

Safetynet version 2, a packet error recovery architecture for vertical handoffs

Henrik Petander¹ and Emmanuel Lochin^{2,3}

¹ NICTA, Sydney, Australia

² CNRS ; LAAS ; 7 avenue du colonel Roche, F-31077 Toulouse, France

³ Université de Toulouse ; UPS, INSA, INP, ISAE ; LAAS ; F-31077 Toulouse, France

Abstract. Mobile devices are connecting to the Internet through an increasingly heterogeneous network environment. This connectivity via multiple types of wireless networks allows the mobile devices to take advantage of the high speed and the low cost of wireless local area networks and the large coverage of wireless wide area networks. To maximize the benefits from these complementing characteristics, the mobile devices need to be able to switch seamlessly between the different network types. However, the switch between the technologies, also known as a vertical handoff, often results in significant packet loss and degradation of connectivity due to handoff delay and also increased packet loss rate on the border of the coverage area of the networks. In our previous work, we have proposed an inter technology mobility management architecture which addresses the packet losses using selective resending of packets lost during the handoff period. In this paper, we extend the architecture to address packet losses due to wireless errors more efficiently by taking advantage of erasure codes to form redundancy packets. We propose to send these redundancy packets over both links. We show that this proposal reduces both the chances of packet loss and the buffering requirements of the original SafetyNet scheme.

1 Introduction

With the proliferation of new wireless access network technologies, mobile users can now access the Internet using multiple types of access network technologies. This heterogeneous network environment provides access through a varying range of network technologies. The characteristics of these access networks vary greatly; Wireless Local Area Networks (WLANs) provide high speed access with a network latency of tens of milliseconds, often at the price of fixed Internet access but with a very limited coverage. Wireless Wide Area Networks (WWANs) on the other hand provide wide coverage but have a significantly lower data rate, higher latencies up to several hundreds of milliseconds and a cost which may several magnitudes larger than that of WLAN networks. Therefore, it is beneficial for a mobile user to be able to switch seamlessly between the different technologies.

Seamless switching between heterogeneous access networks requires carefully managed vertical (inter-technology) handoffs. Protocols, such as Mobile IP [4],

can be used for ensuring the handoff does not break the on-going connections of a mobile node and that the mobile node remains reachable in spite of the handoff. However, the vertical handoff performance of Mobile IP often leads to significant disruption of on-going traffic [1]. Earlier work in the field of mobility management has mostly focused on minimizing the impact of horizontal handoffs, i.e. moving between two networks of the same technology and does not provide optimal performance in vertical handoffs.

In our previous work, we proposed a localized mobility management protocol, *SafetyNet* [2] for minimizing the impact of vertical handoffs. The SafetyNet protocol utilises make-before-break handoffs, in which the mobile node breaks its connectivity with the previous access router (PAR) only after connecting to the new access router (NAR) which makes it possible to perform lossless handoffs. However, upward vertical handoffs, i.e. handoffs from WLAN to WWAN networks are typically performed only when the signal to noise ratio for the WLAN has degraded to nearly unusable. This poor signal strength of the WLAN may result in packet losses due to wireless errors or complete loss of connection with the PAR during the time it takes to prepare the WWAN interface and link layer connection to the NAR. To address this issue, SafetyNet combines the make-before-break handoffs with buffering at the NAR with selective delivery of packets from the buffer, so that any packets lost on the previous link (between the mobile node and PAR) are delivered from the buffer of the NAR at the new link. This allows for the recovery of the packets lost during the handoff period due to wireless errors or loss of connection with the previous access router. This mechanism is described in more detail in the next section.

The SafetyNet protocol showed a significant performance improvement over Fast handovers for Mobile IPv6 [5] in empirical measurements in favorable conditions for a vertical handoff scenario both in terms of TCP performance and in terms of over-the-air overhead. However, the worst case performance of SafetyNet is close to that of Fast Handovers for Mobile IPv6 protocol. In this paper, we target the average and worst case performance of the protocol by replacing resending of packets with use of packets based forward error coding (FEC) [6] both on the link of the previous router and on the link of the new router. We show that by using an adaptive coding scheme we can reduce the cost of recovering lost packets significantly while improving the SafetyNet architecture in terms of processing.

In the next section, we firstly present an overview of the SafetyNet protocol and then discuss its limitations in section 3. Our contribution is given in section 4 where we detail an improved solution and provide a preliminary performance evaluation. Finally we present relevant related work in section 5 and conclude this work in section 6.

