

# We are IntechOpen, the world's leading publisher of Open Access books Built by scientists, for scientists

5,300

Open access books available

130,000

International authors and editors

155M

Downloads

Our authors are among the

154

Countries delivered to

TOP 1%

most cited scientists

12.2%

Contributors from top 500 universities



WEB OF SCIENCE™

Selection of our books indexed in the Book Citation Index  
in Web of Science™ Core Collection (BKCI)

Interested in publishing with us?  
Contact [book.department@intechopen.com](mailto:book.department@intechopen.com)

Numbers displayed above are based on latest data collected.  
For more information visit [www.intechopen.com](http://www.intechopen.com)



---

# Transport Protocol Performance and Impact on QoS while on the Move in Current and Future Low Latency Deployments

---

Eneko Atxutegi, Jose Oscar Fajardo and Fidel Liberal

Additional information is available at the end of the chapter

<http://dx.doi.org/10.5772/intechopen.71779>

---

## Abstract

Transport protocols and mobile networks have evolved independently leading to a lack of adaptability and quality of service (QoS) degradation while running under the variability circumstances present in cellular access. This chapter evaluates the performance of state-of-the-art transmission control protocol (TCP) implementations in challenging mobility scenarios under 4G latencies and low delays that model the proximity service provisioning of forthcoming 5G networks. The evaluation is focused on selecting the most appropriate TCP flavor for each scenario taking into account two metrics: (1) the goodput-based performance and (2) a balanced performance metric that includes parameters based on goodput, delay and retransmitted packets. The results show that mobility scenarios under 4G latencies require more aggressive TCP solutions in order to overcome the high variability in comparison with low latency conditions. Bottleneck Bandwidth and Round-Trip Time-RTT (BBR) provides better scalability than others and Illinois is more capable of sustaining the goodput with big variability between consecutive samples. Besides, CUBIC performs better in lower available capacity scenarios and regarding the balanced metric. In reduced end-to-end latencies, the most suitable congestion control algorithms (CCAs) to maximize the goodput are NewReno (low available capacity) and CUBIC (high available capacity) when moving with continuous capacity increases. Additionally, BBR shows a balanced and controlled behavior in most of the scenarios.

**Keywords:** transport protocols, performance, mobility, 4G, low latency, 5G

---

## 1. Introduction

Mobile broadband (MBB) usage has risen significantly in the last years and so has done the customers' awareness regarding the quality of service (QoS). Thus, the measurement of QoS

---

in MBB networks has become a key issue. In this regard, the proper performance of the transport layer constitutes a critical feature in order to fulfill the QoS requirements of the clients (user equipments—UE) [1, 2]. In MBB, due to the multiple sources of variability (related to the client or self-inflicted effects such as channel quality reporting, propagation and fading pattern alterations due to mobility—and related to the intrinsic features of MBB such as bandwidth sharing, modulation and so forth), the network conditions become more volatile than in fixed networks and therefore the accuracy of transport protocols to adapt their sending rate as close to the available capacity as possible is reduced, impacting on the final performance [1, 2]. Therefore, there is an urgent need to study the relation between the transport protocol performance in MBB and its impact on the actual QoS results.

We consider that the mobility is one of the biggest features that differentiates cellular access from other connectivity schemes. Besides, the wide range of moving possibilities makes every use-case distinct and independent, creating different network conditions for transport protocols. In this regard, it is important to focus the performance-based analysis of current state-of-the-art transport protocol solutions over different mobile network mobility circumstances and understand the implications of the movement in the interaction between the transport protocol and the MBB variability.

Regarding mobile networks, this work analyses different schemes in order to give a wider view of the impact that the performance of transport protocols have in the QoS: 4G scenarios and low latency scenarios targeting assisted by mobile edge computing (MEC-assisted) future 5G deployments [3]. Future 5G networks aim at allowing improved capabilities in terms of achievable capacities, modulation and end-to-end latency among other features. The reduction in the transmission latency is one of the main beneficial evolutions for the suitable performance of transport protocols. It shortens the feedback time between consecutive management decisions in the server, increasing the responsiveness of the transport layer to the fluctuations of the radio side. Therefore, we focus our second scheme in the proximity service of 5G deployments. To that end, we mimic with low latency 4G scenarios a 5G-alike service provisioning. All in all, our evaluation covers the performance of transport protocols in current and future MBB.

Considering that TCP is the predominant transport protocol on the Internet, we focus our study in the evaluation of TCP over distinct mobility circumstances over 4G latencies and low latency deployments. TCP is not a single entity but a family of different congestion control algorithms (CCAs) that manage the outstanding data of the server (clamped by the congestion window—CWND) in a different way based on pre-defined features such as throughput maximize algorithms with loss-based mechanisms, delay-aware implementations or hybrid developments. So far, despite many CCAs being available, none of them have demonstrated to both be easily deployable and appropriately face the variability of MBB fluctuation. This work selects and evaluates the performance of distinct CCAs that count on different features and implementations that in the end result in a different performance outcome in each precise network conditions.

The great success of TCP and user datagram protocol (UDP) have led to the widespread utilization of both of them, either as the selected transport solution or as a substrate to enable

the so-called transport services [4–6]. These transport services are ad-hoc layers that work between the transport layer and the application layer, taking advantage of the substrate transport protocol (mainly TCP and UDP) and gaining some freedom due to its development in the user space of the operating system (OS) and additional functionalities (i.e. congestion control of QUIC over UDP). However, the utilization of TCP and UDP forces the system to stick with the infrastructural characteristics of the selected substrate transport protocol from the beginning of the transmission until it is closed. This limitation has been named “ossification” and it has three main effects: (1) it closes the opportunity to select and modify transport layer protocols at the beginning of a certain transmissions; (2) it leaves little room for transport protocol innovation; (3) it provides limited or non-existent flexibility of the application programming interface (API) [5]. This API serves as the connection point between the application layer and the transport layer. The existence of constraints in the communication between these two layers is directly translated to a standard behavior of the transport protocol with no consideration of the requirements from the application layer. In current MBB, the transport protocols misbehave due to the incapability of adapting to the actual network circumstances and the impossibility of adapting its features to the requirements from upper layers.

In a close future with the implementation of evolved transport services, the API would not only select the best transport protocol based on application requirements, but it would also consider the network status. Even though, this mechanism would require further signaling and interaction, recent advances [7] are evolving in this sense and could provide with a more complex and complete API. Taking into account that each CCA could be more suitable for certain network circumstances or application requirements than others, our work evaluates the best CCA candidate for each combination of conditions. In this regard, out of all possibilities, our analysis covers the study of TCP CUBIC, NewReno, Illinois, Westwood+ and the recently released BBR. Thus, being capable of providing hints in the complex process of improving the behavior in the transport layer and therefore in the resultant enhanced QoS.

The main findings of the chapter are the following. The mobility scenarios under 4G latencies require more aggressive TCP solutions to overcome the high variability in comparison with low latency conditions. Merely focusing on goodput: (1) although BBR provides with the best scalability, it also induces greater mean delays and lost packets; (2) in scenarios that evolve with continuous capacity reductions, Illinois shows the best adaptability to the variable conditions when the achievable capacities are high, whereas CUBIC demonstrates the same for lower bandwidth assignments. Considering the performance with a combination of goodput, delay and retransmissions, CUBIC presents a more balanced behavior with average achieved rates but greater awareness of the self-inflicted delay and retransmissions. With reduced end-to-end latencies, the most suitable CCAs are: (1) NewReno for low available capacity circumstances that moves with continuous capacity increases; (2) BBR as the most balanced CCA that allows both high bandwidth achievement and low delay and retransmissions; (3) CUBIC when scalability is required in presence of big changes between consecutive samples of assigned radio capacity and (4) similar goodput-based performance of all CCAs while moving to worse quality positions.

The chapter is structured as follows: Section 2 introduces the related work in the analysis of TCP in MBB in general and with mobility circumstances in particular. Section 3 shortly describes the utilized CCAs in the analysis. Section 4 covers the methodology with the description of the testbed and the utilized measurement and evaluation process. The analysis and results are explained in Section 5 divided by the 4G latency schemes and the low latency scenarios that mimic 5G deployments. Finally, Section 6 gathers the most important conclusions and proposes future lines.

## 2. Related work

When analyzing how the mobile networks' features have an effect on transport protocol behavior and therefore impact the final QoS, there are several characteristics that have to be mentioned.

