BU+QoS has to be forward to the MAP in order to spare the routers the need to implement the functions to generate BAs (3).[2] In addition, the MAP gets informations about failed BUs that can be used to optimize resource handling in the network. If the MAP receives a negative BU+QoS it generates negative BA-QoS and sends it back to the MN via AP 3 (4). All QoS entities that receive a negative BA-QoS can release the reservations made for this MN’s LCoA (5,6). The old path via AP 2 has been untouched and can be further used by the MN.

If the MN receives a negative BA+QoS it has three opportunities: try the path again with lower level of QoS, choose an other path if any, or stay with the old AP if it is still reachable. In Figure 3.1 the MN has also a link-layer connection to AP 1. So it can try to send a BU+QoS message to this AP (7). This AP can accommodate the QoS requirements and forward the message to the MAP (8). When all QoS entities can met the requirements the MAP performs the handover and sends a positive BA+QoS to the MN (10, 11, 12). When the MN receives a positive BA+QoS it can be sure that the new path can satisfy all QoS-sensitive connections and can immediately use the new path.

If the MN can not find a route that satisfy the QoS requirements no handover takes place.

3.3 Releasing the Reservations

The QoS-conditionalized binding update approach has to take care that all reservations are released when the MN does no longer use a given path. To enable this the protocol can use a combination of a soft-state reservation and an explicit release message.

When the MAP performs a handover it sends a release message via the old path (using the old LCoA), thus all routers in this path can release the reservations for this LCoA. If the MN lost the connection to the AP without performing an handover (it made a global handover or has no link-layer connection) the protocol can not use explicit release messages—it has to give the reservations a lifetime so that the reservations can time out. (Alternatively, inter-domain release messages can in principle be introduced, but that would complicate the approach.)

---

[2] A router capable of generating BA could of course immediately return a failed BA+QoS.