2 Overview of the SafetyNet protocol

In the SafetyNet protocol [2], a Mobile Node (MN) moving from a link connected to the Previous Access Router (PAR) to a link connected to a New Access Router

(NAR), initiates the vertical handoff from PAR to NAR with PAR when it senses that it is about to lose connectivity with the PAR. The handoff process shown in Figure 1 is described briefly below.

- A mobile node which performs a handoff from an old network to a new network will first initiate the handoff with the access router of the previous network (PAR). In a vertical handoff, the mobile node would start activating the network card and preparing a link layer connection with the router of the new network link (NAR) at the same time.
- After a negotiation with NAR, PAR starts delivering copies of packets destined to the old location (care-of address) of the mobile node to the NAR by using tunneling. The PAR labels the packets using a sequence number, so that both the original and the tunneled copy of each packet have the same sequence number. The sequence number is incremented for each new packet. NAR stores these packets in a buffer.
- Mobile node observes the sequence numbers of packets it receives during the handoff at the old location on the link of the PAR. When the new link is ready, it signals the NAR to deliver all the packets from its buffer that the mobile node did not receive during the handoff at the link of the PAR. In Fig. 1, this would be packets starting from sequence number 2.

The details of the protocol are described in more detail in the original article [2].

3 Problem statement

The original SafetyNet architecture has four main issues. One of them deals with the rate at which the next AR (NAR in Fig. 1) buffer is filled. In the case where a mobile node moves from a WLAN to a WWAN, the rate at which packets arrive may often be higher than the rate at which they can be sent in the new network and as a result, the rate at which the buffer is emptied. Indeed, the new WWAN network may often combine a higher delay and a lower bandwidth than the old WLAN network, and this can lead to a significant loss of packets due to buffer overflow. The buffering has a second problem: The buffer space required is the product of the data rate for a user at the old link times the handoff duration. For vertical handoffs from WLAN to WWAN requiring activating a WWAN interface from a sleep state, this buffer space can amount to several megabits which may be a problem, if a large number of users are moving simultaneously. The third issue is the cost of recovering lost packets. The cost of retransmitting packets over a wide area network, such as a UMTS network, is typically significantly higher than the cost of (re)transmitting the same packets over WLAN networks. Additionally, the retransmission over the WWAN link increases the latency of recovering lost packets. Thus, it is beneficial both from a cost and performance point of view to retransmit less packets over the expensive WWAN networks. A possible solution to this is to enforce the reliability of the transmitted packets before the handoff and during the packet duplication process (See Fig. 1).

