

Title	Low-latency communications in LTE using spatial diversity and encoding redundancy
Author(s)	Kucera, Stepan; Yu, Yu; Buddhikot, Milind; Lim, Yuto
Citation	2016 IEEE 84th Vehicular Technology Conference (VTC-Fall): 1-6
Issue Date	2016-09-18
Type	Conference Paper
Text version	author
URL	http://hdl.handle.net/10119/14284
Rights	This is the author's version of the work. Copyright (C) 2016 IEEE. 2016 IEEE 84th Vehicular Technology Conference (VTC-Fall), 2016, 1-6. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.
Description	



Low-Latency Communications in LTE Using Spatial Diversity and Encoding Redundancy

Stepan Kucera*, Yu Yu[†], Milind Buddhikot[†], Yuto Lim[‡]

*Nokia Bell Labs, Dublin, Republic of Ireland

[†]Nokia Bell Labs, Murray Hill, New Jersey, USA

[‡]Japan Advanced Institute of Science and Technology, Ishikawa, Japan

Abstract—Control of data delivery latency in wireless mobile networks is an open problem due to the inherently unreliable and stochastic nature of wireless channels. This paper explores how the current best-effort throughput-oriented wireless services could be evolved into latency-sensitive enablers of new mobile applications such as remote 3D graphical rendering for interactive virtual/augmented-reality overlay. Assuming that the signal propagation delay and achievable throughput meet the basic latency requirements of the user application, we examine the idea of trading excess/federated bandwidth for the elimination of non-negligible data re-ordering delays, caused by temporal transmission failures and buffer overflows. The general system design is based on (i) spatially diverse delivery of data over multiple paths with uncorrelated outage likelihoods, and (ii) forward packet protection based on encoding redundancy that enables proactive recovery of lost or intolerably delayed data without end-to-end re-transmissions. Our analysis is based on traces of real-life traffic in live carrier-grade LTE networks.

Index Terms—LTE, data delivery latency control, spatial diversity, encoding redundancy, random linear codes, forward error correction and protection, multi-path TCP, re-transmissions

I. INTRODUCTION

The first and second generations of cellular networks for mobile communications provided users with basic connectivity using analog and digital technology, respectively. In addition to wireless voice telephony, the third-generation networks such as the CDMA-based HSPA systems enabled also mobile Internet access. The fourth generation of cellular networks represented by the current OFDMA-based LTE systems quickly evolved into a true mobile broadband solution for heterogeneous high-throughput traffic consisting of voice, video, and Internet data.

Each new generation of networking standards offered a technologically new approach to the evolving needs of mobile users but without any quality of service (QoS) guaranteed by the network [1]. Yet the opening gap between the growths of revenue and mobile data traffic [2] brought general interest to enabling new latency-sensitive applications such as steering and control of real and virtual objects (Tactile Internet), remote 3D graphical rendering for virtual/augmented-reality overlay, remote traffic control of self-driving cars and drones, closed-loop control of industrial processes, and gaming [3], [4].

Assuming that the fundamental propagation delay of the infrastructure due to non-zero physical signal propagation is within the user application tolerance, the goal of this paper is to define and evaluate a mechanism for minimizing data re-ordering delays. These are caused by re-transmissions after

failures of basic transmission processes such as decoding, integrity checks, error recovery, segment reassembly, as well as buffer overflows. For example, the reporting of channel quality in LTE networks targets block error rate of 10% at the physical layer, implying a strong default dependence on re-transmissions for error-free data transfer.

More concretely, we examine two approaches to re-ordering delay mitigation - spatial diversity and encoding redundancy.

Spatial diversity refers to the delivery of a data flow over multiple parallel wireless links (paths). The aggregation of multiple independent wireless links, or data delivery paths, into one logical connection theoretically allows increasing the overall throughput and reliability¹ as well as reducing latency by resource pooling - an advantageous feature when single-path connections cannot meet the user demands.