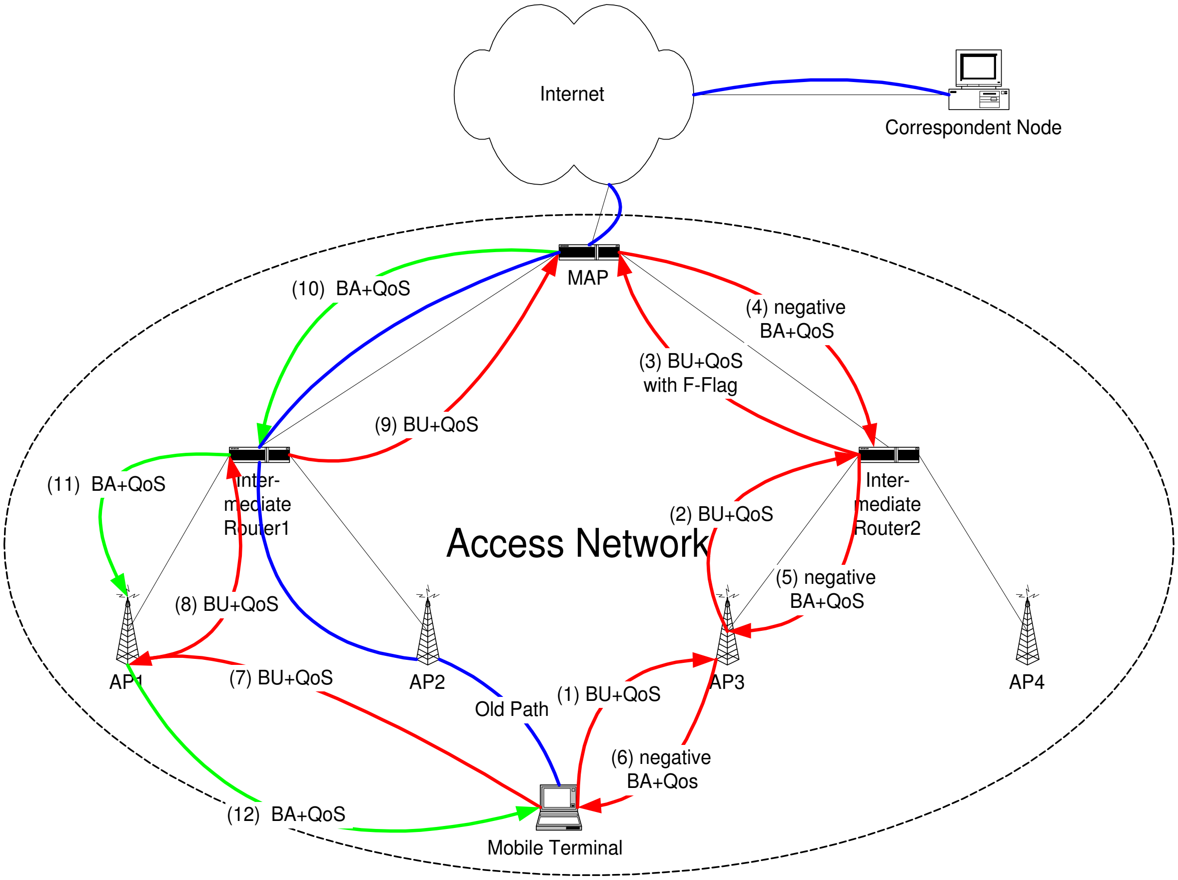

Figure 3.1: Procedure overview of QoS-conditionalized Binding Update

If the network has not enough resources to satisfy all QoS requests, these algorithms have problems. Because the release message is only sent if the BU was successful the old reservation is still valid but superfluously occupies resources. Thus, more MNs can not perform handovers and more resources are needlessly occupied. To avoid this the reservation lifetime could be short but this causes more signaling load. An other solution is to attach a release flag to the BU if the MN has lost the link-layer connection to the old AP. So the MAP can send a release message although the handover was not successful. As it turns out, introducing this explicit release information is crucial to ensure that the network capacity is sustained.

Chapter 4

Evaluation

4.1 Scenarios

4.1.1 Topology

The simulation assumes an access network that allows MNs to have access via radio interfaces only. The principle configuration of the simulation is shown in Figure 4.1. The simulation area is a square with a dimension of 300 to 300 units of length (UL, e.g., meters). This area is covered by nine APs with partially overlapping areas of radio coverage, depending of the range of their radio interfaces. On the top of the hierarchy of the access network resides a MAP that works as gateway router to the Internet, too.

The bandwidth of the radio link is 11 MBit/s and the links between AP and MAP have a bandwidth of 10 MBit/s. All other links have a bandwidth of 100 MBit/s. In the simulation all MNs share the same HA and CN.

4.1.2 Movement model

The MNs move via straight lines across the simulation area using a “Random Waypoint” model. The MN chooses a random destination point. When it arrives at this point it chooses a new destination point. The speed is constant and equal for all MNs. The speed is determined by two parameters: the update frequency and the distance per step.

For the move detection the MNs use the periodic router advertisements sent from the APs. If a MN misses three consecutive advertisements from the current AP it assumes that it has lost the link-layer connection to this AP and has to choose a new one.

4.1.3 Load model

The simulation uses a simple connection-less protocol for the application. The application periodically sends packets with a constant length to the CN and the CN sends packets with the same length and frequency to every MN.

The application sends 50 packets per second with a length of 8000 Bit. This is equivalent to a bandwidth of 400 KBit/s. Because of the IP header length added to each packet the required bandwidth grows to 425 KBit/s. One 10 MBit/s line of the network can hence support 23.5 MNs. The

