Improving the performance of rate adaptation schemes in heterogeneous wireless networks

P. Stathopoulos[1], L. Sarakis[1,2], and N. Mitrou[1]

[1] School of Electrical and Computer Engineering,
National Technical University of Athens, Greece.
E-mails: {pstath,ls6}@telecom.ntua.gr, mitrou@softlab.ntua.gr
[2] Institute of Informatics and Telecommunications, National Center for Scientific
Research "Demokritos", Athens, Greece. E-mail: sarakis@iit.demokritos.gr

Abstract. The emerging heterogeneous networking environment will comprise diverse link layer technologies like 3G, WLANs and WPANs. In this context, multimedia applications should be able to cope with packet losses and delays due to congestion, wireless transmission errors and vertical handovers. This paper addresses the functional requirements from an advanced adaptation architecture, which incorporates enhanced end-to-end adaptation schemes that exploit information delivered by the next-generation terminal. The terminal's management entity should provide information at least for the local access network condition, e.g. estimates of the available bandwidth and wireless link losses, and handover notifications. The adaptation architecture does not count on the existence of QoS infrastructure; however, exploitation of such infrastructure is possible through the collaboration between the terminal and a network management entity. We have extended the LDA and TFRC rate adaptations algorithms to make use of the enhanced functionality. Through simulations we validate the robustness of these extended schemes against vertical handovers.

Key words: Rate adaptation, Vertical handover, LDA, TFRC

1 Introduction

In the next-generation networking environment the user will be always connected, selecting from a variety of available, existing and forthcoming network technologies, ranging from 2.5/3G cellular to *Wireless Local Area Networks* (WLANs), *Wireless Personal Area Networks* (WPANs), in home and ad-hoc networks. This environment will be characterized by the heterogeneity and the diversity of the participating networks in terms of performance, *Quality of Service* (QoS) capabilities, and control and management infrastructure. Furthermore, applications are expected to be pervasive; function gracefully and without interruptions, independently of the underlying network technology. In this context, maintaining a high end-to-end service quality, is crucial. The uncertainty about wide deployment of Internet end-to-end QoS schemes in the near future,

suggests that the applications should be able to adapt to changing network conditions.

Adaptation of streaming multimedia applications comes to play when the network state deteriorates, and it is usually triggered by packet losses. Such losses in homogeneous, single technology wireless environments are caused mainly by congestion and wireless transmission errors. In the next-generation environment, however, inter-system mobility and heterogeneity of the involved networks constitute extra sources of packet losses. The several macro and micro mobility protocols that have been proposed ([1], and references therein), result, in practice, to transient packet losses and increased delay, especially during vertical handovers [2]. Even if a seamless handover is offered by the mobility mechanism, severe congestion may arise in the new network, if this has significantly less resources than the old one (upward vertical handover). In such scenarios the appropriate action is to shield the adaptation scheme from transient mobility and resources-heterogeneity-related losses.

Among the proposed TCP-friendly rate adaptation schemes (see [3] for a survey), the *Loss Delay Adaptation* (LDA)[4] algorithm works in conjunction with the RTP/RTCP protocol stack and exploits the RTCP Receiver Reports, in order to transfer feedback information back to sender, while the *TCP-Friendly Rate Control* algorithm (TFRC)[5] takes into consideration feedback messages sent every *Round Trip Time* (RTT). These schemes require several RTTs and several seconds, respectively, to adjust their sending rate. In case of vertical handovers, this delay could cause buffers to temporarily saturate in the new network. To the best of our knowledge, the effects of vertical handovers on the performance of rate adaptation schemes have not been studied in sufficient depth.

The goal of this paper is twofold: first, to identify the functional entities that are expected to have a key role in an adaptation architecture targeting at next-generation heterogeneous networks, and second, to exploit this enhanced functionality towards making rate adaptation schemes robust against wireless corruption errors and vertical handovers. For evaluation proposes we have chosen to extend the LDA and the TFRC, mainly because they are representative examples of application and transport layer adaptation schemes, respectively.