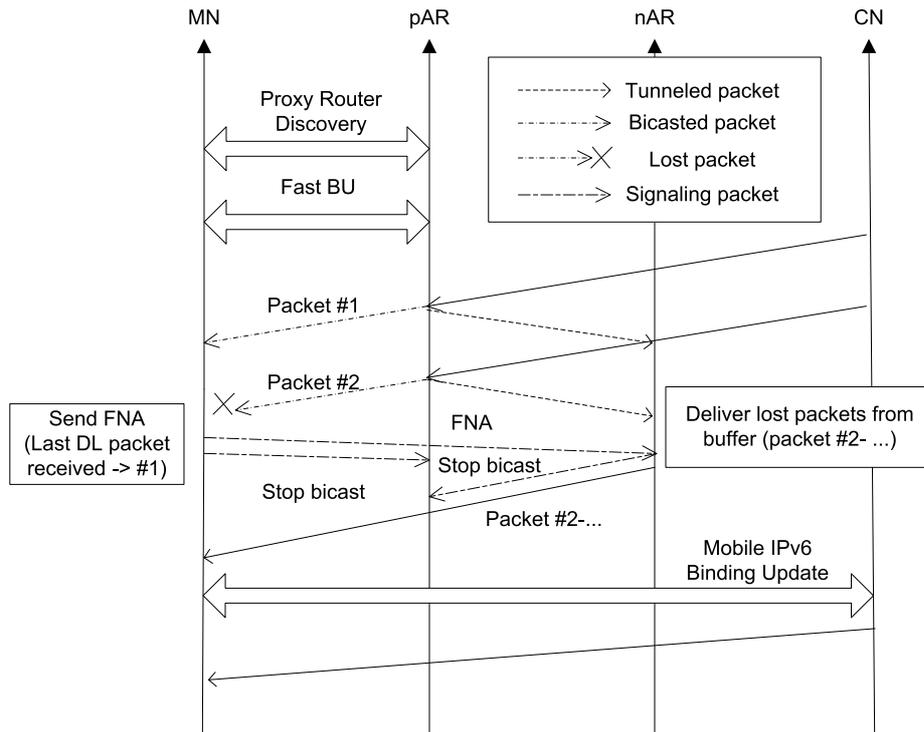


Fig. 1. The SafetyNet protocol operation

Looking at Fig. 1, a naive solution would be to retransmit packets during the bicast process from pAR to MN. Although this could be a solution over a WLAN network, it would be a problem over WWAN networks usually characterized by a high delay, low bandwidth and a higher cost of transmission due to the long delays of the retransmissions. A mobile node performing a vertical handoff from a WLAN would only do so when the connectivity of the WLAN became marginal, i.e. the bit error rate (BER) increased. Therefore, the mobile node would receive the packets with less resending via the nAR and a new link with a lower BER. This is illustrated in Fig. 2 which shows that even with a very high frame error rate (FER) of 25% on the new link, use of SafetyNet is less costly in terms of bandwidth when compared with resending over the IEEE 802.11 WLAN link. Even with a higher cost of communications over the WWAN link, it would still make sense in most cases to use SafetyNet to recover the packets. Therefore, we decided to not tackle this possibility.

To briefly summarize, the main points we would like to improve are:

- to minimize the buffer occupancy at the nAR side in order to avoid dropped and delayed packets;

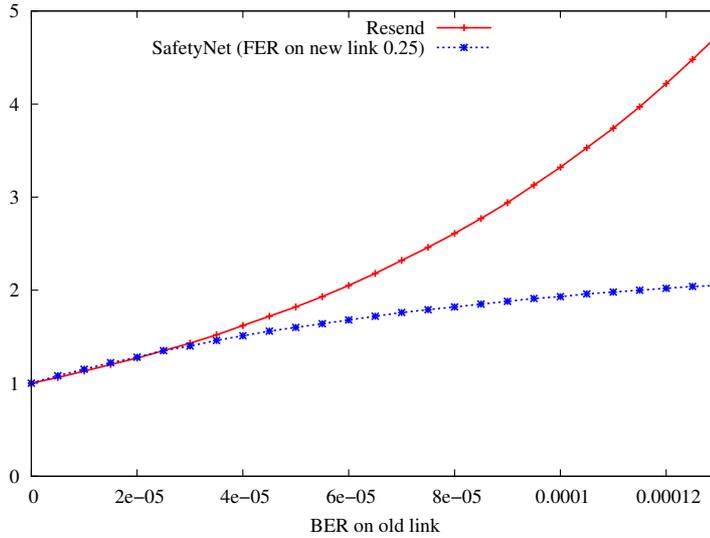


Fig. 2. Comparison of bandwidth cost of resending lost packets over the WLAN link with the use of SafetyNet

- when possible, to enforce the reliability of the transmitted packets during the handoff process in order to decrease the number of the retransmitted packets by the NAR;
- and to achieve this, without introducing overhead in terms of number of packets sent over both networks.

In this paper, we propose to enhance the existing SafetyNet protocol by using an erasure code scheme. The use of such a mechanism leads to potential performance improvements both in terms of reliability during the handoff process and minimization of buffer occupancy. In this context, we propose to use such code both by the PAR and NAR. Indeed, we propose to introduce some redundancy packets on the PAR side during the handoff and the bicast process to decrease the number of lost packets transmitted before handoff and to send only redundancy packets to the NAR. Obviously, by sending only redundancy packets to the NAR, the buffer occupancy will decrease markedly.

4 Redundancy packet bicast scheme

Most of forwarding error codes used over packet erasure channels are block codes [3]. This means that at the encoder side, a set of repair packets (R) is built from a given set of source data (SD) packets and at the decoder side, these repair packets can only be used to recover SD packets from their corresponding set. We denote k the number of SD and $n - k$ the number of R . If too many packets

(among the SD and repair packets) are lost during the transmission, the recovery of the missing SD packets is then not possible (i.e $n - k$ redundancy packets allow to recover k packets among n where the k recovered packets correspond to the original SD packets). On the opposite, if only few packets are lost, some of the repair packets become useless. A solution to this problem, known as Hybrid FEC-ARQ (or H-ARQ) mechanism [3], is to use receiver's feedback to send additional repair packets or to adjust the redundancy level of the FEC to the observed packet loss rate. In our context, retransmissions of redundancy packets are done, if necessary, by the NAR after receiving the FNA message. When the NAR receives the FNA message, containing sequence numbers of any lost packets, it can determine which redundancy packets need to be sent to the mobile node. Before the handover, the mobile node receives the traffic from PAR and moves towards the second network. During the handover, the bicasting procedure duplicates different linear combination of redundancy packets towards both networks. Finally, when the MN arrives in the new network and after sending the FNA feedback packet, the bicasting process ends and the communication is re-routed through the NAR.