Modern mobile devices, such as smart phones and tablets, already support multi-technology multi-band networking: they are commonly equipped with multiple LTE, HSPA, WiFi transceivers operational in multiple licensed and unlicensed frequency bands. A new fifth-generation air interface as well as the allocation of new (shared) frequency bands can be expected in the near future [6]. The ongoing network densification by small-cell deployments improving the overall connectivity also clearly motivates a study in this direction.

By *encoding redundancy* is understood the injection of additional protection data into the actual flow of payload data to compensate for unavoidable errors or outages of the wireless links. In the latency control context, the idea is to establish forward packet protection (FPP) that allows recovering lost or unacceptably delayed payload data without reactive re-transmissions. These increase the delivery delay by at least $1.5\times$ of the connection round-trip time as well as cause throughput drop in protocols with loss-based congestion control. The simplest example of encoding redundancy usage consists in primitive replication of a data flow over multiple parallel paths; techniques more efficient in terms of the overhead are discussed subsequently.

The paper organization and contributions can be summarized as follows. Section II summarizes the considered architecture for minimizing re-ordering delays in wireless links by using forward packet protection on the basis of random linear

¹Next-generation air interfaces may be prone to outages as the user body can efficiently block the propagation of the anticipated mm-wave carriers [5].




Fig. 1. Latency-sensitive LTE-based services are protected by using a parallel WiFi connection for delivering FPP data.

codes. The trade-off between overhead and latency in macro-cell and small-cell LTE networks is quantified in Section III using real-life traffic traces. In Section IV, performance limits of the method are analyzed and design guidelines formulated. Conclusions follow in Section V.

To our best knowledge, this is the first practical study on latency-sensitive services in (LTE) cellular networks. Examples of latency control in proprietary networks include [7], [8]. Within the domain of standardized systems, potential latency reduction gains assuming fundamental modifications of the LTE physical layer are studied in [9]. In terms of medium access, encoding redundancy is used to create novel LTE HARQ mechanisms in [10], [11], [12]. On the network and transport layers, [13] examines opportunistic injections of encoded data into single-path WiFi transmissions. Novel TCP implementation using encoding redundancy principles with delayed block-level decoding is defined in [14]. A faster but difficult-to-configure scheme for proactive TCP re-transmissions of unacknowledged data is discussed in [15].

II. SYSTEM ARCHITECTURE

A. Bandwidth Aggregation

Our approach to latency control in mobile user services over a primary wireless link is based on trading latency for federated L excess bandwidth. More specifically, forward packet protection (FPP) data - to be defined subsequently - is sent over an independent secondary link, simultaneously with the primary payload data. The FPP data is used to recover lost or intolerably delayed payload data to prevent further error propagation, e.g. to congestion control or application layer. When generalizing this atomic setup to multiple links carrying payload and/or FPP data, the multi-path scheduler may cause data re-ordering due to inaccurate feedback [16].




Fig. 2. Formation of FPP packets using random linear codes.

The LTE dual connectivity and carrier aggregation could be used to implement the secondary links. Yet its independence from the primary link, ensuring decoupled failure likelihoods, may be violated by traffic QoS classification by the LTE Service Data Flow processes and their mapping onto bearers.

It is therefore proposed to create the secondary link by using an additional wireless interface of the mobile device. As shown in Fig. 1, WiFi could be considered as the currently best option given its ubiquitous deployment, good performance, and typically free-of-charge access. Regarding link independence, it is to be noted that WiFi access points are often backhauled via cellular networks, especially mobile WiFi in transport vehicles.

The radio interface management in mobile devices can be done by using the network-resident service that we demonstrated in [17]. This service makes central optimized decisions based on network-wide information and operator policies. Its implementation is based modified multi-path TCP, an IETF-standardized backward compatible solution for seamless transport-layer integration of networking technologies [18].

B. Forward Packet Protection

A naïve way of protecting payload data consists in their replication over L parallel links. The fastest link determines the overall latency at the expense of $(L - 1)$ -fold overhead.

To compress the overhead to more reasonable levels, it is proposed to construct each FPP packet as a weighted random linear combination of N payload packets by using binary XOR operation [19]. The multiplicative weights are selected at random in an identical, independent, uniform manner [20].