The rest of paper is organized as follows. Section 2 briefly reviews existing work on rate adaptation. Section 3 identifies the functional components of the adaptation architecture. The extensions to the LDA and the TFRC algorithms are described in Section 4, and are validated against simulation in Section 5. Finally, Section 6 concludes the paper.

2   Related work

For applications running over wireless links the ability to discriminate between losses caused by congestion and losses caused by poor link quality is of great value. For losses caused by poor link quality the most appropriate action is to rely on error control (through link layer retransmissions and forward error control schemes) and adaptation mechanisms, like bandwidth compression using scalable

coding or adaptive packetization. Other approaches propose to use information from several layers to derive the fraction of lost packets in the wireless link. One such mechanism, ([6]), introduces a monitoring agent in the boundary of the wired/wireless segments in order to discriminate between packet drops due to congestion in the wired segment, and losses that are due to wireless transmission errors. Along the same lines, the authors in [7] use a TCP throughput formula to adjust the sending rate. [8] presents two TCP-friendly adaptive transport layer protocols: one for reliable data transmission and one for UDP-based multimedia transmission. Both are based on an "additive increase multiplicative decrease" algorithm with variable parameters. Each protocol represents a single unified approach that can be used equally well for heterogeneous networks with diverse characteristics in terms of delay and packet loss. The results comment on the mean throughput and the fairness achieved in the steady state and do not provide details on the performance of the schemes soon after the handover.

To the best of our knowledge, [9] is the only work that evaluates the performance of TFRC in the presence of vertical handovers (however, it elaborates only on the throughput achieved after the handover). This work, which has apparently been developed independently of ours (the motivation, including an initial architecture and some preliminary results, for our work lies in [10]) shows that the performance of TFRC can severely degrade during handovers and propose two mechanisms to overcome the problem (overbuffering and explicit handover notification).

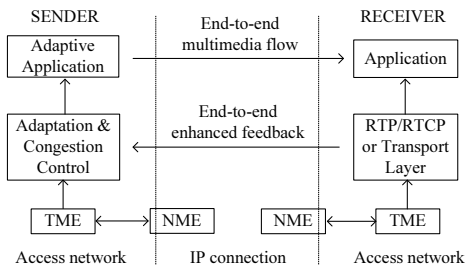## 3   Abstraction of the adaptation architecture

The majority of the adaptations schemes, briefly reviewed in the previous section, assume that certain entities will exist in the next-generation heterogeneous environment and reside both at the terminal and network sides. As far as the former is concerned, the requirements include the presence of a *Terminal Management Entity* (TME) and an appropriate middleware through which enhanced applications will make use of the functionality provided by the TME. The core functionality of TME consists in performing network availability detection, gathering link statistics and performing network condition estimation for vertical handover decision. As far as the network side is concerned, the presence of a *Network Management Entity* (NME) is highly encouraged. In this case, the collaboration between the NME and the terminal can lead to optimal handover decisions.

The proposed architecture combines end-to-end adaptation schemes with feedback messages that are enriched so as to include information provided by both the sender's and receiver's TMEs. The information generated by the TMEs is required to at least correspond to the status of the access networks or, complementary, to a more global network status, should the TMEs cooperate with corresponding NMEs. This enriched feedback information is furnished to the adaptation scheme which, in turn, calculates appropriate adaptation targets in terms of target rate and protection level against wireless errors. Based on these

parameters, the streaming application adapts appropriately by adjusting the encoding properties of the media stream.

At a glance, the sophisticated adaptation architecture incorporates: a) the employment of efficient end-to-end adaptation schemes that do not require any QoS scheme to exist, b) the exploitation of the TME access network monitoring information, both at the sender's and receiver's sides, c) the definition of enhanced feedback messages and d) the exploitation of established QoS mechanisms through the NME and TME interaction. The high level concept is to combine the strengths of the end-to-end schemes, namely, robustness, scalability, network transparency and simplicity, with the minimum set of functionalities that the next-generation networking elements will offer. QoS facilities should not be taken for granted for the entire path; this assures that the adaptation architecture is viable even when only best-effort service is provided.

