



**HAL**  
open science

## A packet error recovery scheme for vertical handovers mobility management protocols

Pierre Ugo Tournoux, Emmanuel Lochin, Henrik Petander, Jérôme Lacan

► **To cite this version:**

Pierre Ugo Tournoux, Emmanuel Lochin, Henrik Petander, Jérôme Lacan. A packet error recovery scheme for vertical handovers mobility management protocols. ICST MobiQuitous - 7th International ICST Conference on Mobile and Ubiquitous Systems, Dec 2010, Sydney, Australia. pp.NC. hal-00582606

**HAL Id: hal-00582606**

**<https://hal.science/hal-00582606v1>**

Submitted on 2 Apr 2011

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

# A packet error recovery scheme for vertical handovers mobility management protocols

Pierre-Ugo Tournoux<sup>1,2</sup>, Emmanuel Lochin<sup>1,2</sup>, Henrik Petander<sup>3</sup>, Jérôme Lacan<sup>4</sup>

<sup>1</sup> CNRS ; LAAS ; 7 avenue du colonel Roche, F-31077 Toulouse, France

<sup>2</sup> Université de Toulouse ; UPS, INSA, INP, ISAE ; LAAS ; F-31077 Toulouse, France

<sup>3</sup> NICTA, Australian Technology Park, Eveleigh, NSW, Australia

<sup>4</sup> ISAE ; 10 av. Edouard Belin - BP 54032 - 31055 Toulouse Cedex 4, 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. In this context, we propose a new handoff framework for switching seamlessly between the different network technologies by taking advantage of the temporary availability of both the old and the new network technology through the use of an “on the fly” erasure coding method. The goal is to demonstrate that our framework, based on a real implementation of such coding scheme, 1) allows the application to achieve higher goodput rate compared to existing multicasting proposals and other erasure coding schemes; 2) is easy to configure and as a result 3) is a perfect candidate to ensure the reliability of vertical handovers mobility management protocols. In this paper, we present the implementation of such framework and show that our proposal allows to maintain the TCP goodput (with a negligible transmission overhead) while providing in a timely manner a full reliability in challenged conditions.

## 1 Introduction

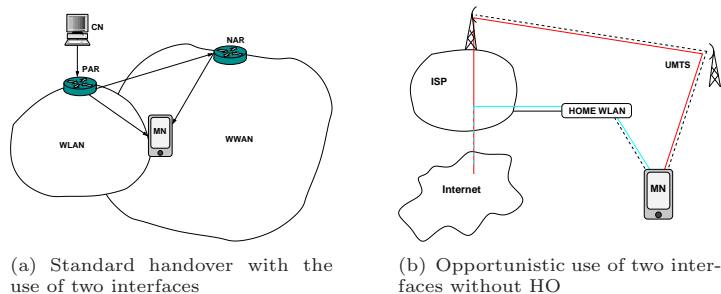
With the proliferation of new wireless access network technologies, mobile users can now access the Internet using multiple types of access network technologies. This heterogeneous network environment provides access through a varying range of network technologies. The characteristics of these access 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 be several magnitudes larger than that of WLAN networks. For obvious cost and performance reasons, smartphone users frequently switch to WLAN when a hotspot is available although the cellular connection is almost always enabled. Therefore the ability to switch seamlessly between these different technologies

allows a user to maximize his data rates and an operator to free resources in more expensive WWAN networks by maximizing the utilization of lower cost WLAN networks.

Seamless switching between heterogeneous access networks requires carefully managed vertical (inter-technology) handovers. Protocols, such as Mobile IP (see RFC 3344), can be used to ensure the handover does not break the on-going connections of a mobile node and that the mobile node remains reachable in spite of the handover. In a Mobile IP vertical handoff, on-going traffic is often disrupted due to protocol deficiencies [2]. Although, more advanced handoff protocols such the Safetynet architecture [13] (and Safetynet v.2 [12], where the use of FEC codes are suggested to mitigate the number of lost packets during the bicasting) can be used to reduce these packet losses, the challenging wireless link conditions triggering the handoff may cause unavoidable packet losses. This is especially the case for upward vertical handovers (i.e. handovers from WLAN to WWAN networks) which are typically performed only when the signal to noise ratio of the WLAN becomes too weak to offer a correct connectivity. This decrease of the signal strength may result in packet losses due to wireless errors or even a complete loss of connectivity with the Previous Access Router (PAR) during the time it takes to prepare the WWAN interface and link layer connection to the Next Access Router (NAR).