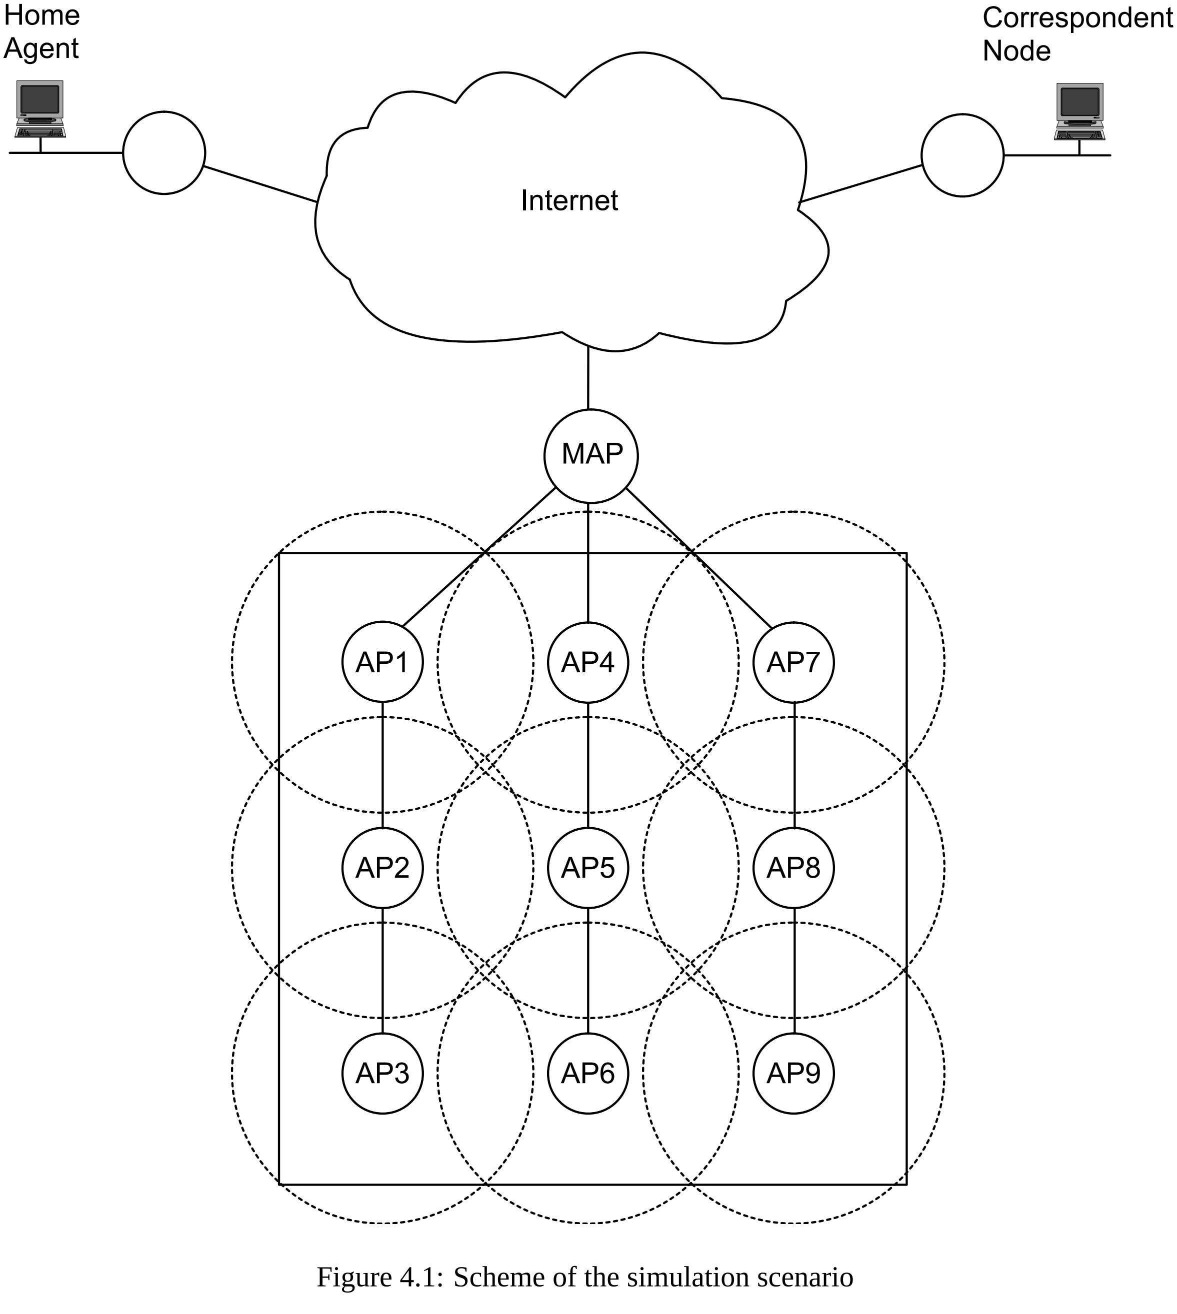

Figure 4.1: Scheme of the simulation scenario

whole access network can then support 70 MN if the MNs are optimally distributed. For this calculation the signaling load is not considered.

These 70 MNs are referred to as the “capacity” of the access network. The load is varied by changing the number of MNs.

In addition to the bandwith limitations imposed by the links between the routers, the routers themselves have a finite queuing capacity; priority queues are installed that can store 500.000 bits for each link. If the queue is full the router discards new packets; however, no signaling packets are lost within the routers in order to ease the implementation of QoS handling (no timer for the intermediate reservation is needed). Nevertheless, the signalling traffic’s bandwidth consumption is correctly accounted for.

4.1.4 Protocols

Each MN can independently choose a mobility protocol. There are three variants available: HMIPv6, the QoS-conditionalized BU approach and QoS-conditionalized BU approach using the release flag (R-Flag), allowing to release superfluous reservations along paths to which the MN has no link-layer connectivity any more.

Since the QoS-conditionalized BU approach is based on HMIPv6 extended by QoS support during handover, the MN uses this approach if it has a QoS-sensitive application. The traffic from this application is referred to QoS traffic. If the MN does not have a QoS application it uses the HMIPv6 approach. The network handles these packets as “best-effort” traffic (“be” traffic). Table 4.1.4 compiles some important parameters for the simulation setup.

Parameter	Value
Advertisement interval of the AP	100 ms
Lifetime of the QoS reservation	10 s
Speed of the MN	6.7 UL per s
Position update	every 600 ms
Packet length	8000 Bit
Inter packet time	20 ms
Radio range of the AP	100 UL
Queue length in the routers	500.000 Bit
Lifetime of the QoS reject	1 s
Dimension of the simulation field	300 x 300 UL

Table 4.1: Basic parameters for the simulation

4.2 Metrics

4.2.1 Handover Latency

A short handover latency is a necessary condition for QoS-sensitive applications because the service is usually interrupted during handover. For the simulation the handover latency is defined as the period from the loss of the IP “connection” to an access network to the reestablishing of the IP connection. The endpoint of this period is easy to measure: It is the time when a positive BA arrives. But the starting point is not measurable in the IP layer. Therefore, the simulation use the link-layer information about the interruption of the radio connection to the current AP.