The adaptation architecture is depicted in Fig. 1. The end-to-end feedback

Fig. 1. Abstraction of the adaptation architecture

messages alone can not provide discrimination between drops due to congestion and losses due to the quality of the wireless link. This desired functionality is provided by the TME. The information from the receiver's TME, combined with end-to-end information, like packet loss, RTT and delay variation of the complete path between the sender and the receiver, is sent back to the sender as an enhanced feedback message. This message together with the information provided by the local TME feed the congestion control and adaptation module. The TME can offer more than just local access network information when it cooperates with an NME.

In greater detail, the TME, by monitoring the present networks, can estimate, even in a coarse manner, the available resources. Furthermore, since the TME participates in network selection and handover decisions, it already has knowledge about imminent vertical handovers. Thus, the TME must calculate/retrieve and report to the adaptation module the following parameters:

• The R_{AVL} and, optionally, the R_{MAX}. R_{AVL} is an estimation of the currently available rate at the IP layer in the current network conditions (in case of handover the estimation can be proactive). R_{MAX} is the nominal value of the maximum rate at the IP layer the current network environment can support.

The granularity in estimating R_{AVL} and R_{MAX} can vary and depends on the capabilities of the TME.

- The l_W. This value denotes the fraction of packets lost due to wireless errors during the last reporting period.
- The boolean flag h. This indicates whether a vertical handover is underway, and should be reported as soon as a vertical handover is decided by the TME.

Both sender's and receiver's TMEs must provide the adaptation module with the information vector $TME_{out} = (R_{AVL}, R_{MAX}, l_w, h)$. The adaptation schemes employed should compensate for inconsistencies in the reported TME_{out} values. The following characteristics determine both the TME and the adaptation scheme requirements. In this respect:

- There is no requirement from the TME *regarding the precision* of the provided R_{AVL} value. This value must be used only as an indication, and the adaptation schemes must be able to recover gracefully from inaccurate estimations.
- The R_{AVL} and R_{MAX} values can correspond to end-to-end or local access network values.
- Updated R_{AVL} and R_{MAX} values are required to be reported in vertical handover events. These values are not required to be periodically reported, and thus reflect dynamic changes inside homogeneous networks, since the end-to-end adaptation schemes should be able to handle such changes.
- End-to-end QoS provisioning must not be required or taken for granted by the adaptation scheme.

Thus, the adaptation schemes used within the context of this architecture should be able to cooperate with TMEs that report TME_{out} information based on the aforementioned requirements.

4 The extended adaptation schemes

Two schemes are proposed in order to fit in the described architecture. The first is based on the LDA and the second on the TFRC (referred to as ext-LDA and ext-TFRC, respectively). These schemes are able to resolve losses due to wireless corruption errors and losses during vertical handovers, whether the latter are due to the mobility mechanism or the heterogeneity of the network's resources.

4.1 The extended LDA

The analysis follows that in [4]. (Due to lack of space we focus only on the differences between LDA and ext-LDA.) In LDA, the sender target rate is estimated by taking into account the fraction of lost packets, l, the round trip time, and the bottleneck bandwidth. The ext-LDA using local TME information (i.e. the vector TME_{out}) isolates losses due to wireless errors. If the fraction of lost packets at the sender's wireless network, l_W^{snd}, and at the receiver's access network, l_W^{rcv}, are known, the fraction of congestion losses, l_c, is given by the formula: $l_c = l - (l_W^{snd} + l_W^{rcv})$.