There are currently two main class of solutions that address the problem from a transport layer point of view. The first one aims to improve TCP tolerance to handover [7, 3] while the second one uses multipath SCTP [14] version to benefit from this multiple connectivity capability [5, 16]. In the present contribution, we show that the proposed coding scheme shares both advantages and allows the use of any kind of transport protocol without modification. In other words, this proposal is completely independent of the end-to-end transport solution deployed.

In this paper, we present the potential use of this “on the fly” coding scheme through an implementation called Tetrys [15], which can be applied to transport and layer-3 mobility management protocols to achieve a so called “soft handover” and significantly reduce the impact of the handover on the application traffic. A soft handover allows a mobile device to connect to multiple networks at the same time and receive coded streams of traffic from multiple routers or base stations at the same time and to combine those coded, partially redundant streams to a single, complete data stream. This allows the handover process to be very smooth since the ratio of data and the level of coding of the different streams can be dynamically adjusted to handle changing packet loss rates. So far soft handovers have been successfully used only in tightly controlled horizontal handovers in CDMA networks in which the traffic is synchronized between the different base stations and the mobile device. This paper explores whether a similar soft handover can be achieved in handovers between IP based WLAN and WWAN networks with the more challenging asymmetric and non-synchronized network conditions. The purpose of this study is not to propose yet another exhaustive mobility management architecture as this generic coding scheme could



**Fig. 1.** Two illustrations of multipath.

be used inside any mobility management protocol (such as Fast Handovers Mobile IPv6 for instance). We rather seek to demonstrate that our adaptive coding scheme can significantly reduce the impact of the challenged network conditions in a vertical handover by using a soft handover like approach. Thus, we evaluate our proposal for vertical soft handovers and the results obtained show that TCP remains close to its pre-handoff bitrate.

## 2 Background and related work

A soft vertical handover differs from a hard vertical handover as no disconnection occurs during the soft handover process. Although both are challenging, they impact the transport and application layers differently and also the methods used to minimize the handover effect at both layers are different.

The main problem in hard handovers, and also failed soft handovers, is the handover delay during which packets are lost. These lost packets cause a number of issues, discussed below. Solutions for reducing this delay, i.e. optimizations for hard handovers, consist of reducing the network detection period, the address configuration interval and the network registration time which reduce (see [2]). The solutions are also valid for reducing the chances of a soft handover turning into a hard handover.

The first and the most pressing issue that a transport protocol, such as TCP, must handle is the loss of its in-flight packets. This handover delay might trigger RTO in TCP and the resulting backoff procedure could lead to a connection stall. This problem can be mitigated by freezing the sending of TCP packets during the handover process [7]. Another challenge in hard and soft handovers is the possible differences in network characteristics between the new and the previous network in terms of propagation delay, bandwidth and packet loss rate (PLR). A plethora of work have been devoted to this problem and have mainly proposed solutions to adapt the TCP congestion windows length or to quickly update the TCP RTO value on the new link (involving active probing on the new link) in order to prevent RTO buffer overflow and/or RTO expiration [7, 3].

As the multiple interfaces present on a mobile host can be enabled at the same time, many proposals provide soft handovers between heterogeneous networks. This can be achieved with IP-level mobility solutions such as MIH, Multihoming MIP (see RFC 4908 [10]) at the cost of a devoted network infrastructure. This provides (see RFC 4980 [11]) improved reliability, load sharing between the different links and bandwidth aggregation.

Several propositions use multihoming to improve the quality of the communication by multicasting the flows via multiple available interfaces. When a handover becomes highly probable, packets are sent both from the PAR and the NAR. This allows a more robust service as it is known that the link signal can significantly decrease during the HO leading to high packet loss rate (PLR) as shown in Fig. 1(a). Instead of copying the same packets on both paths, the authors of [9] choose to send data packets on one path and redundant FEC packets on the other, thus reducing drastically the impact of losses on a video stream compared to the standard multicasting procedure. In [4], the authors obtain a similar result using staggered FEC.