The handover latency strongly depends on the movement detection used by the MN. In the simulation the MNs use the router advertisement sent from the AP to determine the presence of an AP. The APs periodically send router advertisements via the radio interface. If the MN can not receive three consecutive advertisements it assumes that the link-layer connection to this AP is lost. Thus, it needs between two inter-advertisement times and three inter-advertisement times only to detect the movement. In the simulation an inter-advertisement time of 100 ms is used. Therefore a MN needs from 200 ms up to 300 ms for the movement detection. Additional to this time comes the propagation delay and processing time, which is about 5 to 13 ms. So, if nothing disturbs the handover the latency should be in a short interval between 205 ms and 313 ms.

If the load of the network is high some handovers will take a longer time especially if the MN use the QoS-conditionalized BU approach because in some cases the BU will be rejected. Section 5 will provide results for this question.

4.2.2 Packet Success Rate

A packet success rate (PSR) is used as metric for the efficency of the network. It is defined as $PSR = \text{Received Packets in the CN}/\text{Sent Packets from the MN}$.

The packets can be lost during the handover or because of an overload situation in a router. The theoretical maximum of the PSR that a protocol can reach is not 100 % since packets are always lost during the handovers. The portion of packets lost during handover is determined by the number of handovers and the handover latency. The number of handover depends on the radio range, speed of the MN, and the used protocol. With the parameters from Table 4.1.4 a common value is three handovers per MN and minute. This causes, in combination with a mean handover latency of 250 ms in the best case, a maximal PSR of 98.75 %. If handovers are delayed, the PSR will suffer.

4.2.3 BU per successful Handover

If the MN uses the QoS-conditionalized BU approach it can take more than one BU for a successful handover: It send BUs as long as it does not receive a positive BA. If the network has to satisfy high load of QoS traffic the probability of a negative BU is also high.

Another possible metric would be to consider the signaling overhead both with respect to number of messages and processing delay in intermediate routers, with the latter option being the more interesting one. However, as the simulation environment described next did not easily allow to account for processing overhead, we did not pursue this question in detail; implicitly, the bandwidth consumption of signaling messages is of course taken into account.

4.3 Simulation environment

To model these scenarios and protocols, we implemented a simulation environment based on the OMNeT++ simulation tool and the associated IPv4 Suite for OMNeT++. The main extensions are the implementation of (Hierarchical) Mobile IPv6, where the relevant IPv6-specific characteristics are introduced by some simulation-specific packet extensions; on top of this, the QoS-conditionalized

Binding Update is implemented by adding resource constraints to an IP router and the capability to process the QoS options contained in BU messages. More details along with further references can be found in [8].

Chapter 5

Results

5.1 Handover Latency

Figure 5.1 shows a histogram of the handover latency for a MN that uses the QoS-conditionalized BU approach. Most handovers take place in a short interval between 200 ms and 300 ms (see Section 4.2). But a considerable number of handovers take much longer. This happens if the potential new route can not satisfy the QoS requirements of the MN. If a MN receives a negative BU+QoS it tries to use another AP if any is in radio range. Because of the movement detection the path via the old AP is always invalid. If no available AP can fit the MN's QoS needs the MN has to wait that a new AP becomes available or an AP gets free resources.

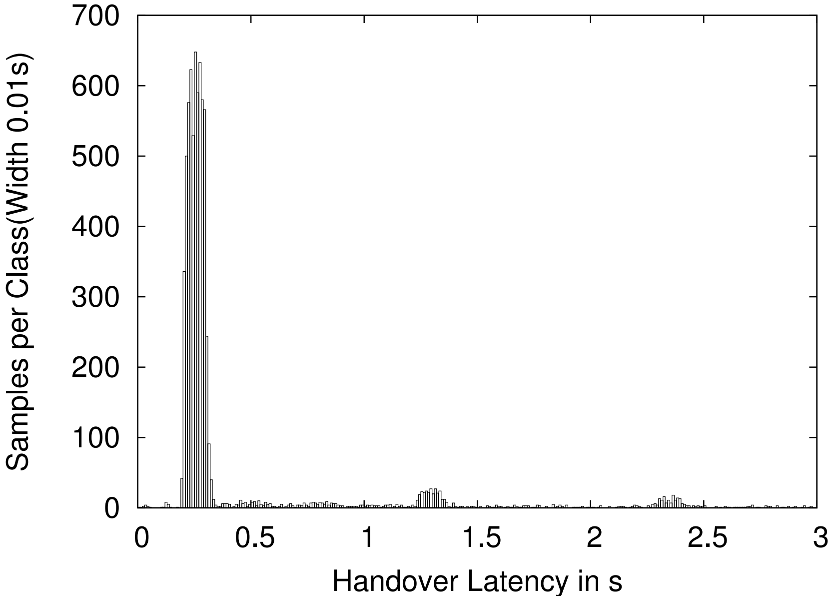

Figure 5.1: Histogram of the handover latency: QoS traffic; 90% load; 100 UL radio range; 50 s QoS-lifetime; update time 0.06s, 0.4 UL per step

These long handovers causes a statistically significant difference between the mean value for the "best-effort" traffic and the QoS-sensitive traffic: Table 5.1 shows means and confidence interval half length for a 95 % confidence level.