Fig. 2 visualizes the simplest protection scheme under which M FPP packets are generated for each N consecutive payload packets. In the figure example, $M = 1$ and $N = 3$.

As depicted by Fig. 3, the decoding of a payload packet missing by a data delivery deadline can be done in real time by using the received generator packets and Gaussian elimination [20], [19]. If the combining weights ensure full rank of the Gaussian system, one FPP packet can be used to recover one missing payload data packet or a part thereof. Thus, the M/N overhead defines the encoding protection strength.

From a system point of view, one could require the overhead data rate not to exceed the payload data rate, i.e. $M/N \leq 1$, as payload data duplication occurs for $N = M$ or simply $N = 1$ (100% overhead). Yet in general, the instability of basic link characteristics such as effective throughput, delay jitter, and

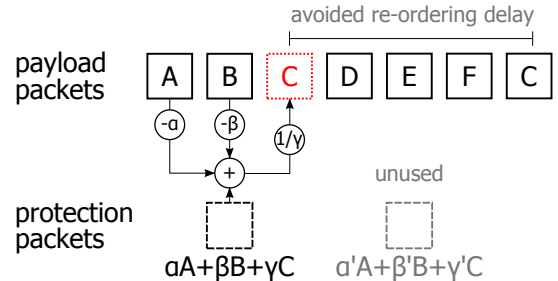


Fig. 3. Recovery of a delayed payload packet using FPP packets.

cross-traffic volume may justify a more complex design based on adaptive joint scheduling of FPP and payload data without an *a priori* primary/secondary characterization [16] but multi-path scheduling issues are out of the scope of this contribution.

The above description implicitly assumes protection against data loss or delay at the transport layer of the OSI model, typically accommodating the connection-oriented TCP and the connection-less UDP. Accordingly, random linear combinations of TCP/UDP packets are sent as the secondary FPP information, and used by the receiver for recovery of data before they are passed on to the TCP/UDP modules.

Encoding protection can be applied to lower OSI layers as well. The advantage consists in smaller size of FPP packets and earlier correction of errors, but only shorter components of the overall end-to-end delay can be eliminated. But higher-layer errors would not be detectable nor correctable.

III. QUANTIFICATION OF LTE DATA RE-ORDERING

This section summarizes the frequency and length of payload data re-orderings as well as the size of affected data blocks as measured in both macro-cell and small-cell LTE networks. The trade-off between the FPP overhead of the secondary link and the achievable re-ordering delay reduction at the primary LTE link is analyzed subsequently.

A. Measurement Scenarios

LTE re-ordering events are quantified based on TCP data captured in live LTE networks in the wider New York area, USA. The choice of transport-layer protocol is irrelevant for tracking re-orderings caused by the LTE protocol stack underneath, the main focus of this study. Yet TCP is generally more suitable than UDP for minimizing re-orderings caused by the transport layer itself. Unlike UDP, TCP relies on explicit flow and congestion control which reduces the probability of data loss due to buffer overflows whenever the sender-side data rate exceeds the fluctuating LTE link capacity. Nevertheless, our measurements also incorporate TCP with capped throughput to account for the case of constant-rate UDP.

Extensive data traces at the order of tens of GB were obtained in the following real-life scenarios:

1) *Macro-cell (MC) Downlink (DL)*: Test user equipment (UE) connects over a commercial MC LTE network of a major carrier to a remote server from both outdoor and indoor locations to download data using FTP. The server is located in physical proximity to the LTE public gateway to minimize the effects of the Internet. Two subscenarios are distinguished:

- *MC DL unlimited* - The LTE data rate is unlimited. Background download/upload traffic is low to moderate (off-peak hours in locations outside New York, USA).
- *MC DL limited* - The LTE data rate is operator-limited to 10% of possible data rates. Background download/upload traffic is high (peak-hour in downtown New York, USA).

2) *Small-cell (SC) Downlink/Uplink (DL/UL)*: Indoor test UEs connect over a carrier-grade SC LTE network in a large private enterprise in the USA to a server to download and/or upload data using FTP. The server is connected directly to the




Fig. 4. Visualization of the characteristics of a re-ordering event.