### 2.1. Delay

It is clear that comparing mobile networks and fixed networks, the former has more variable channel conditions that could lead to achieve a degraded throughput [1, 2]. However, there are effects that from a macroscopic point of view are shared among the distinct networks. For instance, it has been proved [8, 9] that even mobile networks suffer due to the excessive buffering in intermediate queues leading to an increase in the end-to-end delays and dropped packages that severely impact the performance of TCP (bufferbloat effect). Measurements over both 3G and 4G cellular networks of four U.S. providers and Swedish networks have concluded that bufferbloat represent a problem in MBB too. Our study not only considers the achieved capacities but also the induced delay due to the possible impact that may well have as a cross-traffic.

### 2.2. Impact of variability

It has been demonstrated that there are differences among distinct mobile networks. A comparative work of 3.5G and 4G [10] showed that 4G networks are worse in regards to the TCP efficiency due to the superior throughput and variability. This is the higher variability, the worse scenario for TCP due to the lack of rapid adaptability. Garcia et al. [11] carried out measurements in the cellular networks of different Swedish operators so as to analyze the variation of TCP throughput and delay throughout the day. Sudden increases in traffic load leads to bandwidth variability and latency increment [2]. Therefore, TCP happens to drastically reduce its throughput. Additionally, TCP experiences timeouts many times. The timeout events are especially harmful because the CWND is reduced to one segment. In another study, Alfredsson et al. [12] proved that the variable modulation on the 4G link layer is contributing to retransmissions' increment and therefore higher delay and less throughput. Huang et al. [13] carried out a comprehensive study related to TCP throughput and latency estimation over a live long term evolution (LTE) network. In their measurements, they found out similar timeout events. Our work precisely evaluates how different variable mobility circumstances

affect the performance of different CCAs in order to better understand the implications of fluctuations in the channel quality.

### **2.3. Uplink impact and bursty behavior**

Regarding the evaluation of TCP throughput, it has been proved [14] that the uplink performance tends to degrade its performance due to scheduling policies, severely impacting on Acknowledge packet (ACK) arrival and therefore downlink injection ability. Those TCP flavors that merely depend its CWND management upon the reception of ACKs are drastically affected by these ACKs reception, also called ACK-compression. Related to this issue, it has been demonstrated [15] that modern cellular networks' traffic has a tendency to become bursty. For this reason, there can be a large variation in the actual throughput during a short period of time (varying by up to two orders of magnitude within a 10-min interval). This variability could be even harder due to the fact that mobile providers often maintain a large and individual downlink buffer for each UE, provoking high latency instability. Our analysis aims at detecting the most suitable CCA for different MBB mobility circumstances that also implicate distinct bursty conditions over the network.

### **2.4. Impact of speed**

If mobile networks themselves suffer high variability, the channel conditions could be even more variable and challenging for TCP due to the movement of UEs. The movement leads to have distinct propagations and fading patterns over time that at the same time impact on the assigned modulation to the UE, provoking "jumps" between consecutive channel quality reports. Merz et al. [16] studied the performance of TCP in live LTE networks in mobility scenarios with speeds up to 200 km/h. They mainly evaluated the spectral efficiency depending on the modulation and the bandwidth share among the attached users to the eNodeB, together with the ability of those users to make the most of the assigned capacities. Li et al. [17] compared the performance of TCP in static positions with moving scenarios, resulting in harmful RTT spikes, massive dropped packets and eventual disconnections while on the move. Our work complements the mentioned studies by adding the evaluation of multiple state-of-the-art CCAs as well as the inclusion of different MBB schemes with 4G latencies and low latencies.

Taking into account the research and standardization momentum regarding 5G in which most of the work is yet to be fulfilled, there are few works that have considered the performance of TCP in the future 5G MBB. Pedersen et al. [18] demonstrated the potential of using different transmission-time intervals (TTI) in the eNodeB of 5G deployments depending on the metadata related to a certain channel. They showed that shorter TTIs were capable of allowing higher throughputs for short communications, whereas longer TTIs could overall benefit the performance of large transmissions. Sarret et al. [19] study the forthcoming benefit of using full duplex at the radio link layer (RLC) in comparison with the current half-duplex implementation for an improved throughput and delay. Besides, they covered the possible configurations in ultra-dense 5G deployments that could limit the envisioned rates.

Even though several research studies and proposals have reported their concerns and findings regarding the effects between mobile networks and TCP under challenging conditions, none of them have considered the analysis of distinct state-of-the-art CCAs of TCP under different mobility patterns and circumstances over 4G latencies and low latencies targeting proximity MEC-assisted provisioning in 5G.

### 3. Selected CCAs

The studied TCP variants fall into five categories with regard to their employed CCAs: loss-based, combined loss- and delay-based (with or without bandwidth estimation), and delay-based. As examples of loss-based CCAs, we study both NewReno [20] and CUBIC [21]. NewReno was selected due to its prevalence in research and its large implementation base, and CUBIC by the fact that it is the default CCA in Linux. The Westwood+ [22] congestion control was taken as a CCA example of a combined loss- and delay-based with bandwidth estimation technique. In many ways, TCP Westwood and its successor TCP Westwood+ laid the foundation for the work on designing a CCA that is able to distinguish between congestion and non-congestion related packet losses in wireless networks without any support from the wireless MAC layer. Also, Illinois [23] was selected as an example of a combined loss- and delay-based CCA. In contrast to Westwood+, Illinois primarily targets high-speed and long-delay networks. Finally, TCP BBR (Bottleneck Bandwidth and RTT) [24, 25] constitutes a model-based CCA that drives the congestion avoidance management based on two parameters: measured baseline RTT (delay-based) and the timing and rate of ACK packets (bandwidth estimation). A brief overview of the five studied TCP variants is given below.

**TCP NewReno** [20] is the basic TCP implementation that drives its CWND based on the additive increase and multiplicative decrease (AIMD) principle. In this regard, NewReno increases its CWND by one packet for each ACK reception during the Slow Start phase. Instead, during the congestion avoidance phase, it increases the CWND by one segment for each RTT. The increment is performed until a timeout period is consumed or a notification of a loss packet is received (with a triple duplicate ACK—3DUPACK). Depending on the event, NewReno would back-off differently, halving the CWND in case of 3DUPACK and establishing the CWND in one segment when a timeout is detected.

**TCP CUBIC** [21] uses a cubic equation during the congestion avoidance phase to manage the CWND. The closer the CWND is to the previous congestion point in terms of outstanding packets, the slower increment is applied. This function leads to a zero increment while the previous congestion point is achieved. If CUBIC does not detect congestion at that point, it increases the ramp-up pace of the CWND with a convex shape until a new loss event happens. One of the main features of CUBIC in comparison with NewReno is that the CWND is not ACK-clocked and therefore, depend less significantly in the RTT. CUBIC also introduces a modification in the Slow Start phase in order to avoid massive packet

losses at the end of the ramp-up. The modification is called Hybrid Slow Start and tries to transfer the management of the CWND to the congestion avoidance phase prior to overfeed the network. To that end, two exit conditions are added: (1) if a delay increase over a pre-defined threshold is detected and (2) if the ACK train is lengthen. If any of the conditions is met, the Slow Start phase is left and the congestion avoidance phase would follow driving the CWND.

**TCP Westwood+** [22] is a sender-side modification of TCP to allow estimating the available bandwidth by assessing the incoming ACK packets. The measured capacity serves to adjust the CWND during back-off phases after a loss event occurs. The selected CWND tries to establish the potential outstanding packets as close to the maximum capacity as possible but avoid building-up the queue of the bottleneck. To that end, the bandwidth-delay product (BDP) is calculated with the estimated bandwidth and the minimum assessed RTT.

**TCP Illinois** [23] is a loss-based AIMD mechanism that drives the CWND with certain knowledge of the queuing delay and buffer size of the bottleneck. This delay awareness is taken from the RTT measurements and consequently it is updated upon ACK arrival. If no excessive queuing delay is detected, the CWND would increase faster than in conditions of high induced latency. The maximum increment is established in 10 segments per RTT, while the minimum is set to 0.3. When the RTT is close to the maximum, the loss is considered as buffer overflow, whereas in low RTT the loss counts as packet corruption.