	BE traffic	QoS traffic
Handover latency	$0.2603\ s$	$0.82743\ s$
CI half length (95 %)	$\pm 0.0006562\ s$	$\pm 0.034656\ s$

Table 5.1: Handover latency with 95% load; Radio range 100 UL, Reservation lifetime 10 s

In the histogram of Figure 5.1, there are bursts at 1.25 s and 2.25 s. This depends on the "Lifetime of the QoS reject" which is a simulation parameter that determines how long a MN has to wait after receiving a negative BA+QoS before it can send a new BU+QoS to the same AP. In the simulation a MN had to wait 1 s .

5.2 Binding updates per handover

Even at high load, most handovers require only one BU for a successful handover but there are handovers where the MN has to send much more BUs in order to succeed; Figure 5.2 shows a histogram.

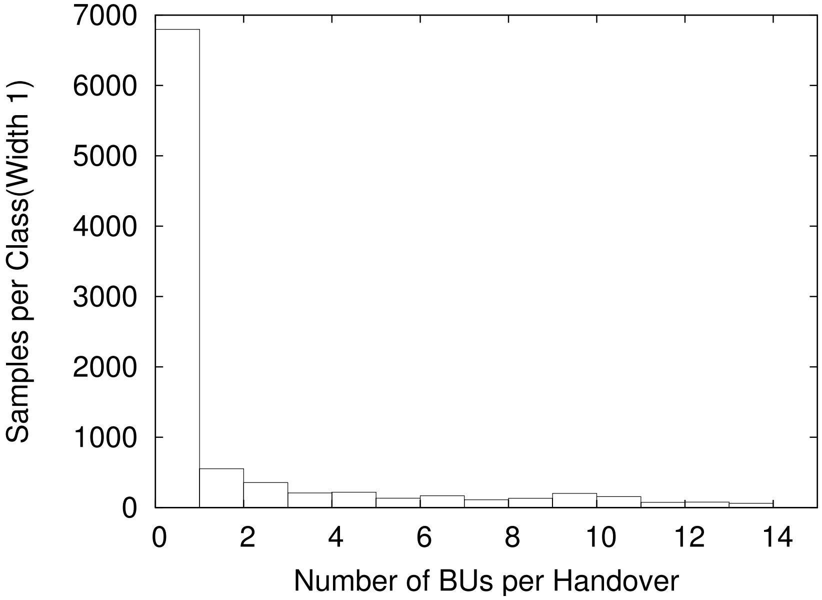

Figure 5.2: Number of BUs for a successful handover: QoS traffic; 95% load; 71 UL radio range; 50 s QoS-lifetime

5.3 Packet success rate

The following figures show the behavior of the access network if the MN uses different protocols. As metric for the efficiency, the PSR is used (see Section 4.2. The types of protocols are not mixed, but the figure shows the results of three independent experiments. For the simulation two parameters are varied: The radio range of the AP and the lifetime of QoS-reservation.

The left part of Figure 5.3 shows the characteristics of the protocols if the radio range is 71 UL and the QoS lifetime is 10 s. The radio range of 71 UL is the lowest value that allows the APs to cover the whole simulation area. In this situation the differences between the protocol alternatives are small. The BE traffic and the QoS traffic with the release flag (R-Flag) have nearly the same curve. The PSR for the QoS traffic without R-Flag is always below the other curves, especially if the offered load converges to the capacity of the access network. This happens because of the superfluous reservations that are kept up to 10 s in the system (see Section 3.3). No protocol reaches the theoretical optimum of 98.75 %. The maximum is reached with the HMIPv6 (92.8 %).

For the right half of Figure 5.3, the radio range was varied from 71 UL to 100 UL. This causes a significant higher PSR of maximal 97.8 % for both QoS-sensitive variants. With the larger radio range more APs become available to a MN, so that it can choose from more APs to find one that has enough resources. Thus, the probability of not being able to perform a handover it lower. As a result the MNs distribute their traffic better over the available APs.

The graph for the BE traffic shows the opposite behavior: The maximal PSR for large radio ranges is significantly lower (87.4 %) then with a radio range of 71 UL (92.8 %). This is due to an overload situation in AP 5 (see Fig. 4.1). In this scenario the probability for a MN to choose a particular AP is not equal for all APs because for some AP the radio coverage is not fully contained in the simulation area. So the border AP have to carry less load then the inner APs. Thus, using a non-QoS protocol it is important to determine the radio range according to the offered load. Using the QoS-conditionalized BU approach the MNs implicitly take care of a proper distribution of the offered QoS-sensitive load.