LTE public gateway, located on the enterprise site. Background traffic is low to moderate and the service LTE data rate is unlimited. Two subscenarios are distinguished:

- *SC UL/DL static* - UEs are static and no handovers occur.
- *SC UL/DL handovers* - UEs are mobile and at least one (typically multiple) handovers occur during each session. UEs move at slow to faster walking speeds.

Four LTE USB modems LG VL600, Qualcomm 9630, LG G7, Huawei e3276s-150 are used to connect FTP clients operated under Ubuntu Linux (14.04.3 LTS) and Microsoft Window 7 to an FTP server based on Ubuntu Linux 14.04.3 LTS as well. The default TCP CUBIC implementation of TCP is used for Ubuntu-based devices. Repetitive large data transfers of 0.1-1GB are used to emulate long-lived sessions. The maximum TCP transmission unit is always 1388 bytes.

B. Evaluation Methodology

The following formalization visualized in Fig. 4 is used to capture the frequently complex LTE re-ordering events.

Let $P(T, S)$ denote a TCP packet received in time T and carrying a range $S(T) = [S_1(T), S_2(T)]$ of TCP sequence numbers ($S_2 - S_1 + 1 = 1388$ bytes in all scenarios). TCP sequence numbers index individual bytes of ordered payload data. Packets arrive in unique times and each sequence number is received only once as TCP discards duplicates.

A packet $P(T, S)$ is *in order* if and only if (i) no sequence number higher than $S_2(T)$ has been received until time T (i.e., no packets expected after time T have been received), and (ii) the union of all sequence numbers of all previously received packets $p(t, s) \forall t < T$ is a contiguous range from 1 to $S_1(T) - 1$, i.e. $\cup_{\forall t < T} \arg_s p(t, s) = [1, S_1(T) - 1]$ (i.e., no packets expected before time T are missing).

A packet $P(T, S)$ is *out of order* if it is not in order.

A *re-ordering event* is defined as a time-wise contiguous sequence of out-of-order packets excluding any in-order packets.

The term *ideal in-order delivery* refers to how packets would have been received if no re-ordering happened (see Fig. 4). Given a re-ordering event characterized by a sequence $P(T^i, S^i) \forall i \in [1, \dots, N]$ of out-of-order packets, the ideal in-order delivery is defined as a sequence of packets $P(T^i, \hat{S}^i)$ that (i) are received in the original times T^i , and (ii) have monotonically increasing sequence numbers from the set $\cup_{\forall i} S^i(T^i)$, i.e. $S_2^i(T^i) + 1 = S_1^{i+1}(T^{i+1}) \forall i \in [1, \dots, N - 1]$.

By *delayed ideal in-order delivery* is then understood a shift of all ideally in-order delivered packets $P(T^i, \hat{S}^i)$ in time by delay D , i.e. a packet sequence $P(T^i - D, \hat{S}^i) \forall i \in [1, \dots, N]$.

Assuming maximum tolerable delay D , an out-of-order packet $P(T^i, S^i)$ is *effectively lost* during a re-ordering event if the corresponding packet $P(T^i - D, \hat{S}^i = S^i)$ in the ideal in-order delivery sequence delayed by D would have been received earlier than the out-of-order packet $P(T^i, S^i)$ itself, i.e. $\arg_T P(T^i - D, \hat{S}^i = S^i) < \arg_T P(T^i, S^i)$ for given i .

The captured LTE traffic traces were analyzed in terms of:

- *average inter-event time separation* - average time interval between two consecutive re-ordering events. Assuming conservatively a limited resolution of the user in distinguishing consecutive events, multiple short re-ordering events occurring within a period of 0.1 second are counted only as a single event.
- *data re-ordering delay* - minimum delay D by which the ideal in-order delivery sequence $P(T^i - D, \hat{S}^i)$ of a re-ordering event must be delayed such that none of the out-of-order packets is considered effectively lost, i.e.

$$D = \arg_D \max_{i \in L} \left[\arg_T P(T^i - D, \hat{S}^i) - \arg_T P(T^i, S^i) \right] \quad (1)$$

for $\hat{S}^i = S^i$ where L is the set of out-of-order packets effectively lost for delay $D = 0$.