Several modifications to the Stream Control Transmission Protocol (SCTP) have been proposed, enabling support for both real-time and non real-time traffic [5, 16] taking benefit of the multihoming capability without the cost of a large network architecture deployment. For instance, the authors in [16] demonstrate that the use of SCTP allowed them to aggregate the bandwidth of different networks, thus providing a video with a better quality and a more robust service even in high mobility scenario.

To illustrate the generic character of our solution, another application where our proposal is of interest is illustrated in Fig. 1(b). In this figure, we represent a mobile node which has subscribed to both ADSL and 3G offer within the same operator and opportunistically uses its home wireless access (or any accessible local wireless spot) conjointly with a 3G access. The benefit of our proposal in this context will be highlighted in Section 4. Our scheme can be deployed both in an end-to-end fashion or from the edge router of the ISP to the terminal host while remaining transparent to the application layer. On the contrary, SCTP needs an end-to-end deployment and the use of specific applications build on top of SCTP socket.

### Which issues have still not been addressed?

Adapting TCP to the context of handover would require changing every transport protocol stack already deployed to support these changes. In addition, defining the parameters of these TCP modifications (new values such as RTO and congestion window: *cwnd*, adapted to the new link) require a probing delay or a certain level of a-priori knowledge on the links characteristics which may not always be available. Finally they do not take advantage of the diversity of the links. The main drawback related to the solutions based on multipath SCTP is that they require modifying the interface between the application and the transport layer. Additionally, to date, the deployment of SCTP has been limited to Unix-like hosts. This requirement would make the deployment of these proposals

harder. The bi-casting proposals, even when they involve FEC coding, result in halving the available bandwidth which might already be limited in the case of WWAN interfaces. Further, with the exception of the bi-casting procedure, none of these proposals seem to perform well in challenged conditions with high PLR common before and after handover due to poor signal quality.

To the best of our knowledge, there currently exists no solution allowing to benefit from the robustness and bandwidth aggregation provided by the multihoming capability of new devices with an application on top of the standard TCP/IP protocol suite (i.e. without any modification of the suite).

All these facts motivate our proposal described in the following section 3.

### 3 Our proposal

We present in this section the architecture and internal mechanisms that define our framework proposal.

#### 3.1 Architecture, coding and handover




Fig. 2. How to plug our coding scheme in the network protocol stack.

Our coding scheme, detailed in the next section, allows recovering in a timely manner from all the losses that occurred on a path, regardless of their distribution. The sole requirement is that in average the amount of redundant packet sent must be greater than the amount of the losses. The key idea is to use our proposal to enforce the part of a path that can be affected by PLR losses during a handover. The whole path can be protected by coding and decoding the packets at both ends of the connection if the Correspondent Node (CN) is aware of the different addresses of the Mobile Node (MN). Otherwise, as Fig. 2 suggests, in the case where multihoming is provided by an IP-level mobility solution, the coding/decoding can also be done between the Access Router (AR) and the MN. We propose to use our coding scheme at layer-3, thus hiding the losses to the transport layer. Fig. 2, describes a possible way to plug our coding scheme

through the use of Divert Socket. In our case, we used the BSD implementation of divert socket which is also provided under GNU/Linux with `ipchains` API.

The sending of a data segment from the CN to the MN works as follows: as a first step, packets travels normally through the TCP/IP protocol stack as the CN is not involved in the coding process. Packets that reach the Access router cross the IP forwarding rules (IPFW) which diverts (second step) the packets destined to the MN to the related Tetrys instance (there might be one instance by MN supported by the AR). Tetrys adds the packets to its encoding window and re-injects them (step 3 and 4) adding a packet sequence number (three bytes might be more than necessary) plus a bit to distinguish redundancy from source data packets inside the IP option field. Redundancy (i.e. coded) packets are injected with such an IP destination address that they go through the WWAN (step 4) or WLAN (step 5) links depending on whether it is an upward or downward vertical handover. The size of these coded packets is equal to the maximum size of the data packets currently in the Tetrys encoding window. Packets reach the MN through the different interfaces and are diverted to Tetrys (step 6) which decodes and rebuilds any lost packets. The whole packets received by the Tetrys encoder at step 2 are re-injected (step 7) without losses and ordered in-sequence. Finally at step 8, packets are transmitted to the transport and application layers in a transparent manner.