If $l_c > 0$, the sender reacts by decreasing the target rate, similar to the original LDA scheme (for simplicity we assume unicast transmissions). If r is the current rate, the new rate r' is given by $r' = r(1 - l_c \cdot R_f)$, where R_f is a reduction factor determining the degree of the sender's reaction to losses. During periods of no congestion or vertical handovers, ext-LDA increases the target rate in the same way that LDA does (i.e. by using the Additive Increase Rate (AIR) parameter). When a vertical handover occurs, the node performing this handover sends, without waiting for the reporting timer to expire, its TME vector to the adaptation module. (The sender uses the interface between the adaptation module and the TME, while the receiver uses feedback messages.) The adaptation module inspects the boolean flag h and, in case of a reported handover, calculates a new target rate r_{new}:

$$r_{new} = \min(R_{AVL}, r_{ext-LDA}),$$

where R_{AVL} is the estimation of the available bandwidth of the new network and $r_{ext-LDA}$ is the ext-LDA rate that would otherwise be calculated at that instance, if vertical handover did not occur. In this way, when the terminal moves to a network with fewer available resources ext-LDA reacts with an immediate rate change. If, on the other hand, the new network has more available bandwidth, the sending rate is increased smoothly according to the ext-LDA algorithm. In any case, after a handover the AIR parameter of the algorithm is set to its initial value.

In the ext-LDA algorithm the sender ignores the first regular RTCP feedback report sent by the receiver after the handover. This is because packet losses reported immediately after a handover cannot provide clear congestion indication. Using this heuristic mechanism, losses that are due to the mobility mechanism are isolated. The value l_W, carried in the information vectors TME_{out} can be used to feed a fidelity module that calculates the protection level against transmission errors.

4.2 The extended TFRC

The original TFRC algorithm ([5]) uses a TCP throughput equation to calculate the streaming rate. The loss event rate is calculated at the receiver and made known to the sender through feedback messages issued once every RTT. Using these feedback messages the sender estimates the value of the RTT and, implicitly, the value of retransmission timeout. The TFRC incorporates a sophisticated algorithm to map packet losses to loss events in order to obtain the loss event rate, p. If $p > 0$ the updated sending rate is the minimum between the rate that has been calculated by the TCP throughput equation and twice the receiver's rate. When there are no losses the sending rate is approximately doubled every RTT.

The ext-TFRC, like the ext-LDA, takes also into account the information vectors TME_{out} that are provided by the sender's and receiver's TMEs. The receiver's TME information is incorporated into augmented TFRC feedback messages. If we relax the condition that the TFRC algorithm reacts according to

the rate of loss events and we allow it to react according to the fraction of lost packets then, by using the l_W information in TME_{out}, it is possible to make it react only to congestion losses (in a fashion similar to ext-LDA).

When the receiver performs a vertical handover it issues an urgent TFRC feedback message notifying the sender about the parameters of the new network through the TME_{out}^{rcv} vector. If case of sender's handover it is assumed that its TME has already knowledge of the parameters of the new network. In any case the new sending rate is given by:

$$r_{new} = \min(R_{AVL}, r_{ext-TFRC}).$$

After the handover, the sender clears the history related to RTT and loss event rate estimates since this does not represent the current network conditions. We have chosen to let ext-TFRC keep the rate r_{new} until the first valid feedback message is received (we have heuristically chosen to discard the first two feedback messages after the handover; therefore, the first valid message is the third one). This is done in order to filter out excessive losses and large RTTs due to the mobility mechanisms.

5 Simulation results

In this section we evaluate the performance of the extended schemes against the original ones and TCP using the *ns-2* simulator. The ext-TFRC is also compared to a variation of TFRC that resets its congestion state (starting eventually from slow start) after vertical handovers. This enhancement has been proposed in [11] and will be referred to as ext-TFRC/SS.