- *re-ordered data block size* - number of out-of-order packets considered lost during a re-ordering event with respect to the non-delayed ($D = 0$) ideal in-order delivery sequence. It is the cardinality of the above set L .
- *physical payload data rate* - average data rate of the ideal in-order delivery sequence during a re-ordering event.

C. Measurement Results

Figs. 5, 6 and 7 show the cumulative distribution functions of measured data re-ordering delays, re-ordered data block sizes, and physical payload data rates for each LTE scenario. Fig. 5 also indicates the average time between two consecutive re-ordering events assuming user resolution of 100 ms.

It is observed that higher average physical data rate of an LTE scenario implies more frequent re-ordering events characterized by generally shorter re-ordering delays.

For example, reducing the LTE macro-cell downlink data rates to around 10% of the achievable rates (“MC DL unlimited” vs “MC DL limited” scenarios) increases the re-ordering event period from 79s to 34s but reduces around 10-times the re-ordering delays. In the small case, uplink data rates are often higher than the downlink rates (due to minimum uplink traffic) but always order-of-magnitude higher than in the macro cell case (better SINR). At the same time, the re-ordering event period is reduced to units of seconds (2–7s) but the overall re-ordering delays are further reduced by nearly another order of magnitude. Small-cell handovers have low impact on the overall data rate but cause major re-ordering delay increase in the faster uplink case.

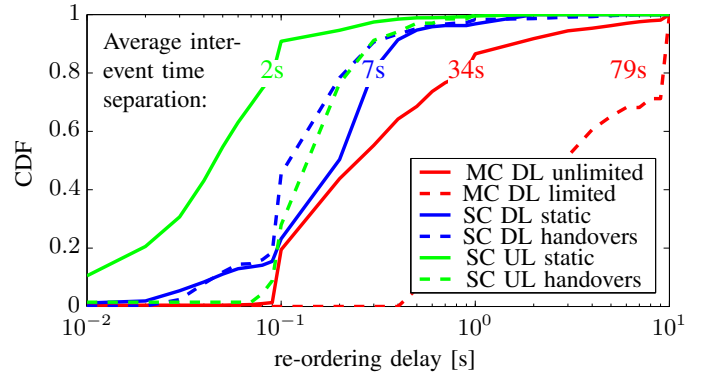


Fig. 5. Cumulative distribution function of data re-ordering delays with indicated average inter-event separation times (100 ms event resolution).

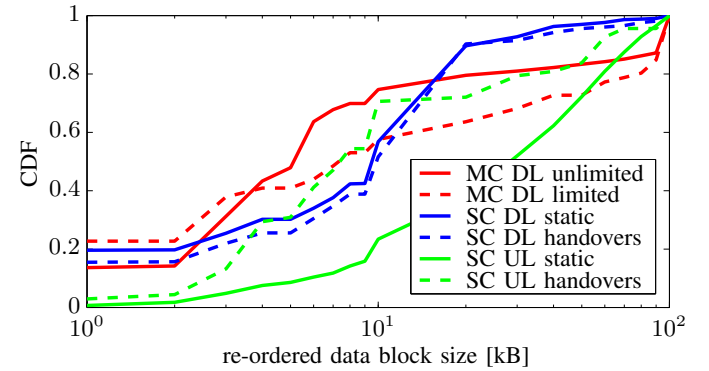


Fig. 6. Cumulative distribution function of re-ordered data block sizes.