**TCP BBR** [24, 25] is the recently developed TCP implementation that bases its CWND management in a model of the bottleneck's BDP. It considers the estimated bottleneck bandwidth and the measured RTT in every update of the model. The estimated bottleneck bandwidth is measured by calculating the timing and rate of receiving ACKs in the sender. The calculated model determines whether the packet injection rate is below or over the capacity of the bottleneck, being able to appropriately adjust to the network requirements. Such an adjustment of the injection rate is carried out following the principle of  *pacing*, either by using Fair Queue packet scheduling or the native and fall-back implementation of  *pacing* developed in the transport layer. Besides the main behavioral features, BBR is handled with a four stages workflow:

- In the Startup stage BBR ramps-up as the Standard Slow Start until it detects that the obtained throughput gain is below the 25% throughout three consecutive RTTs.
- In the draining stage BBR tries to get rid of all the excessive packets in the bottleneck queue.
- In the probing bandwidth stage BBR uses an eight state cycle to cruise at different pacing rates. Throughout six states, BBR injects at the measured bottleneck's BDP rate if no change of the available capacity is detected. The other two states are a bandwidth probing phase with a 25% of the injection rate increment and a draining phase with the ability to drain the excess packets injected in the previous phase if the bottleneck does not tolerate greater throughputs.



- In the probing RTT stage, BBR re-measures the baseline RTT for a proper modeling of the network path. To this end, BBR reduces the CWND to four segments for at least 200 ms and then, re-established the previous CWND.

## 4. Methodology

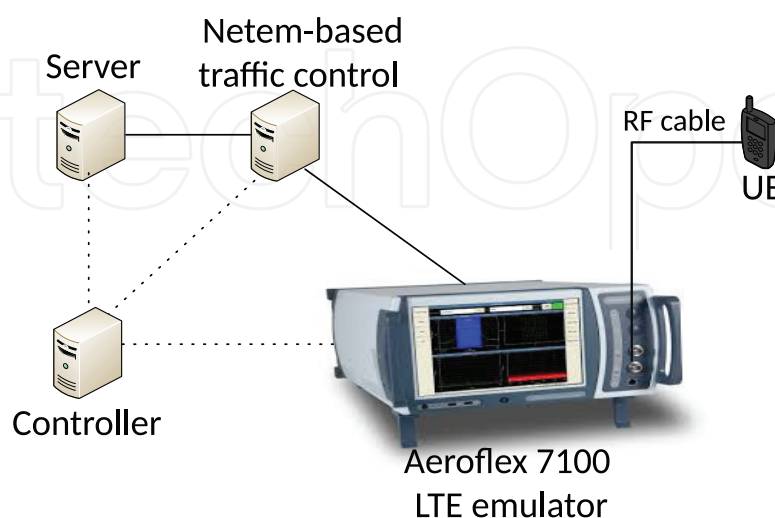
This section presents the measurement testbed with the equipment involved in the assessment process and describes the measurement and evaluation procedure that have led to the results presented in the following analytical section.

### 4.1. Testbed

The LTE deployment uses a digital radio testing emulator or a LTE-in-a-box as the main equipment responsible for the LTE side. The main radio configuration parameters are the utilization of the seventh band of LTE due to its widespread usage and the availability of 100 physical resource blocks (PRB) and 20 MHz channels in order to be capable of performing at the full potential of the cell. This emulator plays among other attributes (full EUTRAN/EPC testbed) the role of the eNodeB, creating the LTE signaling to support the attachment and registration of any LTE device through a radiofrequency (RF) cable.

**Figure 1** shows the experimental testbed and how the LTE-in-a-box is placed and connected to other equipment in the deployment. Apart from the LTE emulator, the testbed is formed by different parts:

A LTE UE is included with the capability of connecting to the network through the RF cable. Such connection is directly done to avoid undesired effects of the environment in the transmission.



**Figure 1.** LTE testbed with Aeroflex 7100 LTE-in-a-box.

A Linux server is placed with a 4.10 kernel that contains the recent advances in the transport layer as well as TCP BBR. Besides, the server contains the required files that are used during the experimental phase. The files are devoted to ensure long experiments or greedy sources so as to appropriately detect differences in the behavior and performance of distinct CCAs. In addition, the server is responsible for gathering the logging files created by the *ss* (socket information) and *tcpdump* (whole pcap file of each transmission) programs that later on will be used for the analysis.

One traffic-control bottleneck server is located between the end-server and the LTE-in-a-box to manage the end-to-end delay with *netem*.

A controller is used to automatize and synchronize the rest of the equipment throughout the experiments. The controller is also responsible for commanding the precise configuration that each part of the testbed should apply. For instance, the controller selects which CCA is used in the end-server, the delay of the bottleneck, the baseline signal-to-noise-plus-interference (SINR) of the channel between the LTE-in-a-box and the UE (either as a fixed value or with variations that could allow emulated mobility) or the fading pattern to be utilized in the mentioned channel.

The emulated testbed enables representative reporting and signaling with real UEs, gives the opportunity to configure the LTE-in-a-box including fading patterns and it is capable of collecting logging traces. All this features make the selected testbed realist enough to be a representation of certain real-world network circumstances with the additional control and parameterization of the measurements outcome that a close testbed provides.

#### 4.2. Measurement and evaluation procedure

In order to experiment with 4G latency MBB scenarios as well as low latency schemes that model the proximity service of MEC-assisted 5G future deployments, we have used different delays in the network path. Due to the privacy and non-disclosure information of operators, there is little data regarding the latencies present in 4G and the ones expected in real-world deployments for 5G. Therefore, our study is based on a report [26] that shows the delay results of four operators being between 68 and 85 ms on average for 4G. Thus, we configure our 4G latency scenarios with a minimum latency of 68 ms and the low latency scheme with the lowest possible value in our testbed, 18 ms.

The emulated effect of movement is obtained by the application of two parameters: (1) a selected fading pattern in the LTE-in-a-box that would affect the channel between the eNodeB and the UE and (2) the external SINR traces that periodically command the baseline SINR. These messages are sent by the controller and applied by the emulator in order to modify the baseline SINR of the channel towards the UE.

Regarding fading, since different maximum achievable rates and variability lead to different challenges for TCP, we have applied two distinct fading patterns in regards to modeling common fading effect under mobility in real deployments. The patterns are the following:

- A mobility scenario with Extended Vehicular A model 60 (EVA60) fading model. This fading and variability tries to mimic the vehicular scenarios at 60 km/h, which is a common limitation in rural roads.
- A mobility scenario with High Speed Train (HST) 300 (HST300) fading pattern. The selected fading models the signal fluctuation of current highspeed trains at 300 km/h.

**Figure 2** shows the distribution of CQIs while applying the combination of a baseline SINR of 20 dB and the selected fadings. It is clear that both patterns depict a variable behavior even under static baseline SINR circumstances. However, both fading patterns result in a completely different environment for TCP. EVA60 demonstrates a variable behavior in high CQI values (high capacities), whereas HST300 shows even more fluctuation in lower CQI values (lower achievable rates).

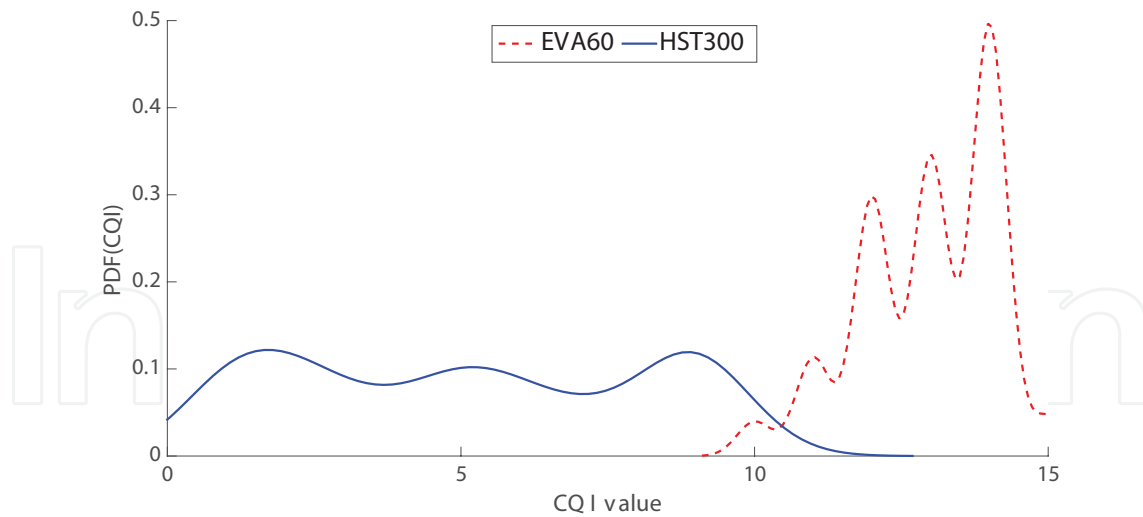
Apart from the variability that the fading traces create in the radio link, in order to run under different mobility patterns, the Controller in the testbed would periodically command the baseline SINR of the channel to the emulator. Our experiments evaluate the responsiveness and suitability of the selected CCAs in two simplified mobility conditions:

- Forward movement or out-cell: this mobility pattern evaluates the performance of different flavors of TCP with a UE moving from the eNodeB towards a constantly worsened radio channel conditions. Therefore, the applied mobility traces would start from the baseline SINR of 20 dB and would go progressively worsening until 5 dBs are reached. The speed of such transition is determined by the scenario under study and thus, the 60 km/h mobility scenario spends five times the time is needed to perform the same at 300 km/h. The main idea is to evaluate TCP under variable conditions that progress with an average continuous capacity reduction.
- Backward movement or in-cell: this mobility pattern evaluates the performance of the selected CCAs with a UE moving in the other way around. The mobility traces would start by applying a baseline SINR of 5 dBs to the channel and the quality would progress increasing in accordance to the selected scenario and its modeled speed. The idea is to assess the performance of TCP in variable channel conditions when the quality of the channel tends to continuously increase its available capacity.

We understand that the simplification of movement into backward movement and forward movement does not directly represent the mobility circumstances in the real-world. However, we stand that many movement patterns can be divided or split into the aforementioned cases.

In regards to the utilized traffic, greedy sources are employed with no cross-traffic. The decision of performing with isolated flows is based on the better understanding of the impact that different mobility patterns and fading traces have, together with the better detection of the CCAs' adaptability.

In order to analyze NewReno, CUBIC, Westwood+, Illinois and BBR in the selected mobility scenarios under 4G latencies and low latencies that model the potential MEC-assisted service



**Figure 2.** CQI distribution while combining a baseline SINR of 20 dBs and the selected fading patterns.

provisioning delay in 5G, we consider different metrics in the evaluation process. Instead of merely measuring the performance only based on the goodput performance, two different approaches are followed:

1. Criterion 1—pure goodput performance in order to assess the adaptability of CCAs and which are the rates capable of achieving.
2. Criterion 2—a performance metric that includes not only the goodput samples but the delay as well as the retransmission events. This way, a single value is able to include several important fields of the performance of a certain CCA. Thus, the evaluation could consider a more balanced parameter that does not rely on a single performance side, avoiding the cyclic dependency of TCP (i.e. in many situations the achievement of greater goodput—positive fact, also means the injection of greater delay—negative impact). In order to build a metric that considers all abovementioned parameters, we decided to create the following one:

$$A = \frac{K}{Kt} \tag{1}$$

$$B = \frac{Dmin}{D} \tag{2}$$

$$C = \frac{BDP - MSS * R}{BDP} \tag{3}$$

$$\alpha = 1; \beta = 1; \gamma = 1 \tag{4}$$

$$Pm = \frac{(\alpha * A) + (\beta * B) + (\gamma * C)}{\alpha + \beta + \gamma} \quad (5)$$

The first parameter (Eq. (1)) measures how much out of the available capacity ( $Kt$ ) is reached in a precise sample ( $K$ ). The second one (Eq. (2)) indicates the growth of current delay ( $D$ ) considering the baseline delay of the transmission ( $D_{min}$ ). The third one (Eq. (3)) takes into account out of the current  $BDP$  ( $BDP$ ), how many bytes are wasted in retransmissions (number of retransmissions— $R$  times the maximum segment size— $MSS$ ).  $\alpha$ ,  $\beta$  and  $\gamma$  (Eq. (4)) are parameters to weight the importance of the three principal. These parameters could be configured for the precise requirements of a certain application (critical with losses or critical with delays among other options). In our evaluation (Eq. (5)), our performance metric ( $Pm$ ) considers every main parameter equally important to get a balanced performance.

The evaluation process aims at deciding the most appropriate options among selected CCAs depending on the selected mobility use-case and application requirement (evaluation criterion). The analysis and evaluation is linked with the current and state-of-the-art possibilities that are being opened in the transport layer so as to select the most suitable protocol at the beginning of every transmission [7]. Our work contributes to the better selection of CCAs under 4G latency and low latency MBB mobility scenarios.

## 5. Analysis

This section covers the overall performance and most suitable selection of TCP under mobility in 4G latencies and under low latency scenarios. In order to better explain the analysis, the section is divided into two subsections: (1) Section 5.1 covers the analysis regarding the performance of TCP under 4G latencies in mobility scenarios, together with the evaluation and selection of most TCP candidate considering goodput requirements and performance metric requirements; (2) Section 5.2 performs the same analysis and evaluation in low latency mobility scenarios. Each subsection covers three main steps in the explanation, it presents the overall performance, the goodput-based evaluation and the assessment based on the performance metric.

Since throughout the analysis we will use a performance metric that also considers the injected delay and number of retransmissions, the representation of the overall performance will take into account three parameters: mean goodput, mean delay and mean retransmissions per time slots of 100 ms. In order to avoid complicated graphs that may well be misinterpreted, the overall performance is depicted in spider plots of three axis, one per performance parameter. Each figure comprises four subplots: the first two subplots on the top depict the results for forward movement pattern, whereas, the bottom line shows another two subplots for backward movement pattern. The two subplots per line are representative of the different speed and fading conditions: from the left to the right, the scenarios at 60 and 300 km/h.

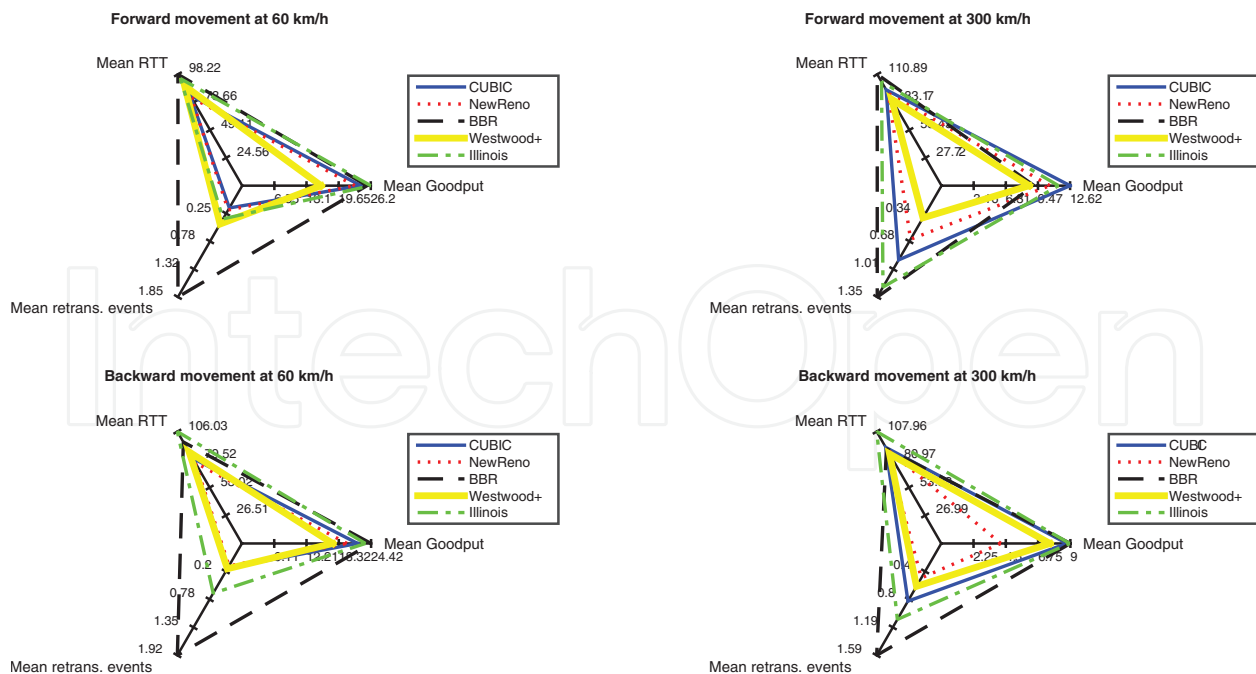
Supported by the spider plot of the overall performance and adding a table that gathers the mean values and the average confidence intervals of goodput in each mobility scenario, the most appropriate CCA is selected. In this sense, the selection may well be utilized by requirement of the applications to pick among all CCAs the option that best maximizes the achieved rates in a precise mobility scheme.