5.1 Simulation topology and performance metrics

The topology that is used is depicted in Fig. 2. There are two access networks, namely Net-A and Net-B, each of which has four hosts ($H = F = 4$). Initially, a mobile host (MH) is served through Net-A. After a warm-up period, MH performs a handover from Net-A to Net-B (time instance t_1). Inter-system handover is supported by the Mobile IP v4 protocol ([12]). The common resources on both networks are shared using a generic round robin scheduler to achieve fairness in bandwidth sharing. The packet size is set to 1kbyte. All paths, except GW-HA and GW-FA have transfer delays equal to $1ms$. The links GW-HA and GW-FA represent the bottleneck links; all other links are over-provisioned. All queues, except HA-GW and FA-GW are $50Kbytes$ long. It is assumed that the MH is able to estimate, prior to handover, the bandwidth that it will get in the new network. This estimate is provided by the TME.

The precision with which the TME estimates the available bandwidth is denoted by the precision factor, f, as the ratio of the true to the estimated bandwidth that is available to MH (i.e. $f = 0.5$ denotes bandwidth overestimation by a factor of two). It is further assumed that the access networks are reliable links and thus no transmission errors occur.
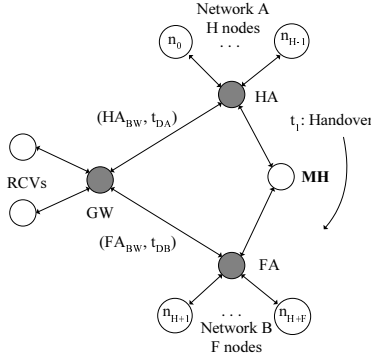
**Fig. 2.** The simulation topology. HA: Home agent, FA: Foreign agent

For every source, the actual sending rate, rs_i, the throughput, r_i, and the fraction of lost packets, d_i, are calculated. The first performance metric is the ratio of the source's sending rate to its fair share of the Net-B bandwidth, given by $\hat{rs}_i = rs_i/r_{fair}$, $r_{fair} \stackrel{def}{=} FA_{BW}/(F+1)$. The second metric is the total utilization $U(r,d)$ of a system with N sources and is given by $U(r,d) = (\sum_i^N U_i(r_i)g_i(1-d_i))/N$, where $U_i(r_i)$ is the utilization function of source i and $g_i(1-d_i)$ is a non linear factor that accounts for utilization reduction due to packet losses d_i ([13], [14]). A utilization function that has the desirable properties is given in [14] by $U_i(r) = log_2(r_i)$. Equivalently, in this paper we use the utilization function $U_i(r) = log_2(\hat{r}_i+1)$, which delivers a utilization equal to one when the distribution of the throughput is fair ($\hat{r}_i$ is the ratio of throughput to its fair share). Similarly to [13], the non linear function $g_i(z)$ is approximated by $g_i(z) = z^2$, $\forall i$.

The third metric is the fairness index of the bandwidth distribution in Net-B after the handover. This index is given by ([15]) $fi = (\sum_{i=1}^N x_i)^2/(N\sum_{i=1}^N x_i^2)$, $x_i = \hat{r}_i$. The last metric of interest is the Coefficient of Variation, CoV, of the sending rate of MH, rs_{MH}. If the average value of the sending rate calculated over a t_w time window after the handover is $\overline{rs}_{MH}^{t_w}$, and the standard deviation is $\sigma(rs_{MH}^{t_w})$, the CoV is given by $CoV(rs_{MH}^{t_w}) = \sigma(rs_{MH}^{t_w})/\overline{rs}_{MH}^{t_w}$.