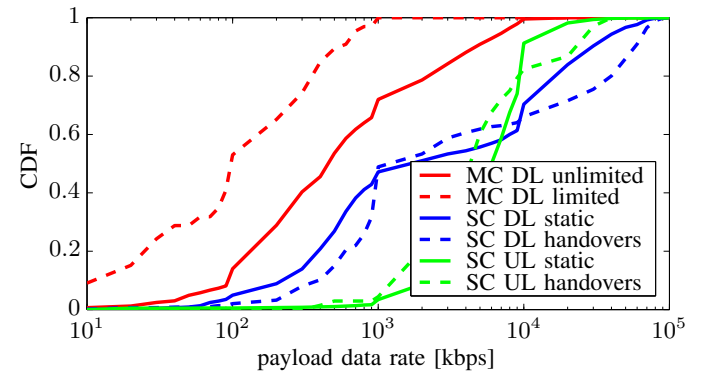


Fig. 7. Cumulative distribution function of physical payload data rates.

It may also be argued that the re-ordered data block size is somewhat smaller for higher data rates too, but the rather monotonically increasing and often overlapping cumulative distributivity functions suggest random re-ordered block sizes.

IV. ANALYSIS OF TRADE-OFF BETWEEN FPP OVERHEAD AND RE-ORDERING DELAY

The trade-off between FPP overhead and primary link re-ordering delay consists in the fact that the replication of primary link data to the secondary link (100% overhead)

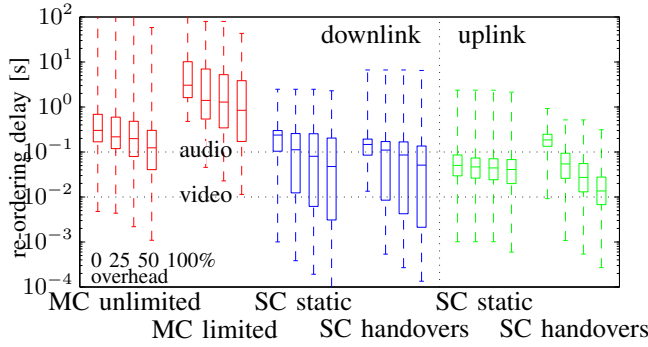


Fig. 8. Distribution of data re-ordering delays as function of constant-rate FPP, defined as a fraction of the physical LTE payload data rate. Range bars delimit the 25-th and 75-th percentiles as well as indicate the median.

minimizes the re-ordering delay while no FPP protection (0% overhead) implies inactive latency control.

To examine how LTE data streams benefit from the FPP, let us examine a scheme with a given constant FPP overhead that attempts to recover *all* the effectively lost packets during each re-ordering event. The reduced delays are computed as the minimum of (i) the data re-ordering delay D experienced during a re-ordering event, and (ii) the time needed for the delivery of FPP data whose size equals to the re-ordered data block size BS given an FPP data rate R^{FPP} , i.e. as

$$D^{\text{reduced}} = \min(D, BS/R^{\text{FPP}}), \quad (2)$$

where R^{FPP} can be defined either as a relative fraction of the physical payload data rate experienced during the re-ordering event (Fig. 8), or in absolute terms (Fig. 9).

It is observed from Fig. 8 that scenarios characterized by low physical data rates (both “MC DL” scenarios) or short re-ordering delays (“SC UL static” scenario) do not benefit from FPP. In the former case, FPP is not sufficiently available due to limited overhead data rate, while in the latter case, re-




Fig. 9. Probability that the FPP overhead data rate of the secondary link is lower than the payload data rate of the primary LTE link.




Fig. 10. Distribution of data re-ordering delays as function of constant-rate FPP from Fig. 8, assuming prevention of TCP buffer overflows. Range bars delimit the 25-th and 75-th percentiles as well as indicate the median.

ordering delay reductions of already short re-ordering delays are negligible. Only the small-cell downlink (“SC DL static” and “SC DL handovers” scenarios) exhibits the right balance of sufficiently high data rates and long-enough re-ordering delays that makes the impact of FPP the most significant.

The benefits of FPP could be clearly enhanced by increasing the data rate of the FPP overhead. In this context, Fig. 9 shows the probability that the FPP data rate of the secondary link is lower than the payload data rate of the primary LTE link. Nevertheless, the notion of primary and secondary link should be revisited when high relative overhead is required (comparable primary and secondary data rates), let alone when the secondary data rate exceeds the primary one.