To assess the right configuration of the FEC block code, an estimation of the PER is needed. However, including a mechanism able to estimate the exact PER during a handover procedure is not realistic and prevent from any real deployment. Furthermore, estimating a PER in WLANs is hard due to the unpredictability especially in indoor environments. Following our experiments [2] and experience, a PER above 25% makes it impossible to maintain TCP connectivity whatever the protection method used as the large amount of lost packets to retransmit or rebuild result inevitably in TCP timeouts. However, the use of a fixed PER allows getting around this problem and dealing with a wide range of PER with little additional overhead. Thus, we propose to set a default peak PER value on the AR device that can be tuned by the wireless network administrator and propose as a default value 20%. For instance, we use a $(k, n) = (4, 5)$ FEC code, which means that one R packet is built from a linear combination of four SD packets, these five packets are sent through the previous link and the bicasting procedure builds and sends another redundancy R' , which is a second linear combination of the fourth SD , to the NAR. Basically, this means that we bicast supplementary linear combinations of redundancy packets to the NAR. As a result, we double the number of redundancy packets and half of them are used to correct losses on the wireless link (the estimated 20%) while the other part is simply stored on the NAR for retransmission purpose.

We previously said that a mobile node can move from either a slower to a faster network or the reverse. Usually, faster networks are WLAN networks which have very low delay compared to WWAN. In the Safetynet/FEC scheme, the rationale to bicast redundancy packets is: 1) to enforce the reliability of the flow on the previous link during the handover in order to reduce the amount of retransmission by the NAR; 2) to ensure all lost packets during the handoff will be rebuilt thanks to supplementary redundancy packets as soon as MN

advertises to the NAR the sequence numbers of any missing packets (through the FNA feedback).

Without bicasting, the probability P to rebuild all packets is given by:

$$P = \sum_{i=k}^n \binom{n}{i} (1-p)^i p^{n-i} \quad (1)$$

This protects the flow to a PLR p when $p < \frac{(n-k)}{n}$. For a given FEC code (k, n) , the bicasting process actually offers (k, n') protection where $n' = 2*n - k$. With the bicasting process, this probability, denoted P_{FEC} becomes:

$$P_{FEC} = \sum_{i=2(n-k)}^n \binom{n}{i} (1-p)^i p^{n-i} \quad (2)$$

and protects $p < \frac{2(n-k)}{n}$ with potential redundancy packet retransmissions. The theoretical results are given figure 3. In this figure, we draw the probability to rebuild a packet as a function of the PLR for a FEC code able to theoretically correct a PLR of 20% with and without bicasting following (1) and (2). The difference between these two curves allows us to assess the probability of redundancy packet retransmissions which is given by the third curve.

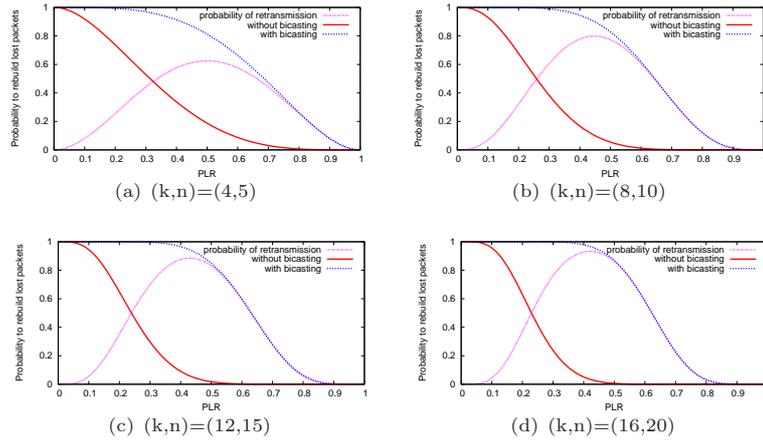


Fig. 3. Probability to rebuild a block as a function of the PLR for various FEC codes