5.2 Performance evaluation of ext-LDA

In the first set of simulations $FA_{BW} = 256kbps$, $t_{d_A} = t_{d_B} = 512ms$, $t_1 = 900s$, the queue length HA-GW corresponds to $512ms$ delay, and the queue length FA-GW is $16Kbytes$. Figure 3 shows that the LDA reacts slowly to the handover event and keeps its high sending rate for a considerable amount of time. The rate is then dropped abruptly as a response to reported packet drops. The result is that for an observation window of $60s$ after the handover the LDA makes little use of its available bandwidth. The ext-LDA with $f = 1$ accurately adapts to its fair share (as it should by assumption), with $f = 0.5$ it drops its sending rate at a
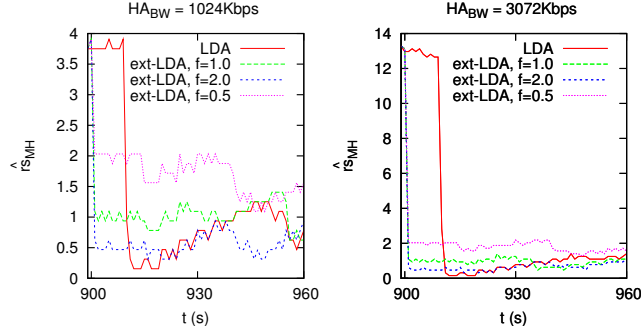
Fig. 3. Normalized sending rate $\hat{rs}_{MH}$ vs. time (LDA competing flows in Net-B)

fairly slow pace, achieving finally a fair share, and for $f = 2$ it delivers a smoother rate increase than LDA. Fig. 4 shows the CoV (covariance) of the sending rate
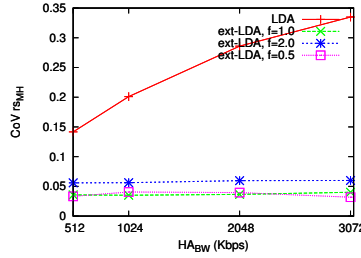

Fig. 4. Covariance of MH's sending rate vs. HA_{BW} (LDA competing flows in Net-B)

for $t_w = 60s$. The LDA has the larger CoV for any HA_{BW}. The mean values of fairness and utilization indices, calculated for $t_w = 30s$, are depicted in Fig. 5 (the figure includes the 95% confidence intervals for the extended schemes). The graph shows that the ext-LDA outperforms LDA, especially when the bandwidth heterogeneity between the two networks is large.

We have also compared the performance of ext-LDA to that of LDA in the presence of competing TCP flows in Net-B (the figures are not included due to lack of space). The results were similar and showed that the ext-LDA schemes preserve their small rate variation and achieve far better utilization and fairness. However, due to the relatively larger rate variation of the TCP flows, the sending rate of ext-LDA delivers slightly larger figures for the CoV.

5.3 Performance evaluation of ext-TFRC

In the second set of simulations $FA_{BW} = 512kbps$, $t_{d_B} = 250ms$, $t_1 = 200s$, the queue length HA-GW corresponds to $384ms$ delay and the queue length FA-GW
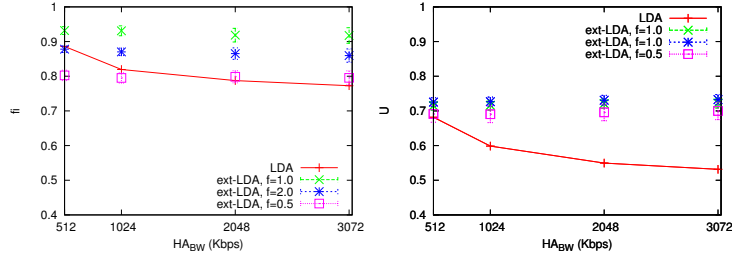
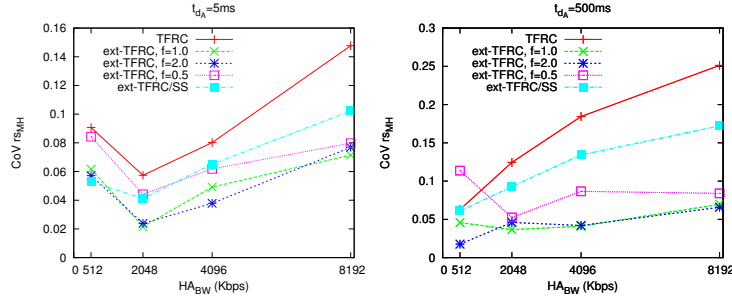**Fig. 5.** Mean value of fairness index (left) and utilization (right) vs. HA_{BW} (LDA competing flows in Net-B)