We observe from the left-bottom part of Fig. 6 that effective losses of few packets that could be attributed to failures of the LTE physical/medium-access layers² account only for 5–20% of all losses in any studied LTE scenario. Such small effective losses can be recovered by using low constant-rate FPP overhead. However, with the notable exception of small-

²LTE coding and modulation for average signal-to-interference-and-noise ratio allows transmitting 1000s of bits over units of LTE resource blocks.




Fig. 11. Correlation of TCP round-trip time (RTT) with re-ordering events as indicated by TCP Selective Acknowledgements (SACK) (“SC DL static”).

cell downlink scenarios, the more typical losses of tens of packets either require a relatively high overhead for payload data recovery or simply make any data recovery impossible.

Major data losses can be explained by the “Unacknowledged RLC mode” which permits to drop unacknowledged data during a handover. Although it was not possible to confirm the RLC configuration in the measured commercial networks, it can be generally stated that this mode is used for broadcast over MCCH and MTCH using MBSFN or for VoIP. The more common TCP/IP traffic is typically handled in the RLC “Acknowledged Mode” under which unacknowledged packets are forwarded from the source eNodeB to the target eNodeB to ensure in collaboration with the LTE PDCP layer an in-order non-duplicate data delivery processing even during handovers.

Traces of eNodeB ingress and egress traffic show that large-scale data re-orderings observed in Fig. 6 are actually caused by buffer overflows in the LTE base stations. More specifically, loss-based mechanisms for TCP congestion control systematically cause buffer overflows by progressively increasing sending data rates to detect the available channel capacity under rapidly fluctuating networking conditions.

Fig. 11 provides exemplary evidence by depicting the correlation between buffer overflow times, indicated by TCP Selective Acknowledgement of out-of-order data, and TCP round-trip time build-up, a quantity directly proportional to queue occupancy in private queues typical for LTE and HSPA. Note that in such queues, cross-traffic cannot directly contribute to buffer occupancy or delay; it only affects the serving data rate.

Fig. 10 shows that preventing buffer overflows caused by the transport layer, e.g. by using delay-based congestion control as known from TCP Vegas or TCP Compound, the efficiency of FPP-driven recovery from unavoidable data losses within the LTE protocol stack is substantially better compared to the case with uncontrolled overflows from Fig. 8. In particular, we observe that substantially lower FPP overhead allows achieving major reductions of re-ordering delays, reaching an order of magnitude in small-cell networks.

V. CONCLUSION

We discussed a networking scenario in which payload data are delivered over a primary link while forward packet protection data are delivered in parallel over a secondary link. The protection data is used to mitigate re-ordering delays of the payload data without the need for lengthy reactive re-transmissions. Extensive measurements of live carrier-grade LTE networks in the wide New York area, targeting both small-cell and macro-cell deployments, are collected for the purpose of detailed performance analysis. The most notable reductions of re-ordering delay are observed in small-cell LTE networks. It is also shown that the transport layer largely determines the overhead-vs-delay trade-off. Forward packet protection in combination with buffer overflow prevention, e.g. via delay-based congestion control, is proposed to achieve even order-of-magnitude reductions of re-ordering delays for low overhead.

VI. ACKNOWLEDGEMENTS

The authors would like to thank to their colleagues Ricardo Garcia, Randolph Cohn, Glenn Steitz, and Edward Grinshpun for their technical expertise and dedicated support.