When our coding scheme is used to improve the link quality (with two interfaces) the source data packets are sent on the fastest (which is also the more lossy one in our experimental scenario) interface while the redundancy (coded) packets are sent through the WWAN.

During upward handover, the PLR is monitored and when it exceeds a given threshold (70% in the experiment), a coded version is sent over the WWAN for each data packet received from the source. When the WLAN is definitely out of range, packets are sent uncoded over the WWAN. During a downward handover, all the source data packets are sent coded over the WWAN and uncoded over the WLAN. When the WLAN PLR decreases below a threshold, the coded packets will be sent according to the redundancy ratio through the WWAN only until the PLR becomes negligible for TCP.

### 3.2 The Tetrys on-the-fly coding scheme

The Tetrys sender uses an *elastic encoding window* (denoted  $W_{sender}$ ) which includes all the source packets sent and not yet acknowledged. Let  $P_i$  be the source packet with sequence number  $i$ . Every  $k$  source packets, the sender sends a (single) repair packet  $R_{(i..j)}$ , which is built as a linear combination (with random coefficients) of all the packets currently in  $W_{sender}$ . The receiver is expected to periodically acknowledge the received or decoded packets, and each time the sender receives an acknowledgment, the acknowledged packets are removed from  $W_{sender}$ . A receiver can decode lost packets as soon as the number of available repair packets is higher or equal to the number of lost packets (the lost packets are detected by the gaps they introduce in the sequence number of the received packets). Fig. 3 illustrates this principle. In the figure  $k = 2$ , which means that a




Fig. 3. Tetrys principle.

repair packet is sent each time two source packets have been sent. The right side of this figure shows the list of packets that are lost and not yet rebuilt, as well as the repair packets kept by the receiver in order to recover them. During this data exchange, packet  $P_2$  is lost. However, the repair packet  $R_{(1,2)}$  successfully arrives and allows to rebuild  $P_2$ . The receiver sends an acknowledgement for packets  $P_1$  and  $P_2$ , in order to inform the sender that it can compute the next repair packets from packet  $P_3$ . Unfortunately this acknowledgement is lost. However, this loss does not compromise the following transmissions and the sender simply continues to compute repair packets from  $P_1$ . After this, we see that  $P_3$ ,  $P_4$  and  $R_{(1,4)}$  packets are also lost. These packets can be rebuilt using  $R_{(1,6)}$  and  $R_{(1,8)}$  since the number of repair packets becomes higher or equal to the number of lost packets.

The acknowledgement path is only used to optimize the encoding process and is not mandatory during a handover of a few seconds. However, these acknowledgements contain information about the PLR experienced by the MN that might be used to update or tune the redundancy ratio when the handover takes a longer time (we consider higher than 30sec). If the role of the coding scheme is to make the transport layer more robust, the redundancy ratio does not require to be frequently adjusted while coding for real-time application would need more accurate adaptation.

Unlike Tetrys, most of the forward error codes (FEC) used over packet erasure channels are block codes [8]. 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. 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.

As a result and compared to block codes:

- Tetrys is tolerant to any burst of source, repair or acknowledgement losses, as long as the amount of redundancy exceeds the PLR;
- the lost packets are recovered within a delay that does not depend on the  $RTT$ ;



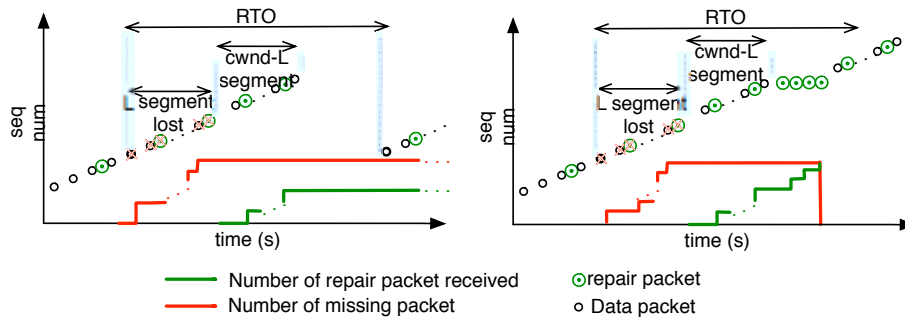


Fig. 4. Mechanism to prevent a blocking TCP window.