As a matter of fact, the possibility to retransmit supplementary linear combination of packets allows to better protect the flow during the handover when the signal becomes poor while decreasing the buffer size of the NAR. Compared to the previous Safetynet proposal, this FEC scheme decreases the number of possible retransmissions as some packets are rebuilt thanks to the FEC redundancy packets. This combination allows to ensure that, even if FEC code is not

able to correct all losses inside a block, a retransmission of redundancy packets via the NAR will allow to correct up to the double the PLR (i.e. up to $\frac{2(n-k)}{n}$) with the price of redundancy packets retransmission delay. The success of this mechanism is obviously linked to the configuration of the FEC code. Indeed, if a full block of packets is lost on the PAR link, the retransmission of supplementary redundancy packets from the NAR will not allow to rebuild the missing block (i.e. if $p > \frac{2(n-k)}{n}$). In this case, a retransmission from the source will be inevitable and would correspond to a wrong setting of the FEC parameters by the administrator. However, Figure 3 shows that in our example, the probability to not recover more than k packets among n with a retransmission (which corresponds to $1 - P_{FEC}$) is nil when the PLR is below 20%, meaning that a rough configuration of the PLR, our scheme should cover this case.

5 Related Work

Forward error correction has been used earlier in soft handoffs in CDMA systems [7]. A mobile node in a soft handoff sends and receives two bit streams via two different base stations which are combined into a single stream at the receiver. The two streams are synchronized, and contain redundant information for handling transmission errors. The use of FEC in IP based communications handoffs has also been proposed earlier for seamless horizontal handoffs for multicast video traffic [8]. A more general solution for horizontal handoffs was proposed by Matsuoka et al. In their proposal which resembles the CDMA soft handoffs, a mobile node receives multiple IP packet streams encoded with a Reed Solomon code via different WLAN access points and combines them into a single stream with redundancy [9]. This work, which is so far the closest to our contribution, proposes to spread the redundancy over two different links. The main difference with our scheme is that the authors split the encoded streams while we propose to duplicate the FEC encoded stream with other linear packet combinations. This allows to increase the error capability correction compare to their scheme.

Finally, these approaches can not be directly applied to vertical handoff scenarios due to the different delay and bandwidth characteristics of the networks in a handoff. Further, they only address the time after the connectivity on the new network has been established. Our work focuses on vertical handoffs and uses a separate stream of selectively delivered redundancy packets to deal with the asymmetric delays and to address the potentially long handoff process itself.

6 Conclusion

In this paper, we presented an improved architecture for localized mobility management providing a mechanism for recovering packets lost during the handoff. The architecture applies forward error correction on packet level to reduce the buffer occupancy of the SafetyNet architecture while providing error recovery.

We provided an initial evaluation on the efficiency of the mechanism which indicates that the architecture may provide recovery of lost packets at a smaller cost than the original SafetyNet proposal which already reduces the resending of packets significantly.

The analytical evaluation of the algorithm suggests that the use of forward error correction may be an interesting strategy for vertical handoffs. As a next step, we are planning to implement the algorithm as a part of the SafetyNet architecture and evaluate the performance empirically using our SafetyNet implementation.

References

1. P. Chakravorty, K. Vidales, I. Subramanian, J. Pratt, J. Crowcroft, *Performance Issues with vertical handovers-experiences from GPRS cellular and WLAN hot-spots integration*, in IEEE PerCom, 2004
2. H. Petander, E. Perera, A. Seneviratne, *Multicasting with selective delivery: A SafetyNet for vertical handoffs*, in Springer Journal on Personal Wireless Communication, vol. 43, issue 3
3. S. Lin, D. Costello, *Error Control Coding: Fundamentals and Applications*, Prentice-Hall, Englewood Cliffs, NJ, 1983
4. C. Perkins, *IP Mobility Support for IPv4*, Request For Comments 3344, 2002
5. D. Johnson, C. Perkins, J. Arkko *Mobility Support in IPv6*, Request For Comments 3375, 2004
6. C. Huitema, *Case for Packet level FEC*, in proceedings of the TC6 WG6.1/6.4 Fifth International Workshop on Protocols for High Speed Networks, 1996
7. D. Bevan et al, *Improved Soft Handoff Method For Uplink Wireless Communications*, International patent no. WO/2004/004155, 2002
8. H. Liu et al., *A Staggered FEC System for Seamless Handoff in Wireless LANs: Implementation Experience and Experimental Study* in IEEE ISM, 2007
9. H. Matsuoka, T. Yoshimura, T. Ohya, *End-to-end robust IP soft handover*, in IEEE ICC, 2003