is $25Kbytes$. The reason for using a larger FA_{BW} is because TFRC makes better use of large available bandwidths. Variable t_{d_A} values are used to simulate networks with diverse bandwidth-delay products and provide insights into the impact the retransmission mechanisms have on the performance of TFRC. Figure 6
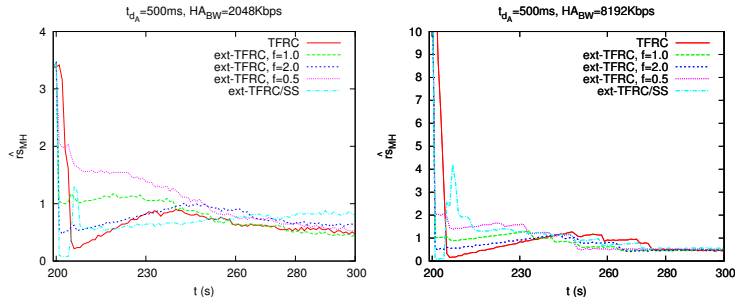

Fig. 6. Normalized sending rate $\hat{rs}_{MH}$ vs. time for $t_{d_A} = 500ms$ (TFRC competing flows in Net-B)

shows the evolution of the TFRC, ext-TFRC and ext-TFRC/SS sending rates after the MH's handover to Net-B, where four TFRC sessions are in progress. All schemes are prevented from converging to their fair share of the bandwidth at a satisfactorily fast pace. This is due to the large value of RTT at Net-A. Simulations with $t_{d_A} = 5, 250ms$ delivered a significantly smaller convergence period (small RTTs give rise to TFRC timers expiration and, consequently, to rate reduction, which is, in the case of upward handover, beneficial). Figure 6 illustrates that the TFRC reacts more slowly to network changes compared to ext-TFRC and ext-TFRC/SS. Shortly after the handover, the ext-TFRC/SS suffers from unstable behavior demonstrated by abrupt rate changes. The ext-TFRC with $f = 1$ and $f = 2$ seems to obtain the best performance demonstrated through small rate variations. The ext-TFRC with $f = 0.5$ cannot sustain a high sending rate but still avoids abrupt rate changes. Figure 7 shows the CoV of the sending

Fig. 7. Covariance of MH's sending rate vs. HA_{BW} for $t_{d_A} = 5,500ms$ (TFRC competing flows in Net-B)
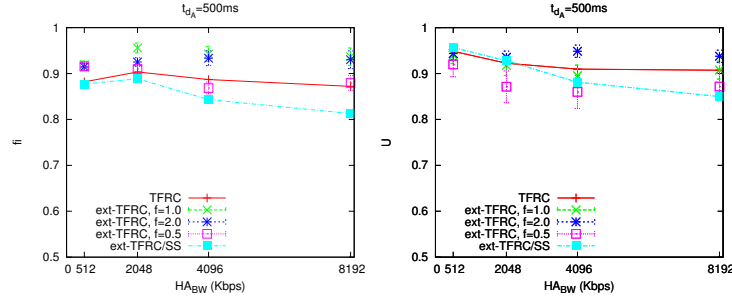
Fig. 8. Mean value of fairness index (left) and utilization (right) for $t_{d_A} = 500ms$ (TFRC competing flows in Net-B)

rate for $t_w = 60s$ and $t_{d_A} = 5,500ms$. In most of the settings, the ext-TFRC schemes outperform the original TFRC and ext-TFRC/SS. These observations still hold true for the mean value of the fairness and utilization indices, depicted in Fig. 8 ($t_w = 30s$). The performance of ext-TFRC has been also studied against TCP flows. The simulations, not depicted due to space limitations, illustrated that both the TFRC and the ext-TFRC were more aggressive and managed to obtain more than their fair share of bandwidth (for several seconds after the handover), especially for large HA_{BW} and small t_{d_A}. Nevertheless, the ext-TFRC schemes with $f = 1$ and $f = 2$ still performed better than the TFRC and the ext-TFRC/SS.