Once selected the best candidates for mobility scenarios merely based on goodput, it is important to carry out a similar task but evaluating the performance of distinct TCP flavor from a point of view that gathers more information about the outcome and possible side-effects. To that end, a figure will depict the empirical cumulative distribution function (ECDF) results of the performance metric in all scenarios. The distribution of subplots is the same as before with forward movement scenarios in the first row and backward movement ones in the bottom line. Each column represents a different speed and fading combination, being from the left to the right the scenarios at 60 and 300 km/h, respectively.

### 5.1. TCP performance with mobility under 4G latencies

This section covers the analysis under 4G latencies of the selected movement patterns of five CCAs. **Figure 3** shows the spider plot results in mobility scenarios for CUBIC (straight line), NewReno (dotted line), Westwood+ (straight thick line), Illinois (dash-dotted line) and BBR (dashed line).

Out of all the results in **Figure 3**, the most important ones are the following:



**Figure 3.** Performance spider plot of five CCAs while moving at different speeds under 4G latencies: forward movement on top; backward movement on bottom line.

- In forward movement pattern: We find two different cases: (1) when the variability happens at high capacities (scenario at 60 km/h) since the movement itself regardless the CCA under use allows having almost throughout the whole experiments packets in-flight, the mean achieved goodput is very similar in most of CCA cases. The best candidates are Illinois and BBR in terms of goodput but they also induce more delay and retransmissions (more significant with BBR) than NewReno or CUBIC; (2) at 300 km/h where the mean fluctuation is harder due to the superior speed, the self-inflicted effects of delay and retransmissions severely impact the goodput performance, dropping the goodput performance of BBR and Illinois and being clearly surpassed by CUBIC; (3) as detected in Ref. [27] Westwood+ suffers due to the excessive reduction in the back-off application after Slow Start, leading to an underuse of the radio capacity.
- In backward movement pattern: (1) both Illinois and BBR has shown better scalability than CUBIC in all scenarios achieving greater goodput rates, but as a drawback, inducing more delay in the network and suffering more retransmissions; (2) the scenario itself due to its low available rates at the beginning of the transmission helps appropriately perform weak CCAs such as Westwood+. In this sense, the deficiency of Westwood+ is minimized, leading to better results in terms of goodput while the delay is kept low; (3) considering goodput, apart from NewReno that demonstrates to underperform in variability circumstances when scalability is required under 4G latencies, the rest of the candidates achieve similar (not equal) results. The scenario at 60 km/h shows bigger differences among the CCA candidates due to the greater capacities present in the scenario, allowing more aggressive CCAs adapt better to the available bandwidth.

We have detected cases in which a similar goodput is achieved but significantly more delay and retransmissions are suffered. These examples, that suppose a difficult performance trade-off to analyze, are the foundation for the evaluation of the protocols based on different points of view in the performance in order to appropriately select the best candidate for each network circumstances but also considering the application requirements.

The overall performance has depicted the goodput performance as one of the parameters in the spider plot. Now instead, **Table 1** covers the performance of the goodput showing the actual average values of the transmission throughout each mobility scenario, together with the average confidence interval as a representation of the differences between independent tests.

**Table 1** shows that depending on the scenario a proper CCA selection could allow the achievement of greater capacities. The best practises regarding the goodput-based evaluation are the following under 4G latencies:

- In backward movement pattern (right side of the table), BBR shows great variability in the behavior. Even with that, it is able to scale better than any other candidate. This ability could be of a great value in other MBB scenarios in which scalability is required, either as a necessary feature or as a response to available bandwidth increments.

Context	Forward movement		Backward movement	
	Mean	Mean CI	Mean	Mean CI
<b>60 km/h with EVA60 under 4G latencies</b>				
CUBIC	24.18	6.59	21.87	6.89
NewReno	22.95	7.95	19.87	7.16
BBR	25.43	18.7	24.42	17.54
Westwood+	16.08	8.01	17.16	7.68
Illinois	26.2	10.7	23.35	14.84
<b>300 km/h with HST300 under 4G latencies</b>				
CUBIC	12.62	7.81	8.76	5.51
NewReno	10.63	7.58	4.15	2.89
BBR	9.09	10.21	9	6.76
Westwood+	8.61	4.13	7.58	4.28
Illinois	11.48	6.8	8.95	5.84

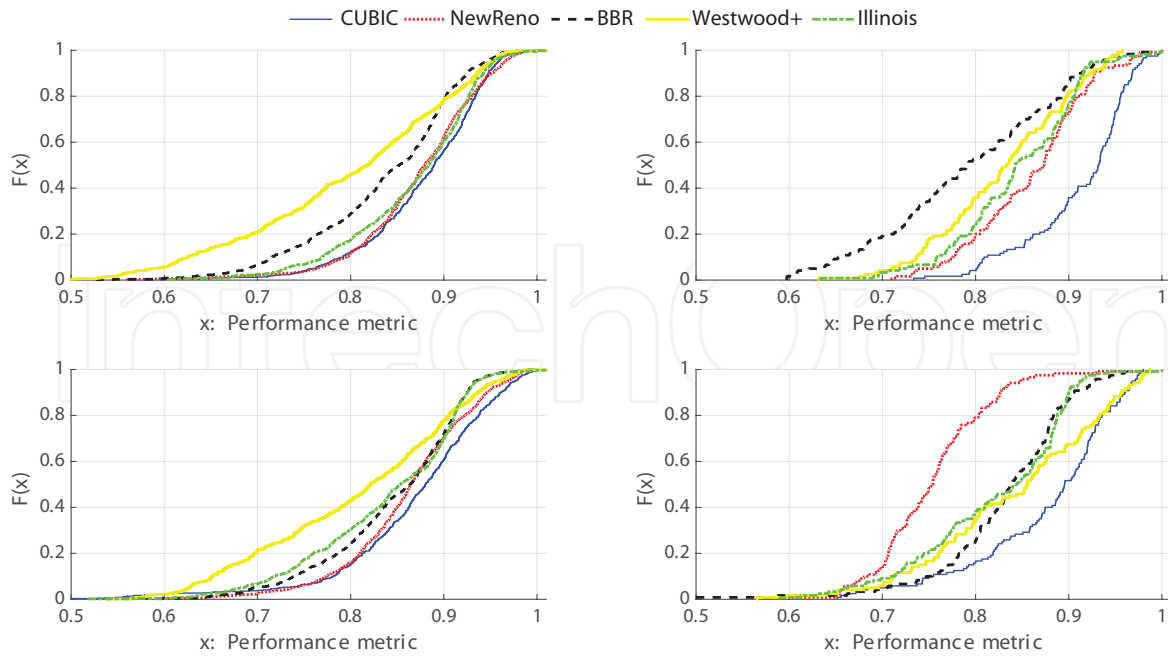
**Table 1.** Goodput-based evaluation of the selected CCAs in mobility scenarios under 4G latencies.

- In forward movement pattern (left side of the table), the two different scenarios require a distinct treatment and therefore CCA candidate. The best practises in forward movement pattern under 4G latencies are the following:
  1. Illinois is selected in scenarios with big bandwidth jumps between consecutive RTTs due to its aggressiveness to handle such variability but also due to its delay awareness. The combination of features provides the best and most consistent (less variation in the behavior than the second candidate according to the results—BBR) goodput performance.
  2. CUBIC is picked in variable scenarios with smaller changes between consecutive available bandwidth samples. The aggressive features of CUBIC are suitable to handle great fluctuation in smaller capacities but do not make the most when it comes to bigger bandwidth jumps.

**Table 1** reveals that in forward movement, an appropriate CCA selection could provide with a great gain in terms of goodput even using a movement pattern that makes easier the task of TCP due to the avoidance of starvation events in the bottleneck buffer.