- the configuration is much easier and more robust to network variation than configuration for a block code. This is a key point in the context of handover; These properties make Tetrys a perfect candidate to reduce packet loss and recovery delay during a handover process.

### 3.3 Redundancy emission/allocation and interaction with TCP

TCP is well-known for bad performance over lossy links as every lost packet is considered by TCP as an indication of congestion. A possible solution to mitigate this effect would be to use a FEC mechanism to correct losses due to error link. However in [1]<sup>1</sup>, the authors show that the joint use of end-to-end FEC with TCP does not solve the problem in case of significant PLR. This is due to the fact that TCP needs in-order delivery of data packets and is also strongly sensitive to RTT variations which trigger spurious timeouts resulting in a decreased throughput. A spurious timeout occurs when a non lost packet is retransmitted due to a sudden RTT increase (typically when the mobile node moves from a WLAN to WWAN) which implies an expiration of the retransmission timer set with a previous, and thus outdated, RTT value.

In a previous work [15], we have already shown that this code protects efficiently real-time traffic such as Voice over IP and video-conferencing over links with high PLR. Even if TCP behaves and performs better above Tetrys than above FEC as previously described, rebuilding a burst of  $L$  lost packets requires receiving at least  $L$  redundant packets. This means that if within a congestion window (and thus a RTT) there were more than  $R \cdot cwnd$  lost packets, the decoder may not be able to rebuild the packets, freezing the connection until the emission of at least  $\frac{L - R \cdot cwnd}{R}$  TCP retransmission. Needless to say, during our experiment we found that over Tetrys, TCP often entered into backoff mode and the connection stalled. To prevent such phenomenons, we propose to increase the redundancy based on a temporal frequency instead of a fixed ratio.

<sup>1</sup> In their scheme, TCP is modified to ignore losses. In our case, we assume a complete separation between the coding layer and the transport protocols.

The sole requirement to correct errors timely while minimising extra transmissions is thus to size the frequency so that  $\frac{1}{RTT} > f_r \geq \frac{L-R \cdot cwnd}{4 \cdot RTT}$ , with  $f_r$  the minimal frequency for the emission of redundancy packet (assuming the RTO is roughly four times the RTT). As explained in Fig. 4, the sending of 1 repair packet every  $k$  source data packets (left subfigure),  $L$  data losses would require  $k * L$  more data packets to be sent by TCP. As TCP cannot send more than  $cwnd - L$  data packets, if  $k * L > cwnd - L$ , RTOs may be triggered and the connection may stall. As the subfigure on the right suggests, this problem can be solved by sending repair packets at a minimal rate  $f_r$  when the TCP window is abnormally stalled (e.g. when the throughput drops below  $k * f_r$ ).

## 4 Evaluation

Our testbed architecture is similar to the one presented in figure 2, except for the WLAN and WWAN links which are emulated with Netem [6] on the top of two Ethernet links. The CN and AR are connected with a 10Mbit/s Ethernet link with a negligible transmission delay. Default settings assume 10Mbit/s and an RTT of 10ms for the WLAN and 4Mbit/s with a RTT of 100ms for WWAN (this is realistic for UMTS HSDPA cellular networks). The default redundancy ratio is set to 20%.

### 4.1 Comparison with FEC over lossy links

PLR	0.0	0.5	2	4	5
	8	10	12	16	20
TCP/Tetrys (4,5)	7.78/0.01	7.81/0.01	7.81/0.01	7.80/0.01	7.81/0.00
	7.82/0.01	7.81/0.01	7.81/0.01	7.18/0.02	4.6/0.26
TCP/FEC (4,5)	7.79/0.02	7.83/0.00	6.48/1.2	3.01/0.6	2.68/2.24
	Timeout				
TCP/FEC (8,10)	7.81/0.01	7.78/0.05	7.81/0.03	6.3/1.48	Timeout
	Timeout				
TCP/FEC (12,15)	7.82/0.02	7.82/0.01	7.79/0.03	7.54/0.10	4.06/5.05
	Timeout				
TCP/FEC (16,20)	7.81/0.01	7.81/0.02	7.82/0.01	7.825/0.01	Timeout
	Timeout				

**Table 1.** Throughput/Std. dev. in Mb/s for  $R = 0.2$ ,  $B_D = 10Mbps$ ,  $B_R = 4.0Mbps$  with data and repair packets sent on WLAN only.