For the simulations that form the results shown in Figure 5.4, the lifetime for the QoS reservation is set to 50 s and the radio range is still at 100 UL. This modification is leading to significantly worse efficiency using the QoS-conditionalized BU approach without the R-Flag. The curves for the other protocols are equal to the curves from the right half of Figure 5.3. In the left figure the PSR for all protocols is equal at an offered load of 80 %. If the load is increased the curves for the BE traffic and the QoS traffic with R-Flag are nearly congruent but the PSR for the QoS traffic without release flag is a bit lower. If the load is less than 95 % the difference is lower than 6 percent points. If the load is increased further, the PSR of the QoS traffic without R-Flag collapses. At 100 % load the difference is 22 percent points. Because the superfluous reservation are kept valid for a long period, more and more MNs can not complete their handovers and therefor the PSR decreases.

In the right figure the same effect is identifiable. But here the curves for the QoS protocols are congruent if the load is lower than 95 %. This indicates that with the longer radio range the rejection of a BU+QoS is a very rare event at this load. But with higher load the protocol suffers from the unreleased reservations, too, and at a load of 100 % the difference is also 18 percent points.

5.4 The Impact of the Radio Range

As shown in the section above the radio range of the APs has an important impact on the efficiency of the QoS-conditionalized BU approach. Varying this range from 71 UL (smallest range for full coverage) to 110 UL, increases the PSR from 83.8 %. Larger radio range does not result in further improvements. This implies that, for the parameter combinations that we used, the load is nearly optimal distributed for a radio range of 110 UL. In Figure 5.5 the radio range is varied from 71 UL to

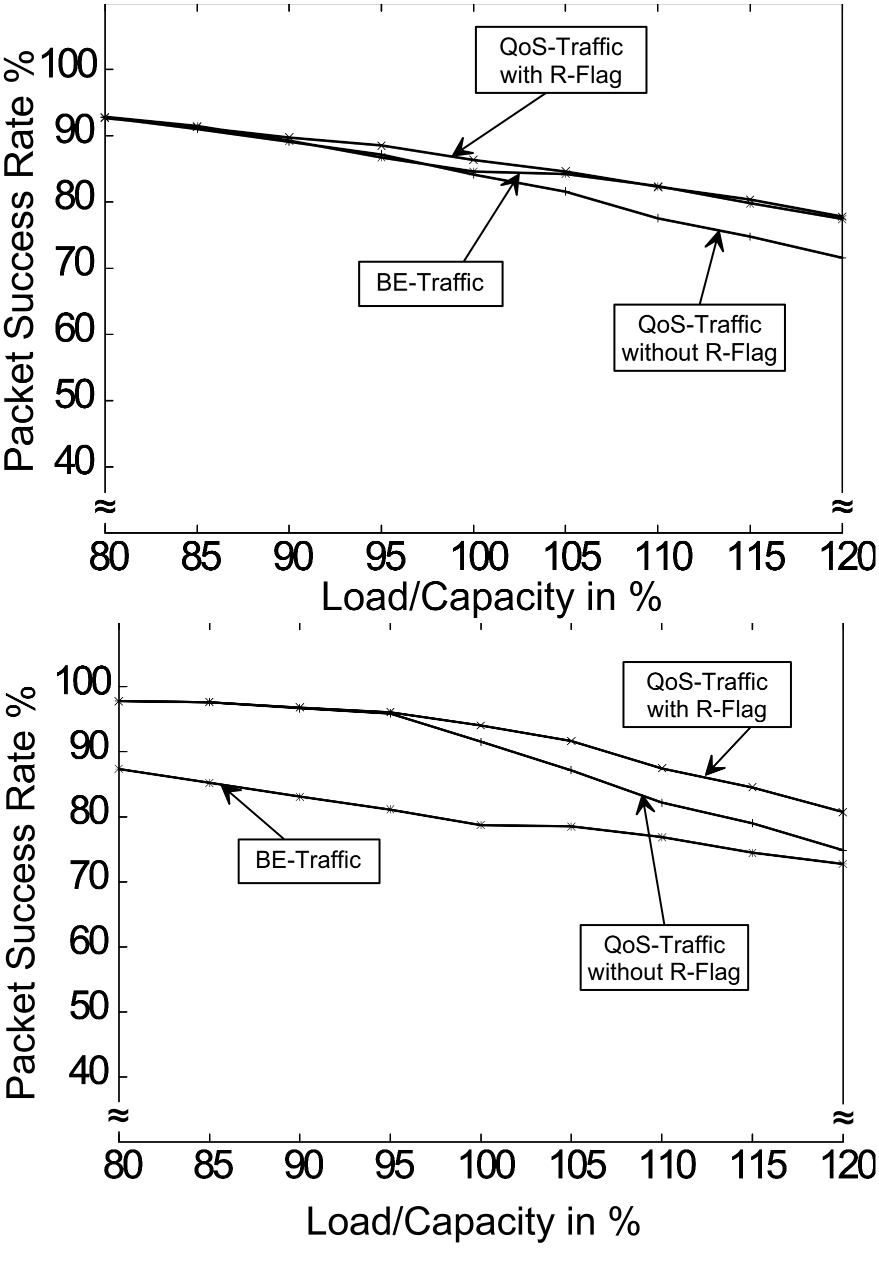

Figure 5.3: Variation of the offered load; 71 or 100 UL radio range; 10 s QoS-lifetime.