Once selected the best candidates for mobility scenarios under 4G latencies merely based on goodput, it is important to carry out a similar task but evaluating the performance of distinct TCP flavor from a point of view that gathers more information about the TCP performance. To that end, **Figure 4** depicts the ECDF results of the performance metric in all scenarios. The



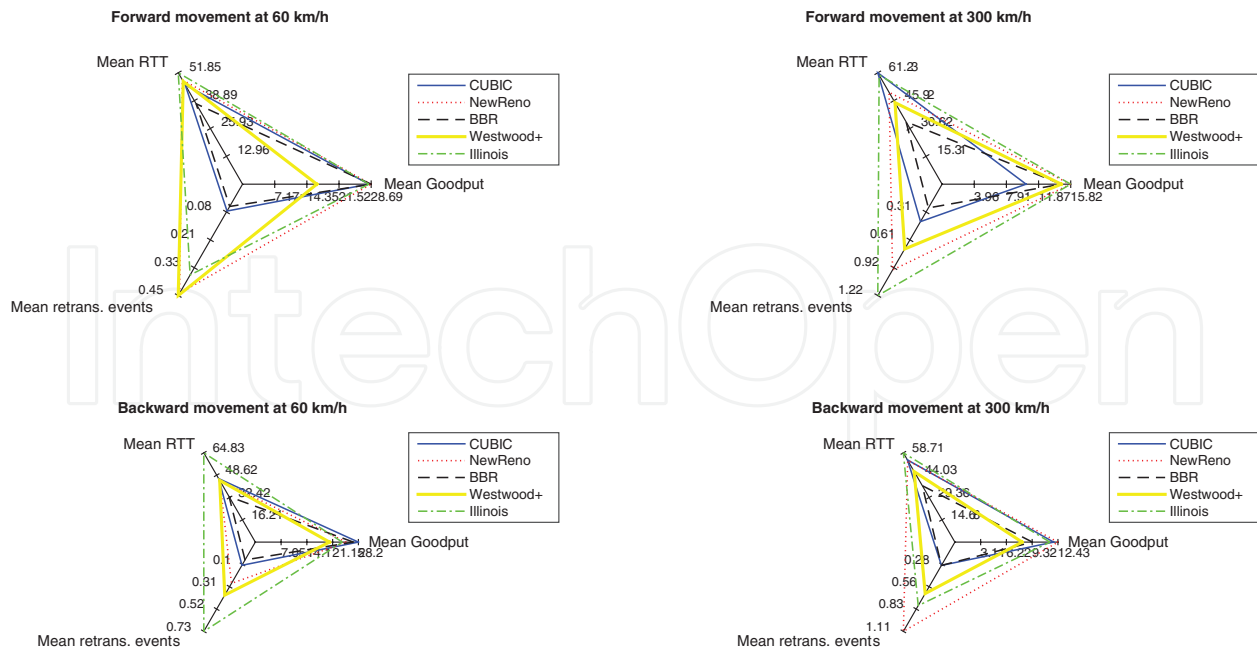


**Figure 4.** Comparison of five CCAs performance metric while movement at different speeds under realistic 4G latencies: forward movement on top; backward movement on bottom line.

distribution of subplots is the same as before with forward movement scenarios in the first row and backward movement ones in the bottom line.

**Figure 4** contrasts with previous goodput-based evaluation in many ways. It shows that sometimes the TCP performance depends on the proportion between parameters, suffering a performance cycle that impacts negatively certain performance fields when others are improved and vice-versa. In general in forward movement, **Figure 4** demonstrates that the performance enhancements are not that clear in one CCA over the other. Maybe one CCA is able to perform better in terms of goodput but at the same time suffering from delay and retransmitted packets. Therefore, taking into account the three selected performance factors, it is not that easy to decide whether we should use one congestion control or the other. If the decision of the CCA is taken upon the evaluation of this performance metric, in the scenarios with similar outcome, the final CCA would be picked based on other circumstances. For instance, the decision is taken with: (a) developer preference; (b) the by-default CCA prevails over the others; (c) if at the point of selection there is one CCA already established and it is among the group with similar outcome, the CCA could remain selected. In contrast and still selecting the most appropriate CCA in forward movement pattern, the 300 km/h scenario is clearly dominated by CUBIC. CUBIC not only achieves a great goodput, but shows a more general and extensive good performance.

Regarding the results in backward movement scenarios, there is a clear pattern that reflects that CUBIC is more appropriate. The outcome contrasts the evaluation based on goodput that suggests the selection of BBR, proposing to pick CUBIC as a candidate that improves the overall performance based on the three selected performance aspects: goodput, delay and impact of retransmissions.



**Figure 5.** Performance spider plot of five CCAs while moving at different speeds under low latencies: forward movement on top; backward movement on bottom line.

## 5.2. Overall performance under mobility under low latency

Once explained the performance of the selected CCAs over the mobility scenarios under 4G latencies and having evaluated the best TCP candidate considering the goodput and the performance metric, this subsection covers the analysis under low latencies. **Figure 5** shows the spider plot results in mobility scenarios for CUBIC (straight line), NewReno (dotted line), Westwood+ (straight thick line), Illinois (dash-dotted line) and BBR (dashed line).

The most important details in **Figure 5** are the following:

- As a general performance and comparing with 4G latencies, Illinois still suffers due to an excessive delay and number of retransmissions, whereas BBR shows a more controlled behavior with a good performance in some scenarios in terms of goodput but also avoiding massive delay and retransmitted packets.
- In forward movement pattern: (1) there still exist huge similarities among CCA from the goodput's point of view. Under low latency circumstances, even the CCAs that in principle are less aggressive (i.e. NewReno in both scenarios or Westwood+ at 300 km/h), demonstrate to perform very similar to more aggressive candidates, thus achieving a great goodput performance but suffering from excessive delay and retransmissions; (2) moving at 300 km/h where the achievable capacities are low, aggressive solutions such as CUBIC obtain worse results than others due to their injection pattern. Even though the mean retransmission events are lower than Illinois or NewReno, the periodicity is different. In each "big" retransmission event of Illinois or NewReno, many retransmitted packets are involved, whereas in the case of CUBIC the "big" retransmission events are more but with less packets involved. Thus, CUBIC due to its aggressiveness, finds

more times “big” retransmission events (increasing the suffered mean RTT), reducing in more occasions the CWND and therefore, being less capable of achieving the available capacity.

- In backward movement pattern: (1) Westwood+ and Illinois are not able to perform better than CUBIC or NewReno in any case. The former primarily due to its poor ability to scale. The latter, due to its self-inflicted effects, suffering by far the greatest delay; (2) in the case of BBR, as mentioned earlier, the behavior under low latency looks more conservative. In this regard, in backward movement pattern that requires scalability, in terms of goodput BBR performs worse than CUBIC in the mobility scenario at 60 km/h and underperforms in the scenario at 300 km/h earning very poor performance results.

Once again and even more noticeable than with 4G latencies, we have detected cases in which a similar goodput is achieved but significantly more delay and retransmissions are suffered by different CCAs. Those examples require an evaluation of the protocols based on different application-layer requirements so as to appropriately select the best candidate whose performance matches the application requirements and network conditions.

**Table 2** only covers the representation of the goodput with average values of the transmission for each mobility scenario and the average confidence interval for different independent tests.

**Table 2** demonstrates that depending on the scenario, the selection of the most appropriate CCA is capable of allowing better performance, achieving greater capacities. All in all, under low latency conditions, the most suitable CCA regarding the goodput-based evaluation are the following:

- In forward movement, CUBIC is damaging in scenarios with drastic fluctuation in low available bandwidths comparing with other candidates and its aggressiveness looks as if does not properly fit for such scenarios. Due to the low end-to-end latency and taking into account that forward movement is an easier movement pattern to handle by TCP (because it allows having almost every time packets in-flight), we can barely decide one specific CCA in each scenario. The selection should be carried out among a group of CCA that have reported a very similar outcome in terms of goodput.
  1. NewReno and Illinois have shown a great performance in variable scenarios with low available capacities (mobility scenario at 300 km/h).
  2. Under variable conditions with big capacities and big jumps between resource assignments, all CCAs but Westwood+ have proven to achieve very similar performance. Comparing with 4G latency scenarios, these similarities clearly come due to the reduction in the end-to-end delay.
  3. In backward movement, even with similarities among the CCAs, it is still clear that one or two candidates prevail over the others. Variable scenarios are better handled by CUBIC when the capacities are higher due to its superior aggressiveness, whereas

NewReno improves the performance of variable scenarios with low available bandwidths. These results demonstrate that, yet, there is value in the performance of by-default CCA (CUBIC) and classic CCA (NewReno) in current and more importantly, future networks with low latency scenarios.