Tab. 1 shows the throughput obtained by TCP on top of Tetrys or FEC when packets are sent over a single lossy link (over WLAN only). We can see that both Tetrys and FEC can handle a low loss rate efficiently maintaining a throughput of 8Mbit/s for the TCP flow. Similarly to previous work on TCP/FEC [1], we observe that with a significant loss rate, the TCP/FEC throughput decreases and the connection often stalls. In contrast to this, TCP over Tetrys is not severely impacted by these loss rates, and in fact the TCP throughput starts to decrease

PLR	0.0	0.5	2	4	5
	8	10	12	16	20
TCP/Tetrys (4,5)	9.54/0.00	7.69/0.18	6.5/0.49	8.18/0.34	9.02/0.05
	8.55/0.7	8.66/0.3	8.45/0.57	7.10/0.31	5.12/0.2
TCP/FEC (4,5)	9.53/0.00	5.96/0.02	3.07/0.08	1.13/0.6	1.25/0.14
	0.17/0.05	Timeout			
TCP/FEC (8,10)	9.54/0.00	7.19/0.08	4.72/0.32	2.71/0.46	1.85/0.40
	Timeout				
TCP/FEC (12,15)	9.55/0.00	7.7/0.11	4.79/1.36	3.73/1.13	Timeout
	Timeout				
TCP/FEC (16,20)	9.53/0.00	7.37/0.65	5.83/0.68	3.73/1.13	0.84/1.18
	Timeout				

**Table 2.** Throughput/Std. dev. in Mb/s for  $R = 0.2$ ,  $B_D = 10Mbps$ ,  $B_R = 4.0Mbps$  with data sent on WLAN and repair sent on WWAN only.

only after  $PLR = 14\%$ . We have to remark that the code rate is fixed during the experiments. As a matter of fact, there would not a decrease of the TCP throughput by dynamically adjusting the code rate as a function of the PLR.

Tab. 2 shows the throughput obtained by TCP on the top of Tetrys or FEC when data packets are sent over the (10Mbit/s, 10ms) WLAN link lossy link and the coded packets over the (4Mbit/s, 100ms) WWAN link.

We can notice that the results of TCP/FEC are even worse than in the one link only experiment (Tab. 1). We can make the same observation for TCP/Tetrys under small PLR (for 0.5% or 2%). This is explained by the delay asymmetry between the two links and the TCP *cwnd* which fits the bandwidth delay product (BDP) corresponding to the WLAN link. When losses occur, their reconstruction requires to wait for the coded packet that arrives 90ms later. During this time, no packet reaches the TCP receiver and thus there is no acknowledgement sent to the TCP sender that would slide the congestion window. The RTT perceived by TCP increases with the PLR and the *cwnd* also increases until it reaches the BDP corresponding to the slowest link. This explains the poor performance<sup>2</sup> of TCP/Tetrys for small PLR and the improvement observed when the PLR is higher. These two facts: 1) the link delay asymmetry, and 2) the few losses not recovered by FEC, impact the TCP throughput over FEC more significantly.

In spite of our testbed not enabling bandwidth aggregation (as in SCTP), these results show that in contrast to previous coding proposals, Tetrys allows transport protocols such as TCP to remain efficient in spite of the deteriorated link conditions during and around a handover.

## 4.2 Illustration of the mechanism during a handover scenario

Fig. 5 shows the results for various WWAN bandwidths (300kbit/s, 1Mbit/s and 4Mbit/s) It takes 0.5 second to the WLAN link to change from "up" state to "down" state and the same for the opposite transition. We can see that even if

<sup>2</sup> This can be solved by artificially delaying the packet at the speed of the slowest interface.




Fig. 5. Handover scenario with various WWAN bandwidth and (10Mb/s, 10ms) WLAN.




Fig. 6. Handover scenario with various redundancy ratio with (10Mb/s, 10ms) WLAN and (4Mb/s, 100ms) WWAN.

these various parameters have an impact on the TCP throughput, they do not significantly impact the amount of time required by TCP to reach its average throughput with Tetrys. Similar results hold for different values of the WWAN RTT.