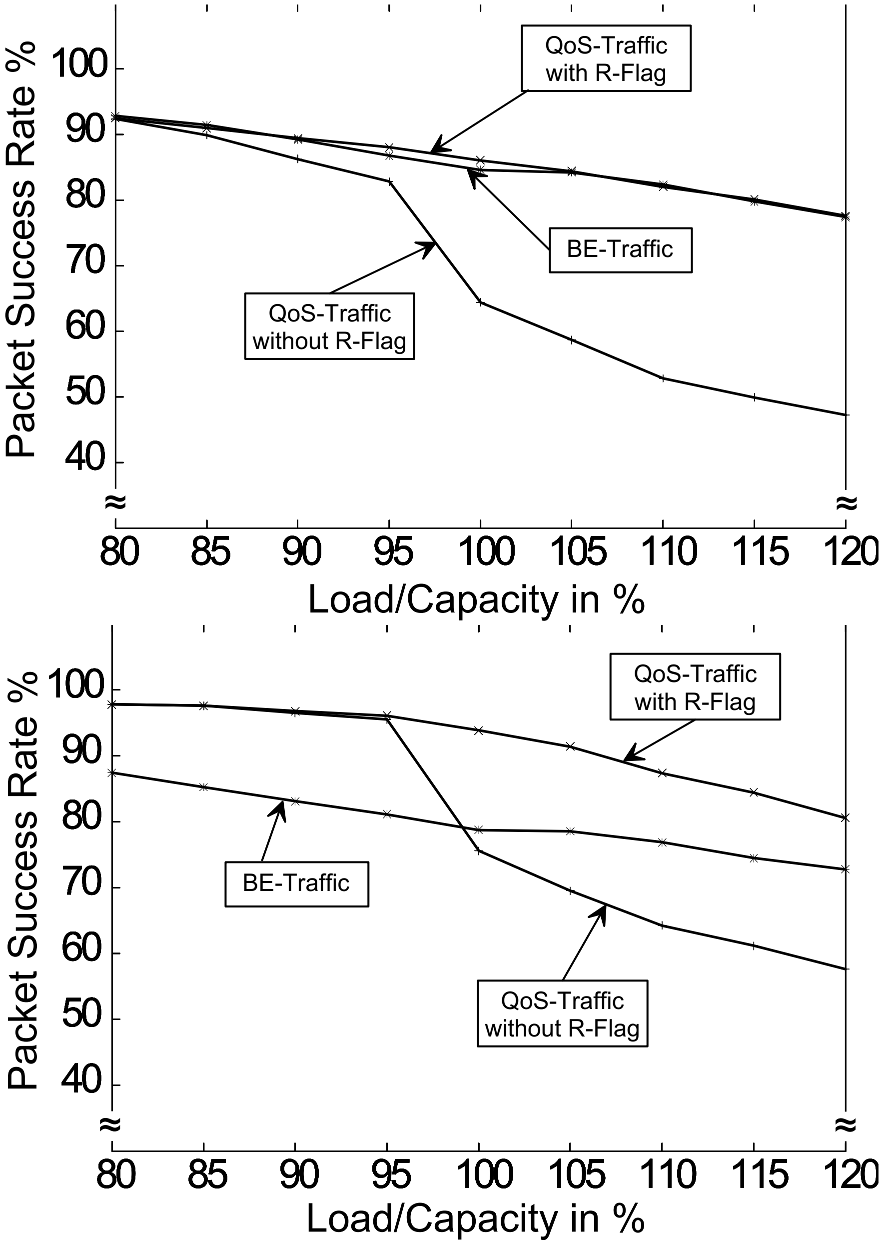

Figure 5.4: Variation of the offered load. 71 or 100 UL radio range; 50 s QoS-lifetime.

150 UL to evaluate this impact. With a radio range of 71 UL, the lowest range where the APs cover the whole area, the PSR is at 83.8 %. Increasing the radio range the PSR is also increased but not continuously. Above 110 UL there is only a very small increase. This implies that, for the parameter combinations that we used, the load is nearly optimal distributed for a radio range of 110 UL .