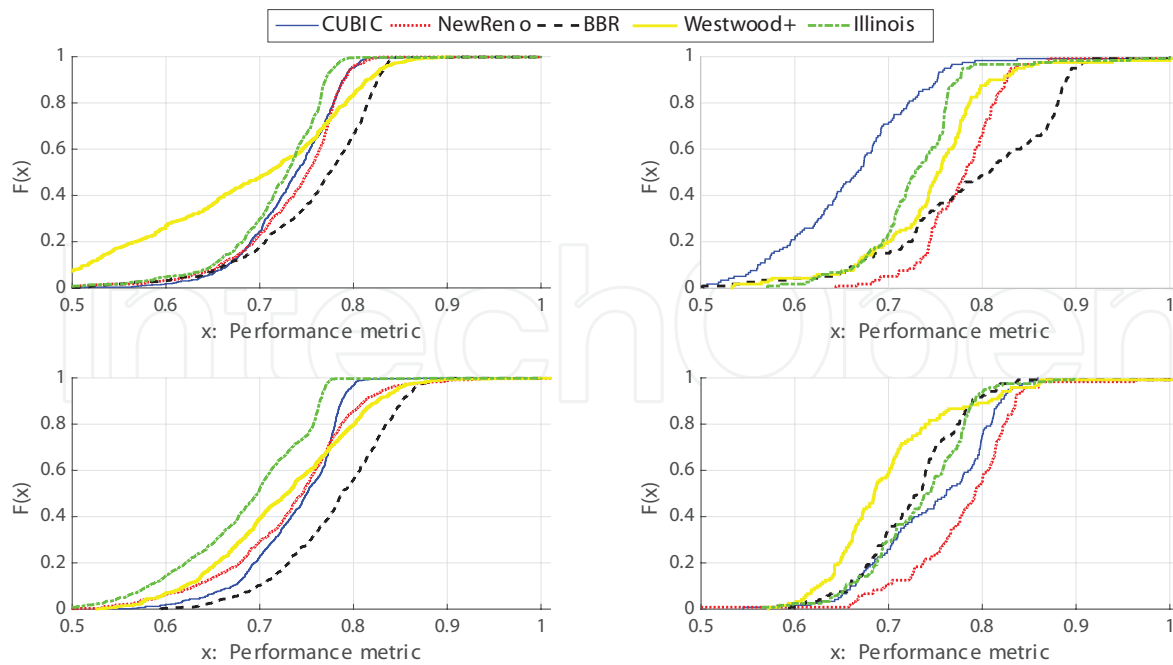
Once we have evaluated and selected based on goodput the best candidates for low latency mobility scenarios, it is crucial to evaluate the performance of the TCP flavors from the point of view of the performance metric. **Figure 6** depicts the ECDF results of the performance metric in all scenarios.

Taking into account that the current evaluation considers three different aspects of the overall performance, the results in **Figure 6** widely differ from the outcome merely based on goodput. There is a clear pattern that confirms BBR as the best candidate in forward movement scenarios. Previous evaluation has shown that based on goodput in all scenarios, BBR is within the group of best performers. Its capacity to reduce the delay close to the baseline delay in a movement pattern that strongly suffers due to the contrary, makes BBR the best candidate considering the performance metric.

In backward movement pattern, there are two clear candidates, one per each speed and fading combination:

Context	Forward movement		Backward movement	
	Mean	Mean CI	Mean	Mean CI
<b>60 km/h with EVA60 under low latencies</b>				
CUBIC	28.18	8.84	28.21	9.02
NewReno	28.68	9.4	24.89	10.36
BBR	28.69	8.29	26.84	8.41
Westwood+	16.57	10.83	20.3	12.27
Illinois	28.19	9.1	23.81	11.78
<b>300 km/h with HST300 under low latencies</b>				
CUBIC	10.33	5.98	11.91	5.4
NewReno	15.79	5.86	12.43	4.97
BBR	14.66	5.38	9.35	5.32
Westwood+	14.66	4.76	8.11	5.58
Illinois	15.82	5.48	11.46	4.92

**Table 2.** Goodput-based evaluation of the selected CCAs in mobility scenarios under low latencies.



**Figure 6.** Comparison of five CCAs performance metric while movement at different speeds under low latencies: forward movement on top; backward movement on bottom line.

1. BBR is able to better handle the required scalability under variable fading conditions with big available capacities and jumps between consecutive assignment samples. As under 4G latencies, BBR demonstrates great scalability. In addition, under low latency circumstances BBR is capable of better managing the induced delay, achieving a considerably better performance than the rest of the TCP implementations in the commented network use-case.
2. The combination of low latency together with low available bandwidths makes the performance playground very suitable for the selection of NewReno, showing great performance both in achieved goodput and based on the performance metric.

Overall the results in **Figure 6** show two important facts in comparison with the mobility scenarios under 4G latencies: (1) the performance of BBR shows a consistent good performance in terms of goodput. The main difference is that in low latency scenarios the behavior of the CCA is more controlled and is able to take full advantage of its features, achieving great goodput while the delay is kept close to the baseline latency. This outcome shows that, the shorter response times (RTT) in each packet, the better for the accurate estimation of BBR in mobility scenarios; (2) based on the performance metric, it is clear that the reduction in the end-to-end latency makes Illinois perform worse in comparison with other candidates. While in 4G latencies, Illinois is eligible in many mobility circumstances that take advantage of its aggressiveness, the same scenarios under shorter latencies do not require its aggressiveness. Actually, this feature only increment the injected delay as well as the number of retransmitted packets.

## 6. Conclusion

After evaluating the QoS of selected mobility scenarios, based on goodput and the performance metric, under both 4G latencies and low latency that mimic the potential delays in 5G networks, we have been able to select the most appropriate CCA for each situation. The selected CCAs are gathered in **Table 3** with an asterisk for the cases in which more than one CCA could be picked.

The mobility scenarios under 4G latencies require more aggressive TCP solutions to overcome the high variability with an increased BDP in comparison with low latency conditions. With reduced end-to-end latencies that model the proximity service provisioning in 5G, the most suitable CCAs are a compendium of classic NewReno for low available capacity circumstances that move in-cell, BBR as the most balanced CCA that allows both high bandwidth achievement and low delay and retransmissions, CUBIC when scalability is required in presence of big changes between consecutive samples of assigned radio capacity and equality in goodput-based outcome in out-cell movement. In this sense, in our pseudo-5G mobility scenarios, it is important to highlight the usability of even weak CCAs such as NewReno in future low latency deployments. The explanation of the selected TCP candidates and their implications in the mobility scenarios is followed CCA-wise.

Illinois has shown an aggressiveness that is suitable for application with goodput requirements in both movement patterns in 4G latency scenarios, but only partially appropriate for forward movement under low latency conditions. The reduction in the end-to-end delay increases the responsiveness of TCP in general and Illinois in particular. This effect makes Illinois suffer greater mean delays and retransmitted packets due to the aggressive injection and ramp-up, leading to deteriorate results (i.e. achieving poor results in backward movement).

BBR has demonstrated good performance in terms of goodput under 4G latencies. However, it also induced a long delay and provokes a great number of transmitted packets, demonstrating an unbalanced behavior in terms of the performance metric. In low latency scenarios instead, BBR has a more controlled and balanced behavior, achieving a sufficiently good performance in goodput but also preserving a low delay.

Context	Low latency		4G latencies	
	Goodput	Metric	Goodput	Metric
<b>60 km/h with EVA60</b>				
Forward	*	BBR	Illinois	*
Backward	CUBIC	BBR	BBR	CUBIC
<b>300 km/h with HST300</b>				
Forward	*	BBR	CUBIC	CUBIC
Backward	NewReno	NewReno	BBR	CUBIC

**Table 3.** Wrap-up of most appropriate CCAs for each mobility scenario and QoS metric.

Westwood+ has demonstrated a poor performance in majority of the scenarios due to its aggressive back-off policy. However, in scenarios with low achievable capacities (at 300 km/h) Westwood+ is capable of slightly bridging the performance gap with other CCAs.

NewReno, under low latency, provides with a TCP implementation able to achieve close to the maximum available capacity, being especially significant its performance and scalability in backward movement pattern with low achievable rates.

CUBIC has shown a balanced performance with low available bandwidth under 4G latencies. Apart from that, under low latency circumstances, CUBIC is appropriate for forward movements with variable big available bandwidths and scenarios that depend upon scalability.