REFERENCES

- [1] G. M. 67, “Rp-150465: Study on latency reduction techniques for lte,” 3GPP Technical Specification Group Radio Access Network, Evolved Universal Terrestrial Radio Access (E-UTRA), Tech. Rep., Mar. 2015.
- [2] GSM Association. (2015) The mobile economy 2015. [Online]. Available: http://www.gsmamobileeconomy.com/GSMA_Global_Mobile_Economy_report_2015.pdf
- [3] N. Nikaen, R. Knopp, A. M. Cipriano, S. Krco, I. Tomic, P. Svoboda, M. Laner, E. Larsson, Y. WU, M. G. Fuertes, J. banos, N. Zeljkovic, and D. Marovic, “Low-latency in wireless communication,” in *Proc. of Vitel*, Slovenia, Apr. 2011.
- [4] G. P. Fettweis, “The tactile internet: Applications and challenges,” *IEEE Vehicular Technology Magazine*, vol. 9, no. 1, pp. 64–70, Mar. 2014.
- [5] C. Gustafson and F. Tufvesson, “Characterization of 60 GHz shadowing by human bodies and simple phantoms,” in *Proc. of IEEE Europ. Conf. on Antennas and Propagation*, Prague, Czech Republic, Mar. 2012.
- [6] 4G Americas. (2015) 5G spectrum recommendations. [Online]. Available: http://www.4gamericas.org/files/6514/3930/9262/4G_Americas_5G_Spectrum_Recommendations_White_Paper.pdf
- [7] D. Szabo, A. Gulyas, F. H. Fitzek, and D. E. Lucani, “Towards the tactile internet: Decreasing communication latency with network coding and software defined networking,” in *Proc. of IEEE European Wireless Conference*, Budapest, Hungary, May 2015.
- [8] M. Weiner, M. Jorgovanovic, A. Sahai, and B. Nikolic, “Design of a low-latency, high-reliability wireless communication system for control applications,” in *Proc. IEEE*, Sydney, Australia, Jun. 2014.
- [9] Ericsson. (2014) LTE latency improvement gains. [Online]. Available: <http://www.ericsson.com/research-blog/lte/lte-latency-improvement-gains/>
- [10] D. Vukobratovic, C. Khirallah, V. Stankovic, and J. S. Thompson, “Random network coding for multimedia delivery services in LTE/LTE-A,” *IEEE Transactions on Multimedia*, vol. 16, no. 1, pp. 277–282, Jan. 2014.
- [11] Z. Lv, K. Xu, and Y. Xu, “A practical HARQ scheme with network coding for LTE-A broadcasting system,” in *Proc. of IEEE Wireless Communications and Signal Processing*, Huangshan, China, Oct. 2012.
- [12] H. Hamdoun, P. Loskot, and T. O’Farrel, “Implementation trade-offs of fountain codes in LTE and LTE-A,” in *Proc. of Conference on Communications and Networking*, Kunming, China, Aug. 2012.
- [13] R. Mahajan, J. Padhye, S. Agarwal, and B. Zill, “High performance vehicular connectivity with opportunistic erasure coding,” in *Proc. of USENIX ATC*, Bostn, USA, Jun. 2012.
- [14] Y. Cui, L. Wang, X. Wang, H. Wang, and Y. Wang, “FMTCP: A fountain code-based multipath transmission control protocol,” *IEEE/ACM Transactions on Networking*, vol. 23, no. 2, pp. 465–478, Jan. 2014.
- [15] C. Paasch, “Improving multipath TCP,” Ph.D. dissertation, Ph.D. Thesis, UC Louvain, Nov. 2014.
- [16] A. Garcia-Saavedra, M. Karzand, and D. J. Leith. (2015) Low delay random linear coding and scheduling over multiple interfaces. [Online]. Available: <http://arxiv.org/abs/1507.08499>
- [17] M. M. Buddhikot and L. Hartung, “Policy driven multi-band spectrum aggregation for ultra-broadband wireless networks,” in *Proc. of IEEE DySPAN*, Stockholm, Sweden, Sep. 2015.
- [18] A. Ford, C. Raiciu, M. Handley, and O. Bonaventure. (2013) RFC 6824: TCP extensions for multipath operation with multiple addresses. [Online]. Available: <https://tools.ietf.org/html/rfc6824>
- [19] S. Katti, H. Rahul, H. Wenjun, D. Katabi, M. Medard, and J. Crowcroft, “XORs in the air: Practical wireless network coding,” *IEEE/ACM Transactions on Networking*, vol. 16, no. 3, pp. 497–510, May 2008.
- [20] M. Karzand, D. J. Leith, J. Cloud, and M. Medard. (2015) Low delay random linear coding over a stream. [Online]. Available: <http://arxiv.org/abs/1509.00167>