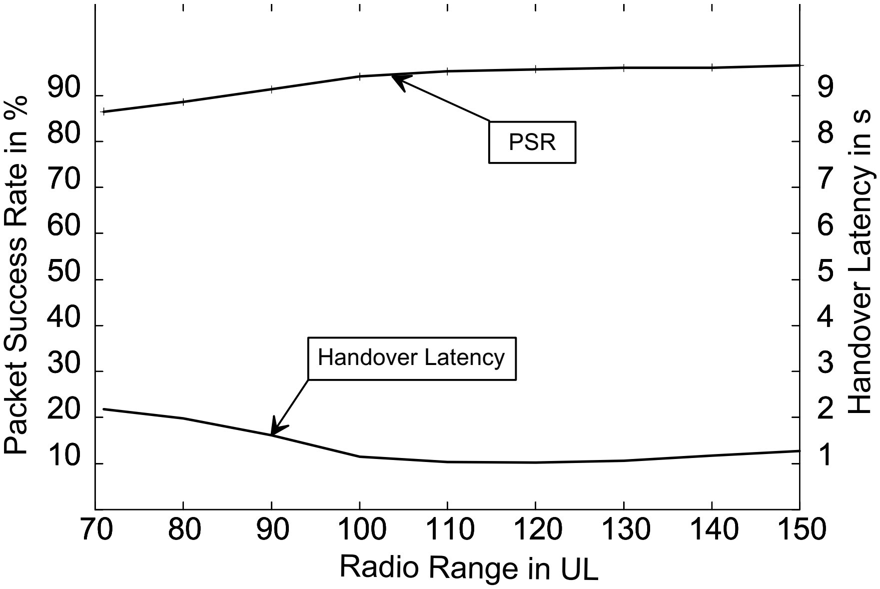

Figure 5.5: Impact of the radio range. load 100%; 10 s QoS-lifetime

Chapter 6

Conclusions and future work

This paper has presented a performance evaluation of the QoS-conditionalized Binding Update approach. We have shown that this approach is able to provide access control and to support handovers in an all-IP network with a proper notion of QoS. The overhead incurred in signalling traffic is acceptable; the total network capacity does not suffer. It is especially suited for a mix of pure best effort traffic and QoS traffic.

In fact, when adding explicit release of old reservations to the original QoS-conditionalized Binding Update proposal, it is competitive or even superior to the behavior of Hierarchical Mobile IPv6. In particular, HMIPv6 is susceptible to an improper choice of radio coverage area: With large cells, HMIPv6 has a tendency to overload some cells, resulting in an increase of packet losses. Our approach, on the other hand, rearranges traffic automatically, making use of unloaded base stations. Hence, our approach not only provides QoS support, it even makes the network more stable against parameter variations.

Further evaluations of this approach would be interesting for a wider range of parameters, e.g., network topologies. Also, a more precise model of processing times for signaling messages in routers would be desirable, preferably supported by measurements. The most interesting conceptual change, however, would be to include link-layer triggers into the mobility management. A link layer that is able to provide QoS information should represent an ideal fit with this approach and with the desired handover semantics. In the context of the SeQoMo project[1] where the QoS-conditionalized BU approach has been developed, also the efficient integration of authentication and authorization into QoS-aware handovers is under investigation.

---

[1] URL: http://www-tkn.ee.tu-berlin.de/research/SeQoMo/

Bibliography

[1] H. Chaskar and R. Koodli. A framework for qos support in mobile ipv6. Internet Draft draft-chaskar-mobileip-qos-01.txt (work in progress), March 2001.

[2] X. Fu, H. Karl, and C. Kappler. Qos-conditionalized handoff for mobile ipv6. In *Proc. of the Second IFIP-TC6 Networking Conf. - Networking2002*, Pisa, Italy, May 2002.

[3] D. Johnson and C. Perkins. Mobility support in ipv6. Internet Draft draft-ietf-mobileip-ipv6-14.txt (work in progress), July 2001.

[4] J. Kempf. Problem description: Reasons for performing context transfers between nodes in an ip access network. Internet draft, work in progress, May 2002.

[5] S. Paskalis, A. Kaloxylos, and E. Zervas. An efficient qos scheme for mobile hosts. In *Proc. of IEEE LCN'01*, Tampa, Florida, November 2001.

[6] C. Shen, W. Seah, A. Lo, H. Zheng, and M. Greis. Using RSVP in Mobile IP v6 networks. Internet Draft (work in progress), July 2001.

[7] H. Soliman, C. Castelluccia, K. El-Malki, and L. Bellier. Hierarchical mipv6 mobility management (hmipv6). Internet draft, work in progress, July 2001.

[8] S. Sroka and H. Karl. An OMNeT++ simulation model of hierarchical mobile ipv6 and qos extensions. `http://www-tkn.ee.tu-berlin.de/research/mobileip/index.html`, November 2002.

[9] L. Zhang, S. Berson, S. Herzog, and S. Jamin. Resource reservation protocol (rsvp) — version 1 functional specification. RFC 2205, September 1997.