Fig. 6 shows the impact of the redundancy ratio ( $R = 1/5, R = 1/6, R = 1/8, R = 1/9$ ) on TCP during a handover. In this case, it takes 10 seconds for the WLAN PLR to switch from 0 to 100% and inversely. The three sub-figures show different runs of the experiment. We can see that the configuration of the Tetrys redundancy ratio does not require to be timely adjusted as there is no impact on the throughput achieved by TCP.

Compared to block codes (characterized by a specific FEC coding configuration) where we would have to dynamically reconfigure the redundancy parameters ( $(k, n)$ ) as a function of the size of the loss burst, Tetrys is resistant to any kind of loss burst patterns and does not need to be dynamically adjusted (as already highlighted Section 3.2 illustrated and Tab. 1, 2). Furthermore, although the increase of the redundancy parameters allows to correct larger burst of losses, they trigger TCP timeout as the decoding process can be longer than the RTO value.

This last result illustrates that compared to block codes (such as a specific FEC coding configuration), Tetrys is resistant to any kind of loss burst patterns.

## 5 Conclusion

In this paper, we evaluate the benefits of using an “on the fly” coding scheme to reduce packet losses during a soft vertical handover due to low signal quality. The experimental evaluation suggests that the use of this type of coding scheme may be an interesting complementary strategy to vertical handover management protocols due to its fast configurability and in the context of multipath communications. Our experiments clearly show that this coding scheme allows to maintain the TCP throughput during a handover by taking advantage of the multiple wireless interfaces present in today’s smartphones. Particularly, results show that it significantly improves the quality of TCP flows in terms of delivery ratio. As a next step, we are planning to integrate the implementation of this error recovery algorithm called Tetrys as a part of the SafetyNet architecture and evaluate the performance empirically using our SafetyNet implementation.

## References

1. T. Anker, R. Cohen, and D. Dolev. Transport layer end-to-end error correcting, 2004. Tech. Rep. School of Engineering and Computer Science, Hebrew University.
2. R. Chakravorty, P. Vidales, K. Subramanian, I. Pratt, and J. Crowcroft. Performance issues with vertical handovers - experiences from GPRS cellular and WLAN hot-spots integration. In *IEEE PERCOM*, 2004.
3. L. Daniel and M. Kojo. The performance of multiple TCP flows with vertical handoff. In *ACM MobiWAC*, 2009.
4. H. Liu et al. A staggered FEC system for seamless handoff in WLANs: Implementation experience and experimental study. In *IEEE ISM*, 2007.
5. M. Fiore and C. Casetti. An adaptive transport protocol for balanced multihoming of real-time traffic. In *IEEE GLOBECOM*, 2005.
6. S. Hemminger. Network emulation with netem. In *Australia’s national Linux conference (LCA)*, Canberra, Australia, 2005.
7. S.-E. Kim and J. A. Copeland. TCP for seamless vertical handoff in hybrid mobile data networks. In *IEEE Globecom*, 2003.
8. S. Lin and D. Costello. *Error Control Coding: Fundamentals and Applications*. Prentice-Hall, Englewood Cliffs, NJ, 1983.
9. H. Matsuoka, T. Yoshimura, and T. Ohya. A robust method for soft IP handover. *IEEE Internet Computing*, 7:18–24, 2003.
10. K. Nagami et al. Multi-homing for small scale fixed network using mobile IP and NEMO. RFC 4908, 2007.
11. C. Ng et al. Analysis of multihoming in network mobility support. RFC 4980, October 2007.
12. H. Petander and E. Lochin. Safetynet version 2, a packet error recovery architecture for vertical handoffs. In *ICST MONAMI*, Santander, Spain, 2010.
13. H. Petander, E. Perera, and A. Seneviratne. Multicasting with selective delivery: a safetynet for vertical handoffs. *Wirel. Pers. Commun.*, 43(3):945–958, 2007.
14. R. Stewart et al. Stream Control Transmission Protocol. RFC 4960, 2007.
15. P. U. Tournoux, A. Bouabdallah, J. Lacan, and E. Lochin. On-the-fly coding for real-time applications. In *ACM Multimedia*, 2009.
16. J. Wang, J., and X. Zhu. Latent handover: A flow-oriented progressive handover mechanism. *Computer Communications*, 31(10):2319–2340, 2008.