The future work will progress in three main lines. Firstly, the analysis will be extended with the measurement of similar mobility scenarios in a selection of real-world use-cases. Secondly, the work would re-measure and compare the evolution of BBR in terms of adaptability and management of losses. Finally, we will evaluate pseudo-random mobility scenarios that combine both forward and backward movement as well as static positions.

## Acknowledgements

This work is partially funded by the European Union's Horizon 2020 research and innovation programme under grant agreement No. 644399 (MONROE) through the open call project MARiL and by the Spanish Ministerio de Economía y Competitividad (MINECO) under grant TEC2016-80090-C2-2-R (5RANVIR). The views expressed are solely those of the author(s).

## Author details

Eneko Atxutegi\*, Jose Oscar Fajardo and Fidel Liberal

\*Address all correspondence to: eneko.atxutegi@ehu.eus

University of the Basque Country (UPV/EHU), Bilbao, Spain

## References

- [1] Koksall CE, Jamieson K, Telatar E, Thiran P. Impacts of channel variability on link-level throughput in wireless networks. In: Joint International Conference on Measurement and Modeling of Computer Systems, ser. SIGMETRICS '06/Performance '06; 2006. New York: ACM; 2006. pp. 51-62
- [2] Nguyen B, Banerjee A, Gopalakrishnan V, Kasera S, Lee S, Shaikh A, Van der Merwe Jacobus J. Towards understanding TCP performance on LTE/EPC mobile networks. In: 4th Workshop on All Things Cellular: Operations, Applications & Challenges; 2014. ACM; 2014. pp. 41-46

- [3] European Telecommunications Standards Institute (ETSI). Mobile Edge Computing (MEC); Technical Requirements [Internet]. 2016 [Updated: 2016]. Available from: [http://www.etsi.org/deliver/etsi\\_gs/MEC/001\\_099/002/01.01.01\\_60/gs\\_MEC002v010101p.pdf](http://www.etsi.org/deliver/etsi_gs/MEC/001_099/002/01.01.01_60/gs_MEC002v010101p.pdf) [Accessed: 2017]
- [4] The Chromium Projects. QUIC, a multiplexed stream transport over UDP [Internet]. Available from: <https://www.chromium.org/quic> [Accessed: 2017]
- [5] Grinnemo KJ, Jones T, Fairhurst G, Ros D, Brunstrom A, Hurtig P. Towards a flexible internet transport layer architecture. In: 2016 IEEE International Symposium on Local and Metropolitan Area Networks (LANMAN); 2016;2016:1-7
- [6] Jones T, Fairhurst G, Perkins C. Raising the datagram API to support transport protocol evolution. In: Workshop on Future of Internet Transport (FIT 2017); 2017. 2017
- [7] Hurtig P, Alfredsson S, Bozakov Z, Brunstrom A, Evensen K, Grinnemo K-J, Fosselie Hansen A., Rozensztrauch T. A NEAT approach to mobile communication. In: The ACM SIGCOMM 2017 Workshop on Mobility in the Evolving Internet Architecture (MobiArch 2017); 2017. ACM Press; 2017
- [8] Jiang H, Liu Z, Wang Y, Lee K, Rhee I. Understanding bufferbloat in cellular networks. In: 2012 ACM SIGCOMM Workshop on Cellular Networks: Operations, Challenges, and Future Design; 2012. New York. ACM; 2012. pp. 1-6
- [9] Alfredsson S, Del Giudice G, Garcia J, Brunstrom A, De Cicco L, Mascolo S. Impact of TCP congestion control on bufferbloat in cellular networks. In: World of Wireless, Mobile and Multimedia Networks (WoWMoM); 2013. IEEE; 2013. pp. 1-7
- [10] Garcia J, Alfredsson S, Brunstrom A. A measurement based study of TCP protocol efficiency in cellular networks. In: 2014 12th International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOpt); 2014; 2014. pp. 131-136
- [11] Garcia J, Alfredsson S, Brunstrom A. Examining TCP short flow performance in cellular networks through active and passive measurements. In: 5th Workshop on All Things Cellular: Operations, Applications and Challenges, ser. AllThingsCellular '15; 2015; ACM; 2015. pp. 7-12
- [12] Alfredsson S, Brunstrom A, Sternad M. Cross-layer analysis of TCP performance in a 4G system. In: Software, Telecommunications and Computer Networks (SoftCOM); 2007; IEEE; 2007. pp. 1-6
- [13] Huang J, Qian F, Guo Y, Zhou Y, Xu Q, Mao ZM, Sen S, Spatscheck O. An in-depth study of lte: effect of network protocol and application behavior on performance. In: ACM SIGCOMM 2013 Conference on SIGCOMM; 2013; ACM; 2013. pp. 363-374
- [14] Zhang L, Okamawari T, Fujii T. Performance evaluation of end-to-end communication quality of LTE. In: Vehicular Technology Conference (VTC Spring); 2012; IEEE; 2012. pp. 1-5
- [15] Xu Y, Wang Z, Leong WK, Leong B. An end-to-end measurement study of modern cellular data networks. passive and active measurement. Lecture Notes in Computer Science. 2014;8362:34-45



- [16] Merz R, Wenger D, Scanferla D, Mauron S. Performance of LTE in a high-velocity environment: A measurement study. In: 4th Workshop on All Things Cellular: Operations, Applications, & Challenges, ser. All Things Cellular '14; 2014; New York. ACM; 2014. pp. 47-52
- [17] Li L, Xu K, Wang D, Peng C, Xiao Q, Mijumbi R. A measurement study on TCP behaviors in HSPA+ networks on high-speed rails. In: 2015 IEEE Conference on Computer Communications (INFOCOM); 2015; 2015. pp. 2731-2739
- [18] Pedersen KI, Niparko M, Steiner J, Oszmianski J, Mudolo L, Khosravirad SR. System level analysis of dynamic user-centric scheduling for a flexible 5G design. In: 2016 IEEE Global Communications Conference (GLOBECOM); 2016; 2016. pp. 1-6
- [19] Sarret MG, Berardinelli G, Mahmood NH, Mogensen P. Impact of transport control protocol on full duplex performance in 5G networks. In: 2016 IEEE 83rd Vehicular Technology Conference (VTC Spring); 2016; IEEE; 2016. pp. 1-5
- [20] Henderson T, Floyd S, Gurtov A, Nishida Y. The NewReno Modification to TCP's Fast Recovery Algorithm [Internet]. 2012 [Updated: 2012]. Available from: <https://tools.ietf.org/html/rfc6582> [Accessed: 2017]
- [21] Ha S, Rhee I, Xu L. CUBIC: A new TCP-friendly high-speed TCP variant. SIGOPS Operating Systems Review. 2008;42(5):64-74
- [22] Mascolo S, Casetti C, Gerla M, Sanadidi MY, Wang R. TCP Westwood: Bandwidth Estimation for Enhanced Transport over Wireless Links. In: 7th Annual International Conference on Mobile Computing and Networking; 2001; New York. ACM; 2001. pp. 287-297
- [23] Liu S, Başar T, Srikant R. TCP-Illinois: A loss and delay-based congestion control algorithm for high-speed networks. In: 1st International Conference on Performance Evaluation Methodologies and Tools, ser. Valuetools '06; 2006. New York: ACM; 2006
- [24] Cardwell N et al. BBR: Congestion-based congestion control. Communications of the ACM. 2017;60(2):58-66
- [25] Cardwell N, Cheng Y, Yeganeh S, Jacobson V. BBR Congestion Control [Internet]. 2017 [Updated: 2017]. Available from: <http://www.ietf.org/internet-drafts/draft-cardwell-iccr-g-br-congestion-control-00.txt> [Accessed: 2017]
- [26] fiercewireless. 3G/4G wireless network latency: How did Verizon, AT&T, Sprint and T-Mobile compare in Q3 2015? [Internet]. 2016 [Updated: 2016]. Available from: <http://www.fiercewireless.com/special-report/3-g-4-g-wireless-network-latency-how-did-verizon-at-t-sprint-and-t-mobile-compare-q3> [Accessed: 2017]
- [27] Atxutegi E, Liberal F, Grinnemo K-J, Brunstrom A, Arvidsson Å, Robert R. TCP behaviour in LTE: impact of flow start-up and mobility. In: IFIP Wireless and Mobile Networking Conference - WMNC; 11-13 July 2016; Colmar, France