6 Conclusions

In this paper we presented an adaptation architecture targeting at next-generation heterogeneous networks and extended the LDA and TFRC rate adaptation schemes to fit in this architecture. The most essential feature of the architecture is the cooperation between a management entity, located at the terminal, and end-to-end adaptation schemes. This management entity provides enhanced

feedback to the extended adaptation schemes, including estimates of available bandwidth and losses due to poor link performance, and handover notifications. This approach targets at a low implementation overhead and high deployment scalability; most tasks are handled by end point nodes, and no special entities, like RTP monitoring agents, are required in the wired/wireless networks boundaries. The simulation results showed that the extended versions delivered better performance in terms of responsiveness, rate stability, utilization and fairness. The improved performance was maintained even when terminal management entities did not report perfectly accurate available-bandwidth figures, provided that these were not heavily overestimated. Future work will address a more detailed evaluation of the implementation overhead of the proposed schemes and study their performance in simulation scenarios featuring more complex topologies with a greater number of hosts.

References

1. Reinbold, P., Bonaventure, O.: A Comparison of IP Mobility Protocols, Tech. Rep. Infonet-TR-2001-07, University of Namur, Infonet Group (2001)
2. Stemm, M., Katz, R.H.: Vertical handoffs in wireless overlay networks. Mobile Networks and Applications **3** (1998) 335–350
3. Widmer, J., Denda, R., Mauve, M.: A survey on TCP-friendly congestion control. IEEE Network **15** (2001) 28–37
4. Sisalem, D., Schulzrinne, H.: The loss-delay based adjustment algorithm: A TCP-friendly adaptation scheme. In: Proc. of NOSSDAV, Cambridge, UK. (1998)
5. Floyd, S., Handley, M., Padhye, J., Widmer, J.: Equation-based congestion control for unicast applications. In: Proc. of SIGCOMM, Stockholm, Sweden (2000) 43–56
6. Yoshimura, T., Kawahara, T., Ohya, T., Etoh, M.: QoS Control Architecture with RTP Monitoring Agent for Mobile Multimedia Streaming. In: IPSJ Symposium on Multimedia, Distributed, Cooperative and Mobile Systems. (2001)
7. Yang, F., Zhang, Q., Zhu, W., Zhang, Y.Q.: End-to-end tcp-friendly streaming protocol and bit allocation for scalable video over wireless internet. IEEE JSAC **22** (2004) 777–790
8. Akan, O.B., Akyildiz, I.F.: ATL: An adaptive transport layer suite for next-generation wireless internet. IEEE JSAC **22** (2004) 802–816
9. Gurtov, A., Korhonen, J.: Effect of vertical handovers on performance of tcp-friendly rate control. ACM Mobile Computing and Communications Review **8** (2004) 73–87
10. Stathopoulos, P., Papageorgiou, P., Kouis, D., Zoi, S., Mitrou, N.: A rate adaptation scheme for media streaming over heterogeneous networks. International Journal of Wireless and Mobile Computing (2005) to appear in issue 5.
11. Kohler, E., Handley, M., Floyd, S.: Datagram Congestion Control Protocol (DCCP), Internet Draft, IETF (2004)
12. Perkins, C.: Mobility support for IPv4, RFC 3220, IETF (2002)
13. Bajaj, S., Breslau, L., Shenker, S.: Uniform versus Priority Dropping for Layered Video. In: Proc. of SIGCOMM. (1998) 131–143
14. Mo, J., Walrand, J.: Fair end-to-end window-based congestion control. IEEE/ACM Transactions on Networking **8** (2000) 556–567
15. Jain, R.: The Art of Computer Systems Performance Analysis. Wiley (1991)